


Upgrading Your Skills to MCSA Windows Server 2012 R2



Exam Ref

70-417

Exam Ref 70-417: Upgrading Your Skills to Windows Server 2012 R2

J.C. Mackin

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

Copyright © 2014 by J.C. Mackin

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: 2014936139
ISBN: 978-0-7356-8440-9

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 Book Support at mspinput@microsoft.com. Please tell us what you think of this book at <http://www.microsoft.com/learning/booksurvey>.

Microsoft and the trademarks listed at <http://www.microsoft.com/en-us/legal/intellectualproperty/Trademarks/EN-US.aspx> are trademarks of the Microsoft group of companies. All other marks are property of their respective owners.

The example companies, organizations, products, domain names, email addresses, logos, people, places, and events depicted herein are fictitious. No association with any real company, organization, product, domain name, email address, logo, person, place, or event is intended or should be inferred.

This book expresses the author's views and opinions. The information contained in this book is provided without any express, statutory, or implied warranties. Neither the authors, Microsoft Corporation, nor its resellers, or distributors will be held liable for any damages caused or alleged to be caused either directly or indirectly by this book.

Acquisitions Editor: Anne Hamilton

Developmental Editor: Karen Szall

Editorial Production: Box Twelve Communications

Technical Reviewers: Mitch Tulloch and Brian Svidergol

Cover: Twist Creative • Seattle

Contents

| | |
|---|------------|
| Introduction | ix |
| <i>Microsoft certifications</i> | <i>x</i> |
| <i>Acknowledgments</i> | <i>x</i> |
| <i>Errata, updates, & book support</i> | <i>xi</i> |
| <i>We want to hear from you</i> | <i>xi</i> |
| <i>Stay in touch</i> | <i>xi</i> |
| Preparing for the Exam | xii |
| Chapter 1 Install and configure servers | 1 |
| Objective 1.1: Install servers | 1 |
| Minimum hardware requirements | 2 |
| Features on Demand | 2 |
| Objective summary | 5 |
| Objective review | 6 |
| Objective 1.2: Configure servers | 7 |
| Installing roles and features | 7 |
| Converting a server with a GUI to or from Server Core | 20 |
| Minimal Server Interface | 22 |
| NIC teaming | 24 |
| Objective summary | 28 |
| Objective review | 29 |
| Objective 1.3: Configure local storage | 30 |
| Introducing Storage Spaces | 30 |
| Objective summary | 37 |
| Objective review | 38 |
| Answers | 40 |

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

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

www.microsoft.com/learning/booksurvey/

| | | |
|------------------|---|-----------|
| Chapter 2 | Configure server roles and features | 45 |
| | Objective 2.1: Configure servers for remote management. | 45 |
| | Managing multiple servers with Server Manager | 46 |
| | Using Group Policy to enable remote management | 57 |
| | Using Remote Server Administration Tools for Windows 8 and Windows 8.1 | 59 |
| | Objective summary | 60 |
| | Objective review | 61 |
| | Answers. | 64 |
| | | |
| Chapter 3 | Configure Hyper-V | 67 |
| | Objective 3.1: Create and configure virtual machine settings | 67 |
| | Hyper-V Module in Windows PowerShell | 68 |
| | Generation 1 and generation 2 virtual machines | 69 |
| | Enhanced session mode | 72 |
| | Dynamic Memory | 74 |
| | Resource Metering | 77 |
| | Non-uniform memory access (NUMA) | 79 |
| | RemoteFX | 79 |
| | Objective summary | 80 |
| | Objective review | 81 |
| | Objective 3.2: Create and configure virtual machine storage. | 82 |
| | New VHDX disk format | 82 |
| | Virtual Fibre Channel adapter | 86 |
| | Storage Quality of Service (QoS) | 88 |
| | Objective summary | 89 |
| | Objective review | 89 |
| | Objective 3.3: Create and configure virtual networks | 90 |
| | Virtual switch extensions | 91 |
| | Network isolation | 93 |
| | Single-root I/O virtualization (SR-IOV) | 95 |
| | Bandwidth management | 97 |
| | Advanced features for virtual network adapters | 99 |
| | Objective summary | 100 |
| | Objective review | 100 |
| | Answers. | 103 |

| | | |
|------------------|--|------------|
| Chapter 4 | Install and administer Active Directory | 107 |
| | Objective 4.1: Install domain controllers | 107 |
| | Installing domain controllers in the GUI | 108 |
| | Installing domain controllers with Windows PowerShell | 113 |
| | Install from Media (IFM) option without defragmentation | 119 |
| | Installing a domain controller in Windows Azure | 120 |
| | Objective summary | 127 |
| | Objective review | 128 |
| | Answers. | 131 |
| | | |
| Chapter 5 | Deploy, manage, and maintain servers | 133 |
| | Objective 5.1: Monitor servers | 133 |
| | Virtual machine resource pools | 134 |
| | Monitoring servers through Windows PowerShell | 136 |
| | Reviewing older monitoring features | 137 |
| | Objective summary | 139 |
| | Objective review | 139 |
| | Answers. | 141 |
| | | |
| Chapter 6 | Configure network services and access | 143 |
| | Objective 6.1: Configure DirectAccess | 143 |
| | What is DirectAccess? | 144 |
| | Understanding IPv6 and DirectAccess | 144 |
| | Understanding the DirectAccess connection process | 146 |
| | Understanding DirectAccess infrastructure options | 147 |
| | Installing and configuring DirectAccess | 153 |
| | Objective summary | 170 |
| | Objective review | 171 |
| | Answers. | 174 |
| | | |
| Chapter 7 | Configure a network policy server infrastructure | 177 |
| | Objective 7.1: Configure Network Access Protection | 177 |
| | How NAP works | 178 |
| | Configuring NAP | 181 |
| | SHV multi-configuration | 183 |
| | Objective summary | 188 |

| | |
|--|------------|
| Objective review | 188 |
| Answers..... | 191 |
| Chapter 8 Configure and manage Active Directory | 193 |
| Objective 8.1: Configure domain controllers..... | 193 |
| Cloning domain controllers | 193 |
| Objective summary | 200 |
| Objective review | 201 |
| Objective 8.2: Maintain Active Directory..... | 202 |
| Restoring deleted objects in Active Directory | 203 |
| Objective summary | 209 |
| Objective review | 209 |
| Answers..... | 212 |
| Chapter 9 Configure and manage Group Policy | 215 |
| Objective 9.1: Configure Group Policy processing..... | 215 |
| Remote Group Policy update | 216 |
| Windows PowerShell cmdlets for Group Policy | 222 |
| Group Policy caching | 224 |
| Objective summary | 226 |
| Objective review | 227 |
| Answers..... | 230 |
| Chapter 10 Configure and manage high availability | 233 |
| Objective 10.1: Configure failover clustering..... | 234 |
| Cluster storage pools | 234 |
| Cluster shared volumes (CSVs) | 236 |
| Virtual hard disk sharing for guest clusters in Windows Server 2012 R2 | 239 |
| Dynamic quorum | 240 |
| Dynamic witness in Windows Server 2012 R2 | 241 |
| Node drain | 241 |
| Cluster-aware updating (CAU) | 242 |
| Active Directory-Detached Clusters in Windows Server 2012 R2 | 246 |
| Configuring Cluster Properties in Windows PowerShell | 248 |
| Objective summary | 250 |
| Objective review | 250 |

| | |
|--|-----|
| Objective 10.2: Manage failover clustering roles | 251 |
| Creating a Scale-Out File Server (SoFS) | 251 |
| Assign role startup priority | 253 |
| Virtual machine application monitoring | 254 |
| Objective summary | 259 |
| Objective review | 259 |
| Objective 10.3: Manage virtual machine (VM) movement | 260 |
| Live migration | 261 |
| Storage migration | 274 |
| VM network health protection in Windows Server 2012 R2 | 276 |
| Objective summary | 278 |
| Objective review | 278 |
| Answers | 281 |

Chapter 11 Configure file and storage solutions 285

| | |
|--|-----|
| Objective 11.1: Implement Dynamic Access Control | 285 |
| Introduction to Dynamic Access Control | 286 |
| Configuring claims-based authentication | 287 |
| Configuring file classification | 291 |
| Configuring access policies | 302 |
| Objective summary | 307 |
| Objective review | 308 |
| Answers | 311 |

Chapter 12 Implement business continuity and disaster recovery 313

| | |
|---|-----|
| Objective 12.1: Configure and manage backups. | 313 |
| Certificate requirements for Windows Azure Backup | 314 |
| Performing Windows Azure Backups in Windows PowerShell | 324 |
| Objective summary | 326 |
| Objective review | 327 |
| Objective 12.2: Configure site-level fault tolerance | 329 |
| Configuring Hyper-V physical host servers | 330 |
| Configuring VMs | 333 |
| Performing Hyper-V Replica failover | 342 |
| Extending replication to a third site in Windows Server 2012 R2 | 347 |
| Using Hyper-V Replica in a failover cluster | 348 |

| | |
|---|------------|
| Objective summary | 351 |
| Objective review | 352 |
| Answers..... | 355 |
| Chapter 13 Configure network services | 359 |
| Objective 13.1: Deploy and manage IPAM..... | 359 |
| What is IPAM? | 360 |
| Installing and configuring IPAM | 361 |
| Managing address space | 372 |
| Role-based access control for IPAM in Windows Server 2012 R2 | 381 |
| Objective summary | 383 |
| Objective review | 383 |
| Answers..... | 386 |
| Chapter 14 Configure identity and access solutions | 389 |
| Objective 14.1: Implement Active Directory Federation Services (AD FS)..... | 389 |
| AD FS scenarios | 390 |
| How AD FS Works | 391 |
| Active Directory Federation Server Configuration Wizard | 392 |
| AD FS management console | 395 |
| Workplace Join | 403 |
| Windows PowerShell cmdlets for AD FS | 406 |
| Objective summary | 407 |
| Objective review | 408 |
| Answers..... | 411 |
| <i>Index</i> | 415 |

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

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

www.microsoft.com/learning/booksurvey/

Introduction

This book is written for IT professionals who want to earn the MCSA: Windows Server 2012 certification by passing the Microsoft exam “Exam 70-417: Upgrading Your Skills to MCSA Windows Server 2012.” Note that in January 2014, this exam was updated to cover the recent technology updates in Windows Server 2012 R2.

Exam 70-417 serves as a path to the Windows Server 2012 MCSA for those who have already earned the Windows Server 2008 certification that is named “MCITP: Server Administrator” and “MCSA: Windows Server 2008.” The book is therefore written specifically for IT professionals who have already earned this Windows Server 2008 certification and maintain the associated level of expertise in Windows Server 2008 or Windows Server 2008 R2.

Exam 70-417 also serves as an upgrade path to the Windows Server 2012 MCSA from certifications other than the Windows Server 2008 MCSA. These other certifications include MCITP: Virtualization Administrator, MCITP: Enterprise Messaging Administrator, MCITP: Lync Server Administrator, MCITP: SharePoint Administrator, and MCITP: Enterprise Desktop Administrator certifications. However, the assumed knowledge for readers of this book is only MCSA-level expertise in Windows Server 2008 or Windows Server 2008 R2.

One of the first things you need to understand about the 70-417 exam is that it is a condensed version of three other exams: Exam 70-410, Exam 70-411, and Exam 70-412. This set of three exams allows you to earn the Windows Server 2012 MCSA from scratch, without any prior certification. Together, these three exams include 18 domains of broader skills and 62 more specific objectives. Because the exams are intended for individuals who haven’t yet earned Windows Server certification, the exams test new features in Windows Server 2012 as well as older features that haven’t changed since Windows Server 2008 or even earlier.

On the 70-417 exam, only 14 of the original 18 domains and 22 of the original 62 objectives have been adopted from these three source exams. This smaller subset of material corresponds generally to the new features in Windows Server 2012. Approximately 75 percent of the questions on the 70-417 exam will assess your knowledge of new Windows Server 2012 features in some way. Approximately 25 percent of the questions will be “review” questions about features that have not changed since Windows Server 2008—questions you could have seen when you earned your existing certification. *The questions that comprise this 25 percent can be taken from any of the 62 original objectives on exams 70-410, 70-411, or 70-412.*

In order to create a book that is a manageable study tool, we’ve focused on covering the 75 percent of material that is new to Windows Server 2012 and that forms the core of the 70-417 exam. After all, the remaining 25 percent of what’s covered on the exam draws upon the knowledge you already have already demonstrated when you earned your Windows Server 2008 certification. However, it’s possible you will need to review some of these older topics,

so we've provided guidance throughout the book to help you identify any topics that might require further review.

This book covers every exam objective, but it does not cover every exam question. Only the Microsoft exam team has access to the exam questions themselves and Microsoft regularly adds new questions to the exam, making it impossible for us to cover specific questions. You should consider this book a supplement to your relevant real-world experience and other study materials. If you encounter a topic in this book that you do not feel completely comfortable with, use the links you'll find in the book to find more information—and then take the time to research and study the topic. Valuable information is available on MSDN, TechNet, and in blogs and forums.

Microsoft certifications

Microsoft certifications distinguish you by proving your command of a broad set of skills and experience with current Microsoft products and technologies. The exams and corresponding certifications are developed to validate your mastery of critical competencies as you design and develop, or implement and support, solutions with Microsoft products and technologies both on-premise and in the cloud. Certification brings a variety of benefits to the individual and to employers and organizations.

MORE INFO ALL MICROSOFT CERTIFICATIONS

For information about Microsoft certifications, including a full list of available certifications, go to <http://www.microsoft.com/learning/en/us/certification/cert-default.aspx>.

Acknowledgments

I'd like to thank Anne Hamilton and Karen Szall at Microsoft for their long-standing support; Travis Jones, Adnan Ijaz, and Osama Sajid at Microsoft for answering low-level questions about Remote Management; Jeff Riley at Box Twelve Communications for his steady management and flexibility; and Mitch Tulloch and Brian Svidergol for their world-class technical review.

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/ER417R2>

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 is 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>.

Preparing for the Exam

Microsoft certification exams are a great way to build your resume and let the world know about your level of expertise. Certification exams validate your on-the-job experience and product knowledge. While there is no substitution for on-the-job experience, preparation through study and hands-on practice can help you prepare for the exam. We recommend that you round out your exam preparation plan by using a combination of available study materials and courses. For example, you might use the training kit and another study guide for your "at home" preparation, and take a Microsoft Official Curriculum course for the classroom experience. Choose the combination that you think works best for you.

Note that this training kit is based on publically available information about the exam and the author's experience. To safeguard the integrity of the exam, authors do not have access to the live exam.

Install and configure servers

The Install and Configure Servers content area (or “domain”) originates from the 70-410 exam. Unlike that exam, the 70-417 upgrade exam strongly emphasizes *new features* in Windows Server 2012 and Windows Server 2012 R2—like Features on Demand, Server Core-GUI convertibility, and Storage Spaces—that relate to the initial configuration of Windows Server or of server hardware. In fact, it’s important to know that as many as 75 percent of the questions on the 70-417 exam in general relate to new features.

The most sensible strategy for this book is therefore to focus on these new features. These topics make up the bulk of the exam, and you have already demonstrated your understanding of the remaining 25 percent by passing the exams for your Windows Server 2008 certification. Even so, make sure that as part of your preparation plan for 70-417, you take the time to review any unchanged features that you don’t feel confident about.

Objectives in this chapter:

- Objective 1.1: Install servers
- Objective 1.2: Configure servers
- Objective 1.3: Configure local storage

IMPORTANT

Have you read page xii?

It contains valuable information regarding the skills you need to pass the exam.

Objective 1.1: Install servers

“Installing servers” might sound like an easy topic that you don’t need to study for, but this objective is more challenging than you might think. Yes, you should certainly review the hardware requirements for Windows Server 2012 and Windows Server 2012 R2, but you also need to understand a new concept that you’re likely to see on the 70-417 exam: Features on Demand.

This section covers the following topics:

- Minimum hardware requirements for Windows Server 2012 and Windows Server 2012 R2
- Features on Demand

Minimum hardware requirements

You already know you won't see questions on any Microsoft exam that would ask you, for example, "What are the processor requirements for Windows?" But sometimes hardware requirements sneak into exam questions indirectly. For example, you might see a scenario in which a new feature that is available only in Windows Server 2012 is needed and the existing server hardware (based on, for example, an x86 processor) requires an upgrade to support the new operating system. Fortunately, this time around, the hardware requirements are easy to learn: The minimum hardware requirements for Windows Server 2012 and Windows Server 2012 R2 are, in fact, the same as those for Windows Server 2008 R2. Here's a recap:

- Processor: 1.4 GHz 64-bit processor
- RAM: 512 MB (allocate more for the Chinese version)
- Disk space: 32 GB

Don't miss the obvious here: Windows Server 2012 and Windows Server 2012 R2 require a 64-bit processor, unlike Windows Server 2008 (but like Windows Server 2008 R2). This fact could easily form the basis for a test question. If a question states you need to upgrade to Windows Server 2012 on an existing server, make sure the server has a 64-bit processor. If not, you need to replace the hardware. If the hardware is compatible, you can perform an in-place upgrade (as opposed to a fresh installation) from Windows Server 2008 SP2.

Features on Demand

A copy of the binary files for all features and roles that are installed during Windows Setup is stored in a directory called the *side-by-side store*, located in `Windows\WinSxS`. Keeping a copy of the feature files available on disk in this way allows you to add a role or enable a feature after Windows Server installation without needing to access Windows Server media. The side-by-side store first appeared in Windows Server 2008 and Windows Server 2008 R2, and in those previous versions of Windows Server, the features files remained on disk for the life of the operating system. The disadvantage of this approach was that these files took up space on the disk even if you never wanted to install its associated feature or role. In addition, you weren't able to reduce the size of the installation image, as you might want to do when creating custom installation media for your organization.

Beginning with Windows Server 2012, you can minimize the footprint of your installation by deleting the files for features you're not using from the side-by-side store. This ability to delete feature files is called *Features on Demand*. To later reinstall a role or feature for which files have been deleted, you need to have access to the source files.

To completely remove all files for a role or feature from disk, use the `Uninstall-WindowsFeature` cmdlet of Windows PowerShell and specify the name of the feature using the `-Remove` parameter. For example, to delete the DHCP server binaries from server storage, run the following Windows PowerShell command:

```
Uninstall-WindowsFeature DHCP -Remove
```

To remove from disk all the feature files that are not currently installed on the local server, run the following Windows PowerShell command:

```
Get-WindowsFeature | Where-Object -FilterScript { $_.Installed -Eq $FALSE } | Uninstall-WindowsFeature -Remove
```

Note that if you want to remove the feature files for a single role or feature, you can perform that same function at an elevated command prompt by using the Dism utility with the /Remove option:

```
Dism /Online /Disable-Feature /FeatureName:DHCPServer /Remove
```

(The Dism utility is covered in more detail later in this chapter.)

NOTE Roles and features in Windows PowerShell are referred to by their command names specific to Windows PowerShell, not by their common display names. The 70-417 exam covers Windows PowerShell more thoroughly than its counterpart exams did in Windows Server 2008 and Windows Server 2008 R2, so it's a good idea to familiarize yourself with many of these Windows PowerShell command names. You can do this by typing **Get-WindowsFeature** at a Windows PowerShell prompt and reviewing the output.

Figure 1-1 shows the result after this procedure when you run the Get-WindowsFeature cmdlet. The DHCP Server install state is described as Removed.

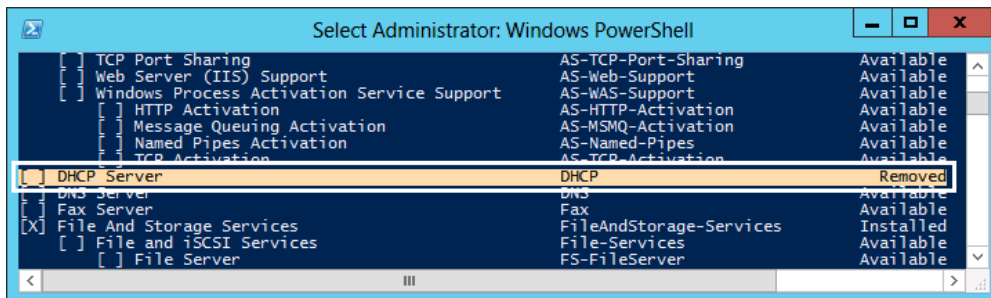


FIGURE 1-1 Removing feature files

You can reinstall these feature files at any point. To install a role or feature for which the binaries have been deleted, you can use the Install-WindowsFeature cmdlet in Windows PowerShell with the -Source parameter to specify any of the following:

- A local path to the \sources\sxs directory on the product DVD.

For example, to install the DHCP feature after its feature files have been removed from disk and to specify the product DVD (on the D: drive) as a source of those feature files, type the following at a Windows PowerShell prompt:

```
Install-WindowsFeature DHCP -Source D:\sources\sxs
```


- A path to a local WIM file, for example, on the product DVD.

The path for a WIM file should be in the following format: WIM:[drive letter]:\sources\install.wim:[image index], for example, WIM:e:\sources\install.wim:4.)

- A UNC path to a network share that contains the WinSxS folder for the appropriate version of Windows Server 2012 or Windows Server 2012 R2.
- A UNC path to a WIM file on a network share, using the "WIM:" prefix before the path.

If you do not specify a –Source option, Windows will attempt to access the files by performing the following tasks in order:

1. Searching in a location that has been specified by users of the Add Roles And Features Wizard or Deployment Image Servicing and Management (DISM) installation commands.
2. Evaluating the configuration of the following Group Policy setting: Computer Configuration\Policies\Administrative Templates\System\Specify Settings For Optional Component Installation And Component Repair.
3. Searching Windows Update. (This can be a lengthy process for some features.)

Alternatively, you can reinstall the feature using Server Manager. When you get to the final page of the Add Roles And Features Wizard, choose the option to specify an alternate source path, as shown in Figure 1-2. Then, provide a path to source files when prompted.

The source path or file share must grant Read permissions either to the Everyone group (not recommended for security reasons) or to the computer account of the destination server; granting user account access is not sufficient.



EXAM TIP

Remember how to reinstall features whose feature files have been removed from disk.

MORE INFO For more information on Features on Demand, visit <http://technet.microsoft.com/en-us/library/jj127275.aspx>.

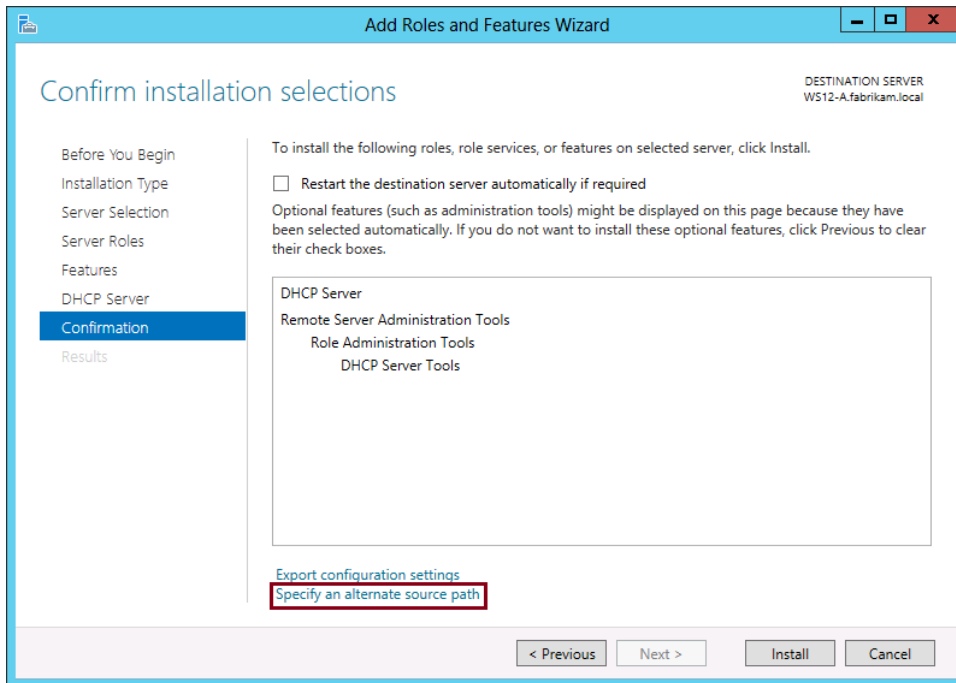


FIGURE 1-2 Reinstalling feature files that have been removed

Objective summary

- The minimum hardware requirements for Windows Server 2012 and Windows Server 2012 R2 are the same as those for Windows Server 2008 R2: a 1.4 GHz 64-bit processor, 512 MB of RAM, and 32 GB of storage.
- The `Uninstall-WindowsFeature` cmdlet uninstalls and removes specified roles, role services, and features from a computer that is running Windows Server 2012 or Windows Server 2012 R2, or an offline VHD that has Windows Server 2012 or Windows Server 2012 R2 installed on it.
- You can reduce the storage footprint of your Windows Server 2012 or Windows Server 2012 R2 installation by removing from disk the files for unused roles or features. To remove feature files, use the following Windows PowerShell command:

```
Uninstall-WindowsFeature feature name -Remove
```

- To reinstall a feature for which files have been removed from the local disk, use the following Windows PowerShell command:

`Install-WindowsFeature feature name [-Source path to the \source\sxs directory on the product DVD, to a WIM file, or to a share containing an WinSxS folder from an appropriate installation of Windows Server 2012 or Windows Server 2012 R2]`

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the “Answers” section at the end of the chapter.

1. You work for a large company named Contoso.com. A server in the finance department named Server1 is running Windows Server 2008. The server includes a 2.0 GHz 32-bit CPU and 4 GB of RAM.

Management has issued the requirement that every server should be reduced to a minimal footprint, and the files of all uninstalled features should be completely removed from server storage. What should you do? (Choose all that apply.)

- A. Keep the existing server and install Windows Server 2012 R2.
- B. Replace the existing server and install Windows Server 2012 R2.
- C. Run the following command at an elevated Windows PowerShell prompt:

```
Get-WindowsFeature | Where-Object -FilterScript { $_.Installed -Eq $FALSE } |  
Uninstall-WindowsFeature -Remove
```

- D. Run the following command at an elevated Windows PowerShell prompt:

```
Get-WindowsFeature | Where-Object -FilterScript { $_.Installed -Eq $TRUE } |  
Uninstall-WindowsFeature
```

2. You want to reduce the amount of space taken up by Windows Server 2012 R2 for a Server Message Block (SMB) file server named Server1. Server1 is a member of the Contoso.com domain but doesn't perform any function beyond that of an SMB file server. Which of the following commands, entered at a Windows PowerShell prompt, are acceptable methods to reduce the size of the Windows Server 2012 R2 installation on Server1? (Choose all that apply.)
 - A. `Uninstall-WindowsFeature Web-Server -Remove`
 - B. `Dism /online /disable-feature /featurename:iis-webserverrole /remove`
 - C. `Uninstall-WindowsFeature FS-FileServer -Remove`
 - D. `Dism /online /disable-feature /featurename:File-Services /remove`

3. Web1 is a web server on your network connected to the Internet. You have used the Uninstall-WindowsFeature cmdlet in Windows PowerShell to remove from disk the feature files for Active Directory Domain Services on Web1. Which of the following commands provides a valid method to reinstall these feature files, if you insert the product media into the D: drive? (Choose all that apply.)
- A. Install-WindowsFeature –Source WIM:D:\sources\install.wim:1
 - B. Install-WindowsFeature –Source D:\sources\install.wim:1
 - C. Install-WindowsFeature –Source WIM:D:\sources\install.wim
 - D. Install-WindowsFeature –Source D:\sources\sxs

Objective 1.2: Configure servers

Within this objective, there are three major feature changes that were introduced in Windows Server 2012 and one major feature introduced in Windows Server 2012 R2.

First among the improvements introduced in Windows Server 2012 is the process of adding or removing server roles and features: You can now perform these functions either locally or remotely, and through either the GUI or by using Windows PowerShell. The second new feature is the ability to switch between a Server Core installation of Windows Server 2012 and an installation of Windows Server 2012 that includes a graphical user interface. Third, Windows Server 2012 introduced network interface card (NIC) teaming, a fault resiliency feature that you are likely to configure soon after installation.

The new configuration option introduced in Windows Server 2012 R2 is *Windows PowerShell Desired State Configuration (DSC)*, which is a new code framework for ensuring servers are properly configured.

This section covers the following topics:

- Installing roles and features
- Configuring online and offline images with the DISM.exe utility
- Converting between Server Core and full graphical user interface (GUI)
- Minimal Server Interface
- NIC teaming

Installing roles and features

You already know you can use Server Manager to add or remove roles or features locally. As we now have seen in the last objective, you can also now use the new Install-WindowsFeature and Uninstall-WindowsFeature cmdlets to achieve these same tasks in Windows PowerShell.



EXAM TIP

Add-WindowsFeature is an alias of the **Install-WindowsFeature** cmdlet and **Remove-WindowsFeature** is an alias of the **Uninstall-WindowsFeature** cmdlet. You can see all of these versions on the 70-417 exam.

Even more interestingly, you can now use either Windows PowerShell or Server Manager to perform these tasks remotely.

Deploying features and roles on remote servers through Windows PowerShell

Beginning in Windows Server 2012, you can deploy roles and features on remote servers. This feature represents an important new functionality that is likely to be tested on the 70-417 exam.

NOTE For the following procedures, it is assumed that the remote computer is configured to allow remote management (this is the default configuration) and that both the source and destination computers are located in the same Active Directory Domain Services domain.

MORE INFO For information on how to manage remote servers from Server Manager in a workgroup environment, see “Add Servers to Server Manager” at <http://technet.microsoft.com/en-us/library/hh831453>.

To install roles and features on a remote server by using Windows PowerShell:

1. Type **Get-WindowsFeature** and then press Enter to view a list of available and installed roles and features on the local server. If the local computer is not a server, run **Get-WindowsFeature -ComputerName <computer_name>**, where *computer_name* represents the name of a remote computer that is running Windows Server 2012 or Windows Server 2012 R2. The results of the cmdlet contain the command names of roles and features that you add to your cmdlet in step 4.

Note that the output of **Get-WindowsFeature** is long and unwieldy if you are looking just to see which features are already installed. To see only the features on the local machine that are already installed, type the following at a Windows PowerShell prompt:

```
Get-WindowsFeature | Where-Object -FilterScript { $_.Installed -Eq $TRUE }
```

Again, you can target a remote server with the `-ComputerName <computer_name>` parameter. For example, to see the list of installed features on a remote server named DC1, type the following:

```
Get-WindowsFeature -ComputerName DC1 | Where-Object -FilterScript { $_.Installed -Eq $TRUE }
```

2. Type **Get-Help Install-WindowsFeature** and then press Enter to view the syntax and accepted parameters for the Install-WindowsFeature cmdlet.
3. Type the following and then press Enter, where *feature_name* represents the command name of a role or feature that you want to install (obtained in step 1) and *computer_name* represents a remote computer on which you want to install roles and features. Separate multiple values for *feature_name* by using commas. The `-Restart` parameter automatically restarts the destination server if required by the role or feature installation.

```
Install-WindowsFeature -Name <feature_name> -ComputerName <computer_name> -Restart
```

Figure 1-3 shows an example of using this cmdlet to install a feature (NFS-Client) on a remote server (WS12R2-B).

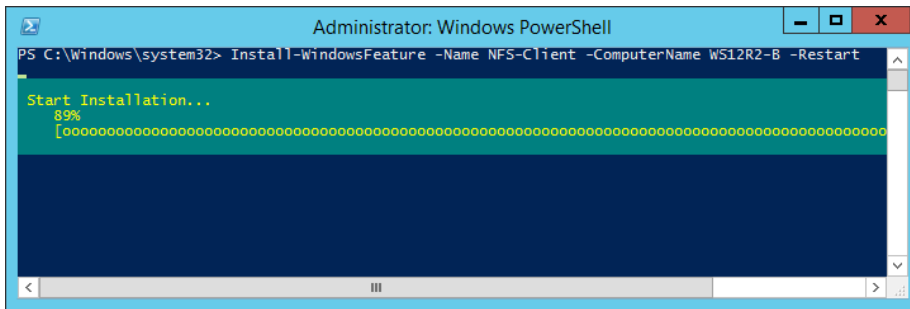


FIGURE 1-3 Installing a feature on a remote server

Deploying features and roles on remote servers with Server Manager

If you prefer to use Server Manager to deploy roles and features to a remote server, you first need to add the remote server to the Server Manager server pool.

To add a remote server in Server Manager, follow these steps:

1. From the Manage menu, select Add Servers, as shown in Figure 1-4.

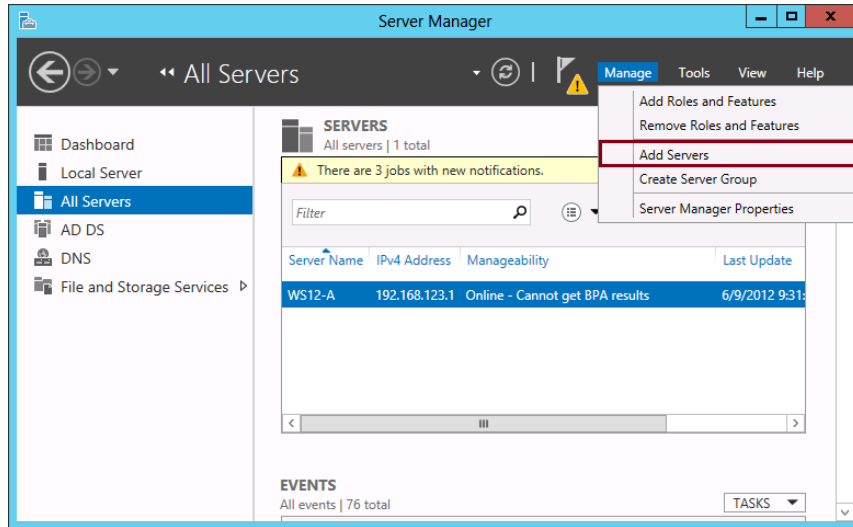


FIGURE 1-4 Adding a remote server to manage in Server Manager

2. Do one of the following:
 - On the Active Directory tab, select servers that are in the current domain. Press Ctrl while selecting multiple servers. Click the right-arrow button to move selected servers to the Selected list.
 - On the DNS tab, type the first few characters of a computer name or IP address and then press Enter or click Search. Select servers that you want to add and then click the right-arrow button.
 - On the Import tab, browse for a text file that contains the DNS names or IP addresses of computers that you want to add, one name or IP address per line.
3. When you are finished adding servers, click OK.

The new server will appear in Server Manager when you select All Servers in the navigation pane, as shown in Figure 1-5.

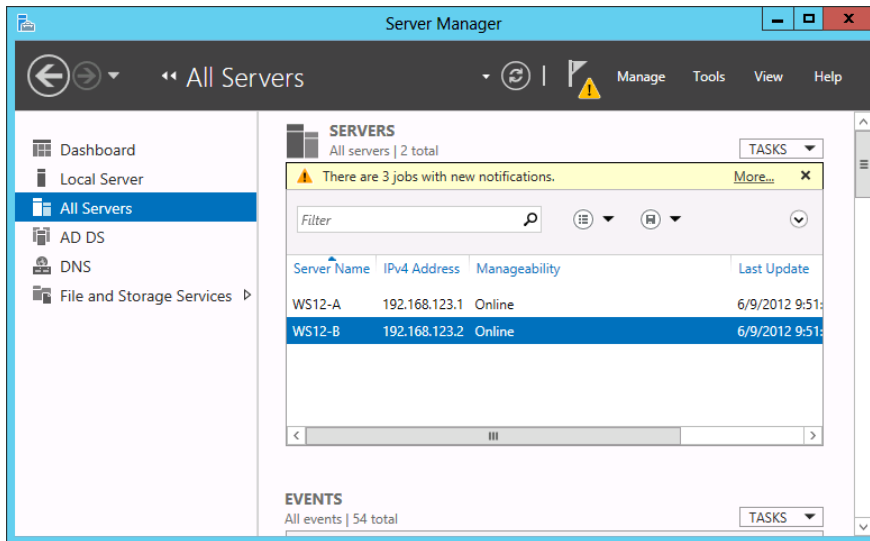


FIGURE 1-5 The remote server WS12-B has been added in Server Manager

After you have added the remote server to the server pool, you can deploy features to it as you would to the local server.

To install roles and features on a remote server by using Server Manager:

1. From the Manage menu of Server Manager, select Add Roles And Features.
2. On the Before You Begin page, click Next.
3. On the Select Installation Type page, select Role-Based Or Feature-Based Installation to install all parts of roles or features on a single server, or Remote Desktop Services Installation to install either a virtual machine–based desktop infrastructure or a session-based desktop infrastructure for Remote Desktop Services. The Remote Desktop Services Installation option distributes logical parts of the Remote Desktop Services role across different servers as needed by administrators. Click Next.
4. On the Select Destination Server page, select a server from the server pool. After you have selected the destination server, click Next.
5. Select roles, select role services for the role if applicable, and then click Next to select features.
6. On the Confirm Installation Selections page, review your role, feature, and server selections. If you are ready to install, click Install.

You can also export your selections to an XML-based configuration file that you can use for unattended feature installations with Windows PowerShell. To export the configuration you specified in this Add Roles And Features Wizard session, click Export Configuration Settings, as shown in Figure 1-6, and then save the XML file to a convenient location.

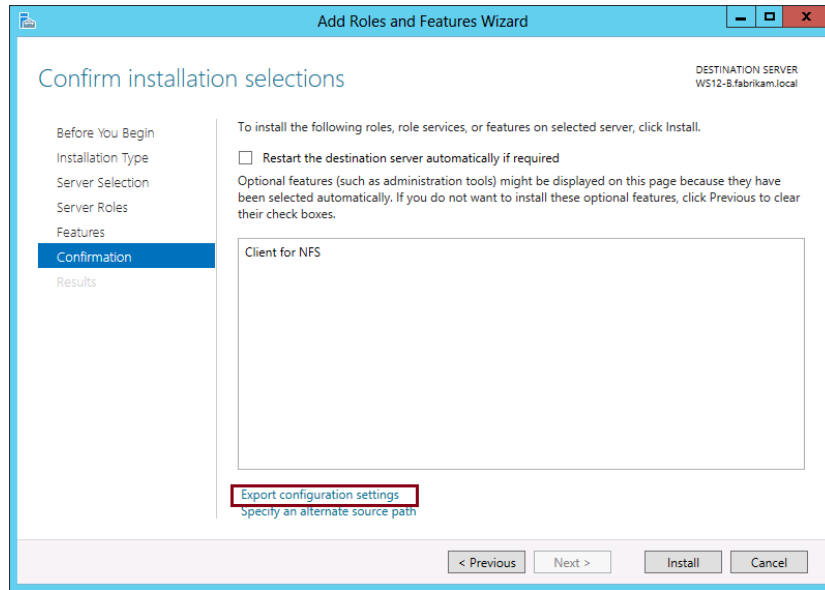


FIGURE 1-6 Exporting an XML configuration file for use with Windows PowerShell



EXAM TIP

You can use the Add Roles And Features Wizard to create server configuration files that can later be used to perform unattended feature installations with Windows PowerShell.

7. After you click Install, the Installation Progress page displays installation progress, results, and messages such as warnings, failures, or postinstallation configuration steps that are required for the roles or features that you installed. In Windows Server 2012 and Windows Server 2012 R2, you can close the Add Roles And Features Wizard while installation is still in progress, and view installation results or other messages in the Notifications area at the top of the Server Manager console. Click the Notifications flag icon to see more details about installations or other tasks that you are performing in Server Manager.

Windows PowerShell Desired State Configuration (DSC)

DSC is an installable feature and server configuration framework that is new to Windows Server 2012 R2. The primary advantage of this new feature is that it lets you ensure that software components of selected servers are present and properly configured.

DSC works through Windows PowerShell. To install the feature, you can use the Add Roles And Features Wizard to add the component of the Windows PowerShell feature named "Windows PowerShell Desired State Configuration." To add the feature from an elevated Windows PowerShell prompt, type **Install-WindowsFeature DSC-Service**.

What's new and different about DSC is that it's a *declarative* management system. This means that, instead of performing particular tasks on one or more target servers, you declare through DSC a desired end-state for those target servers. If the servers already meet the desired configuration, no errors are generated. If the servers do not meet the desired configuration, they individually and automatically make the changes needed to bring themselves into compliance with the desired end-state.

You can use DSC to manage target servers running only Windows Server 2012 R2 by default. However, if you install Windows Management Framework 4.0 on servers running Windows Server 2008 R2 and Windows Server 2012, you can manage these servers through DSC as well.

There are three steps to using DSC after you install it on the management and target servers:

1. Define the desired configuration by creating a new Configuration function in the Windows PowerShell ISE tool. You give this function a name of your choice and use it to invoke DSC resources such as WindowsFeature, File, Group, or Service. You use the syntax of each DSC resource to define the specific components you want to ensure are configured properly on the target server or servers.
2. Enact the configuration by calling (typing) your new function at a Windows PowerShell prompt. This step turns your desired configuration into a Managed Object Format (MOF) file that is usable with the Start-DscConfiguration cmdlet. (Start-DscConfiguration is the cmdlet you use to apply configurations to servers through DSC.) The MOF file is created beneath the current directory in a new subdirectory with the name of your new function.
3. Run the following Windows PowerShell command:

```
Start-DscConfiguration -Wait -Verbose -Path .\FunctionName
```

where FunctionName is the name of your Configuration function. Note that the `-Wait` and `-Verbose` parameters are merely recommended and not required. Also note that if the target server is remote, you first need need to copy the MOF file to the remote server before you can run Start-DscConfiguration.

The following is an example Configuration function named EnsureIIS. Its purpose is to ensure that the IIS role is installed on the server named Web1. Read the code along with the comments to get an idea of the basic syntax you might need to understand for the 70-417 exam.

```
Configuration EnsureIIS
{
# First you have to specify one or more servers with the keyword "Node."
# A node is a server on which the configuration will take place.
  Node "Web1"
  {
# After specifying the node, you have to specify one or more DSC resources.
# WindowsFeature is one of the many built-in DSC resources.
# Each DSC resource has specific syntax for configuring components.
    WindowsFeature IIS
    {
        Ensure = "Present" # You would set this to "Absent" to uninstall the role.
        Name = "Web-Server"
    }
  }
}
```

After you run this code, you would type EnsureIIS at the same prompt. This step creates an MOF file in a new directory named EnsureIIS beneath your current directory.

Finally, you would run the following command at the same prompt (or another prompt pointed at the same location in the file structure):

```
Start-DscConfiguration -Wait -Verbose -Path .\EnsureIIS
```

This final step would install IIS if it were not already installed.

What do you need to know about DSC for the 70-417 exam? First of all, you need to understand the scenarios in which DSC is useful. The key concept here is that DSC provides *declarative management*. This means that with DSC you can enforce a desired configuration on remote servers without generating an error if one or more of those target servers is already in compliance. So, if you see a question scenario that calls for a way to manage servers by specifying a target state for those servers, you should suspect that DSC is at least part of the answer.

Next, you should also remember that DSC includes resources for supporting the following configuration components:

- Enabling or disabling server roles and features
- Managing registry settings
- Managing files and folders
- Starting, stopping, and managing processes and services
- Managing local user and group accounts
- Deploying new software packages
- Managing environment variables
- Running Windows PowerShell scripts

Third, you should understand the basic steps required to configure a server through DSC. For example, you might be presented with a scenario in which you have already created a Configuration function called MyConfig, and you will be asked which step to take next. (Answer: Type MyConfig. Don't use Start-DscConfiguration yet).

Finally, it is possible that you will need to understand some very simple DSC code, especially the basic syntax for DSC resources. Remember above all that the line Ensure = "Present" in a resource block of code ensures that a feature is installed and Ensure = "Absent" ensures that a feature is *not* installed.

DSC RESOURCES

Aside from the WindowsFeature DSC resource in the code example above, you should also learn to understand basic syntax for the File, Group, and Service resources. The following is an example of using the File DSC resource, which could appear as a resource block in a Configuration function beneath "Node":

```
File DirectoryCopy
{
    Ensure = "Present" # You can also set Ensure to "Absent."
    Type = "Directory" # Default is "File"
    Recurse = $true # Ensure presence of subdirectories, too.
    SourcePath = "C:\Users\Public\Documents\DSCDemo\DemoSource"
    DestinationPath = "C:\Users\Public\Documents\DSCDemo\DemoDestination"
}
```

The following is an example of the Group resource:

```
Group GroupExample
{
    # This will remove TestGroup, if present.
    # To create a new group, set Ensure to "Present."
    Ensure = "Absent"
    Name = "TestGroup"
}
```

And the following is an example of the Service resource:

```
Service ServiceExample
{
    Name = "TermService"
    StartupType = "Manual"
}
```

MORE INFO You can view the syntax for all the built-in DSC resources at <http://technet.microsoft.com/en-us/library/dn249921.aspx>.

PARAMETRIZED FUNCTIONS

You can write your Configuration function to accept parameters, which lets you apply your configuration to any server. For example, before the line beginning with "Node" in the function EnsureIIS, you could include the following line:

```
param ($MyTargetNodeName)
```

Then instead of specifying a particular server ("Web1") in the line beginning with "Node", you would include this line:

```
Node $MyTargetNodeName
```

When you called the function in step 2, you would specify the target computer this way:

```
EnsureIIS -MyTargetNodeName Web1
```

The advantage of this code is that it would also let you target another server, such as Web2. So, using a parameter, you could quickly generate an MOF file to configure any server in the way specified in the Configuration function.

CONFIGURING FEATURES IN A SEQUENCE WITH THE REQUIRES KEYWORD

Sometimes you have to configure features on a server in a particular order. For example, you have to install IIS before you can configure its default website to stop. In this case, you need to use the Requires keyword in a resource block. The following is an example of a website resource defined in a code block that uses the Requires keyword in just such a way:

```
# Stop the default web site.
Website DefaultSite
{
    Ensure = "Present"
    Name = "Default Web Site"
    State = "Stopped"
    PhysicalPath = "C:\inetpub\wwwroot"
    Requires = "[WindowsFeature]IIS"
}
```



EXAM TIP

Don't be too intimidated by DSC. The most important thing to remember about it for the exam is *what it is used for*. It's quite possible that if you see a question about DSC on the exam, the question will test only your understanding of the purpose of the feature. That said, to be well-prepared for the exam and to gain a deeper understanding of the feature, you really need to try it out in a test environment. You will find that it isn't as hard to understand as it looks.

Windows PowerShell Web Access (PSWA)

Windows PowerShell Web Access (PSWA) is a feature that first appeared in Windows Server 2012. PSWA allows administrators to connect to a Windows PowerShell prompt on a remote server through a web site hosted on that remote server. For security, users are blocked from accessing the PSWA site unless they are specifically given access to it through authorization

rules. These authorization rules also specify the computers from which authorized users are allowed access to PSWA.

To configure PSWA, you first need to install the feature, which is a component of the Windows PowerShell feature. (Installing the PSWA component automatically installs Internet Information Services [IIS] if it isn't already installed.) Then, you need to take the following steps:

1. At a Windows PowerShell prompt, type the command **Install-PswaWebApplication**. PSWA is a web application whose name on the exam may be referred to as the "Windows PowerShell Web Access gateway." Running the `Install-PswaWebApplication` cmdlet will install and configure the PSWA gateway on Default Web Site in IIS. The URL for the gateway is `https://servername/pswa`, where *servername* represents the name of the server.
2. Bind a valid SSL certificate to port 443 on the Default Web Site in IIS. (For testing purposes, you can avoid this step if you use the `-UseTestCertificate` parameter with the `Install-PswaWebApplication` cmdlet. The `-UseTestCertificate` parameter generates a self-signed SSL certificate and binds it to Default Web Site.)
3. At a Windows PowerShell prompt, use the `Add-PswaAuthorizationRule` cmdlet to configure authorization rules for specific users and computers. Use the `-UserName` parameter to specify authorized users, the `-ComputerName` parameter to specify authorized computers from which connections are allowed, and the `-ConfigurationName` parameter specifies the Windows PowerShell session configuration file on the local server for which these authorization rules apply.

Here is an example authorization rule, which authorizes `Contoso\JohnM` and `Contoso\NancyB` to connect from `Srv2.contoso.com` when the default Windows PowerShell session configuration file is active on the local server:

```
Add-PswaAuthorizationRule -UserName Contoso\JohnM, Contoso\NancyB -ComputerName Srv2.contoso.com -ConfigurationName Microsoft.PowerShell
```



EXAM TIP

Expect to see a question about how to configure PSWA on the exam. Remember especially the significance of the `Install-WebApplication` and `Add-PswaAuthorizationRule` cmdlets.

MORE INFO For more information on configuring PSWA, see "Install and Use Windows PowerShell Web Access" at <http://technet.microsoft.com/en-us/library/hh831611.aspx>.

Deployment Image Servicing and Management

If you received your MCSA certification for Windows Server 2008 before the release of R2, you might have missed hearing about the *Deployment Image Servicing and Management (DISM)* utility. DISM is an image configuration tool that first appeared in Windows 7 and

Windows Server 2008 R2, and its functionality was expanded in the first release of Windows Server 2012. DISM replaced several deployment tools that were used in Windows Server 2008 and Windows Vista, including PEimg, Intlcfg, ImageX, and Package Manager.

In Windows 8 and Windows Server 2012 and later, DISM helps you service Windows Imaging (WIM), VHD, and the new VHDX file types.

You can use DISM with .wim files to do the following:

- Capture and apply Windows images
- Append and delete images in a .wim file
- Split .wim files into several smaller files

You can use DISM with .wim, .vhd, or .vhdx files to do the following:

- Add, remove, and enumerate packages
- Add, remove, and enumerate drivers
- Enable or disable Windows features
- Upgrade a Windows image to a different edition
- Prepare a Windows PE image

An important thing to know about DISM is that you can use it to service online images as well as offline images. *Servicing the online image is essentially the same as configuring the local running installation of Windows.*

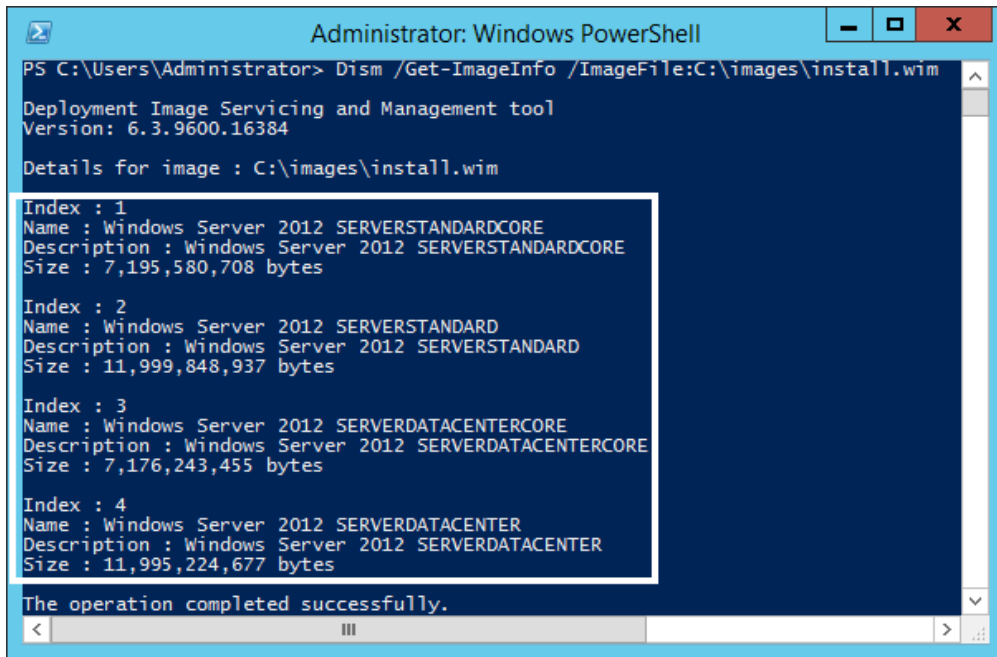
NOTE Beginning in Windows 8 and Windows Server 2012, Windows PowerShell includes a new module for DISM. You can review the cmdlets in this module by typing the command `Get-Command -Module Dism` at a Windows PowerShell prompt. For more information about this module, visit <http://technet.microsoft.com/en-us/library/hh852126>.

ADD FEATURES TO AND REMOVE FEATURES FROM AN OFFLINE IMAGE WITH DISM

Before you can service an offline image, you need to mount the image in the file structure, specifying the image by index or name. In Windows Server 2012 and Windows Server 2012 R2, you can first find the image names and indexes within an image file by using DISM with the `/Get-ImageInfo` switch. For example, to see the listed images within an image file named `Install.wim` that is stored in `C:\images`, type the following:

```
Dism /Get-ImageInfo /ImageFile:C:\images\install.wim
```

The output of this command is shown in Figure 1-7.



```
Administrator: Windows PowerShell
PS C:\Users\Administrator> Dism /Get-ImageInfo /ImageFile:C:\images\install.wim
Deployment Image Servicing and Management tool
Version: 6.3.9600.16384

Details for image : C:\images\install.wim

Index : 1
Name : Windows Server 2012 SERVERSTANDARDCORE
Description : Windows Server 2012 SERVERSTANDARDCORE
Size : 7,195,580,708 bytes

Index : 2
Name : Windows Server 2012 SERVERSTANDARD
Description : Windows Server 2012 SERVERSTANDARD
Size : 11,999,848,937 bytes

Index : 3
Name : Windows Server 2012 SERVERDATACENTERCORE
Description : Windows Server 2012 SERVERDATACENTERCORE
Size : 7,176,243,455 bytes

Index : 4
Name : Windows Server 2012 SERVERDATACENTER
Description : Windows Server 2012 SERVERDATACENTER
Size : 11,995,224,677 bytes

The operation completed successfully.
```

FIGURE 1-7 Obtaining image information from an image file

Once you know the name or index of the desired image, you can mount it in a specified directory. For example, use the following command to mount the image with index 2 in the C:\images\offline directory:

```
Dism /Mount-Image /ImageFile:C:\images\install.wim /index:2 /MountDir:C:\images\offline
```

At this point, you can use the /Get-Features switch if needed to determine the command name of the relevant features or to determine which features are enabled on the image and which are not:

```
Dism /Image:C:\images\offline /Get-Features
```

Finally, you can use DISM to point to the mounted image and enable a desired feature. You can use the /All argument to enable all of the parent features in the same command. For example, to enable the Remote-Desktop-Services role and all parent features, type the following:

```
Dism /Image:C:\images\offline /Enable-Feature /FeatureName:Remote-Desktop-Services /All
```

If you want to remove or disable a feature from an offline image, use the /Disable-Feature switch. For example:

```
Dism /Image:C:\images\offline /Disable-Feature /FeatureName:Remote-Desktop-Services
```




EXAM TIP

You can use Dism and the /Add-Package option to apply to an image an update in the form of a .cab or .msu package. Use the /IgnoreCheck option if you don't want to verify the applicability of each package before installing. Use the /PreventPending option to skip the installation of the package if a system restart is required.

MORE INFO For more information on DISM in Windows Server 2012, visit <http://technet.microsoft.com/en-us/library/hh825236.aspx>.

Converting a server with a GUI to or from Server Core

As in Windows Server 2008 and Windows Server 2008 R2, Windows Setup in Windows Server 2012 and Windows Server 2012 R2 allows you to choose either of two installation types: Server Core Installation or Server With A GUI (also called a *full installation*), as shown in Figure 1-8. One of the more interesting features first introduced in Windows Server 2012 is the ability to convert a full installation to a Server Core Installation and vice-versa.

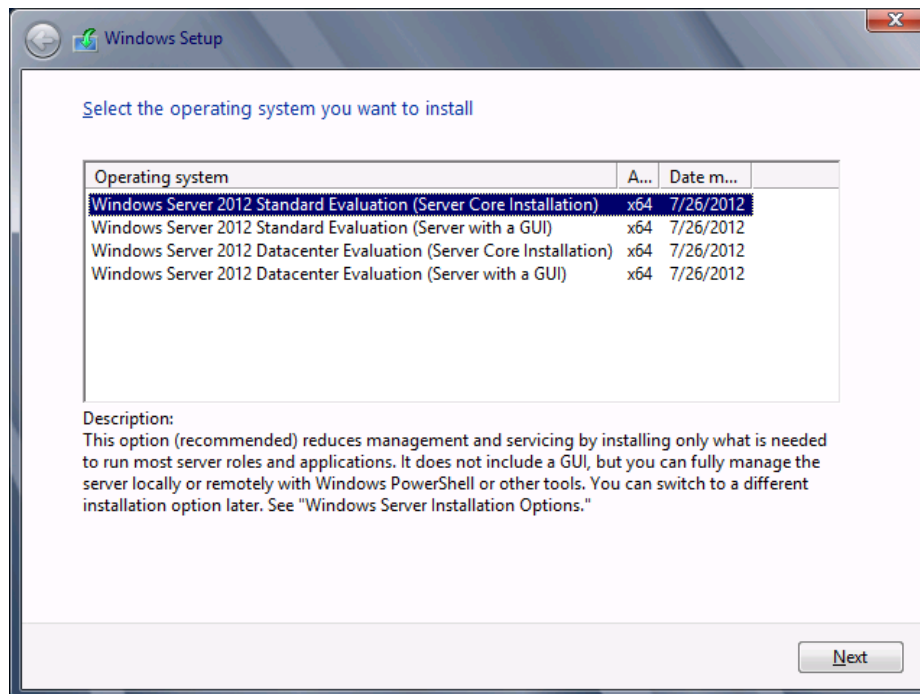


FIGURE 1-8 Windows Server 2012 and Windows Server 2012 R2 include a Server Core Installation option and a Server With A GUI option

You can switch between a Server Core installation and a full installation in Windows Server 2012 and Windows Server 2012 R2 because the difference between these installation options is now contained in two specific Windows features that can be added or removed. The first feature, Graphical Management Tools and Infrastructure (Server-Gui-Mgmt-Infra), provides a minimal server interface and server management tools such as Server Manager and the Microsoft Management Console (MMC). The second feature, Server Graphical Shell (Server-Gui-Shell), is built on this first feature and provides the rest of the GUI experience, including Windows Explorer. In Figure 1-9, you can see these two features in the Add Roles And Features Wizard, on the Select Features page, beneath User Interfaces And Infrastructure.

To convert a full installation to Server Core, simply remove these two features in Server Manager. Note that removing the first feature will automatically remove the second, dependent feature.

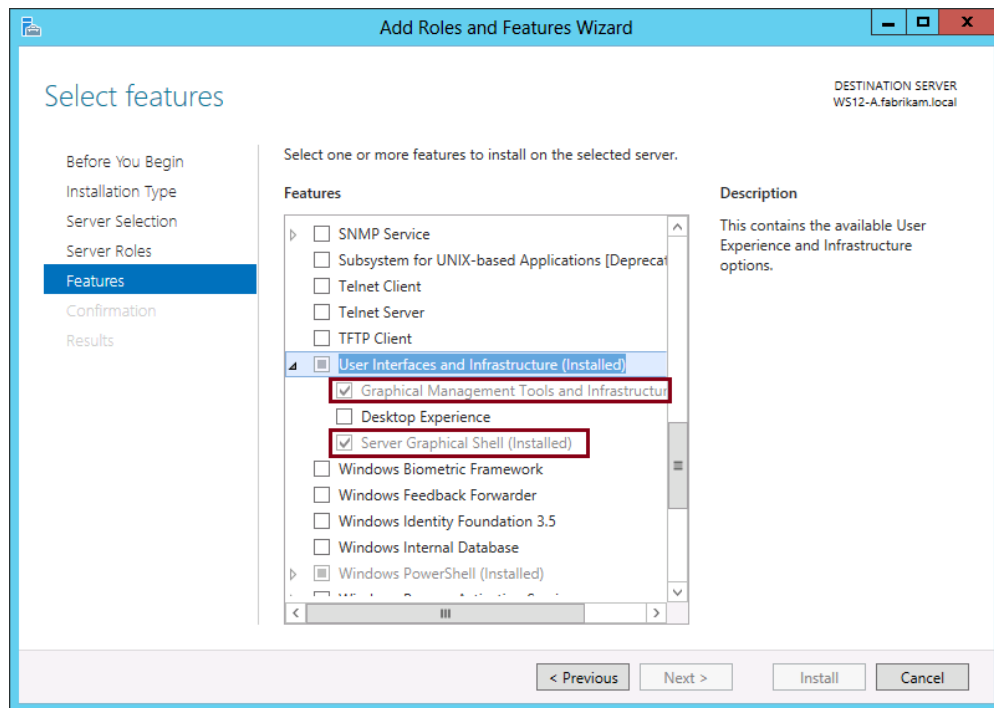


FIGURE 1-9 Two features are responsible for the difference between the full installation and Server Core

NOTE As shown in Figure 1-9, Desktop Experience is a third available GUI feature. It builds on the Server Graphical Shell feature and is not installed by default in the Server with a GUI installation of Windows Server 2012 or Windows Server 2012 R2. Desktop Experience makes available Windows 8 and Windows 8.1 client features such as Windows Media Player, desktop themes, and photo management.

You can remove these graphical interface features also in Windows PowerShell or by using the DISM utility at a normal command prompt. If you have deployed a GUI installation of Windows Server 2012 or Windows Server 2012 R2 and want to convert it into a Server Core installation, run the following Windows PowerShell command:

```
Uninstall-WindowsFeature Server-Gui-Mgmt-Infra -Restart
```

To perform the same step by using DISM, type the following:

```
Dism /Online /Disable-Feature /Featurename:ServerCore-FullServer  
/Featurename:Server-Gui-Shell /Featurename:Server-Gui-Mgmt
```

Remember that you need only to specify Server-Gui-Mgmt-Infra for removal in order to remove both this feature and Server-Gui-Shell. Once the graphical management tools and graphical shell have been removed, the server restarts. When you log back on, you are presented with the Server Core user interface.

The process can be reversed by adding both features back. You can do this from a remote server through the Add Roles And Features Wizard in Server Manager. You can also do it locally by running the following Windows PowerShell command:

```
Install-WindowsFeature Server-Gui-Mgmt-Infra,Server-Gui-Shell -Restart
```

Note that when you install these two features from Windows PowerShell, you must specify them both.

To use the DISM utility to add the GUI features back, type the following at an elevated prompt:

```
Dism /Online /Enable-Feature /Featurename:ServerCore-FullServer  
/Featurename:Server-Gui-Shell /Featurename:Server-Gui-Mgmt
```

NOTE If you just want to configure basic settings in a Server Core installation of Windows Server 2012 or Windows Server 2012 R2 as opposed to adding or removing entire features, you can use the Sconfig utility. This utility appeared in Windows Server 2008 R2 and allows you to set the domain/workgroup, computer name, Remote Desktop, network settings, date and time, Windows activation, Windows Update, and other similar settings.

Minimal Server Interface

The server with a GUI option is made up of two cumulative features built on top of Server Core in Windows Server 2012 and Windows Server 2012 R2. You have the option of installing only the first of these graphical features: Graphical Management Tools And Infrastructure (Server-Gui-Mgmt-Infra). Doing so results in what is called the Minimal Server Interface, shown in Figure 1-10. This form is not available when you install Windows Server 2012 or Windows Server 2012 R2, but you can configure it through Server Manager or Windows PowerShell. To configure a server with the Minimal Server Interface in Server Manager, begin

with a full installation and then simply remove the Server Graphical Shell feature by using the Remove Roles And Features Wizard. In Windows PowerShell, you can either begin with a full installation and remove only the Server-Gui-Shell feature or you can begin with a Server Core installation and add only the Server-Gui-Mgmt-Infra feature.

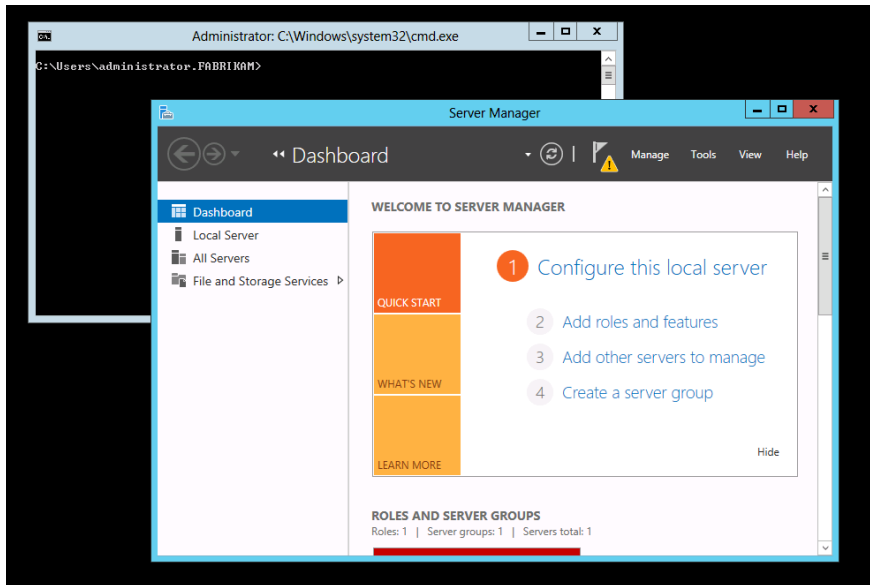


FIGURE 1-10 The new Minimal Server Interface option makes Server Manager and other administrative tools available without a desktop or Start screen

The relationship between the Minimal Server Interface and the Server With A GUI installation levels is illustrated in Figure 1-11.




FIGURE 1-11 Server-Gui-Shell is the difference between Server With A GUI and Minimal Server Interface

When you configure the Minimal Server Interface, the following elements are removed from the full installation:

- Desktop
- Start screen
- Windows Explorer
- Windows Internet Explorer

However, the following management tools *are* available in the Minimal Server Interface:

- Server Manager
- Microsoft Management Console (MMC) and snap-ins
- Subset of Control Panel

The Minimal Server Interface is a good option if you want to reduce the footprint of your installation but prefer not to be restricted to command-line-based management.



EXAM TIP

Expect to see questions on the 70-417 exam about converting between a Server Core installation, a Server With A GUI installation, and a Minimal Server Interface installation. Be sure to remember the command names of the features Server-Gui-Mgmt-Infra and Server-Gui-Shell, as well as how to remove the GUI by using Server Manager, Windows PowerShell, or the Dism utility.

MORE INFO For more information about converting between installation options in Windows Server 2012, see “Server Core and Full Server Integration Overview” at <http://technet.microsoft.com/en-us/library/hh831758.aspx> and “Windows Server Installation Options” at <http://technet.microsoft.com/en-us/library/hh831786.aspx>.

NIC teaming

NIC teaming, also known as *Load Balancing and Failover (LBFO)*, is a feature that was introduced in Windows Server 2012 and that enables multiple network adapters on a server to be grouped together into a team. NIC teaming has two purposes:

- To help ensure the availability of network connectivity if one adapter fails
- To aggregate network bandwidth across multiple network adapters

Before Windows Server 2012, implementing network adapter teaming on Windows Server required using third-party solutions from independent hardware vendors. However, network adapter teaming is now built into the Windows Server operating system and can therefore work across different NIC hardware types and manufacturers.

Windows NIC teaming supports up to 32 network adapters in a team in three modes:

- **Static Teaming** Also called *Generic Teaming*, the *Static Teaming* mode is based on IEEE 802.3ad draft v1 and is supported by most server-class Ethernet switches. It requires manual configuration of the switch and the server to identify which links form the team.
- **Switch Independent** The *Switch Independent* mode allows each NIC in a team to connect to different switches.
- **LACP** Also called *Dynamic Teaming*, the *LACP* mode is based on IEEE 802.1ax and is supported by most enterprise-class switches. It allows teams to be automatically created through the *Link Aggregation Control Protocol (LACP)*. LACP dynamically identifies links between the server and a specific switch. To use this mode, you generally need to enable LACP manually on the port of the switch.

NIC teaming can be enabled from Server Manager or by using Windows PowerShell. In Server Manager, you can begin by right-clicking the server you want to configure and selecting Configure NIC Teaming, as shown in Figure 1-12.

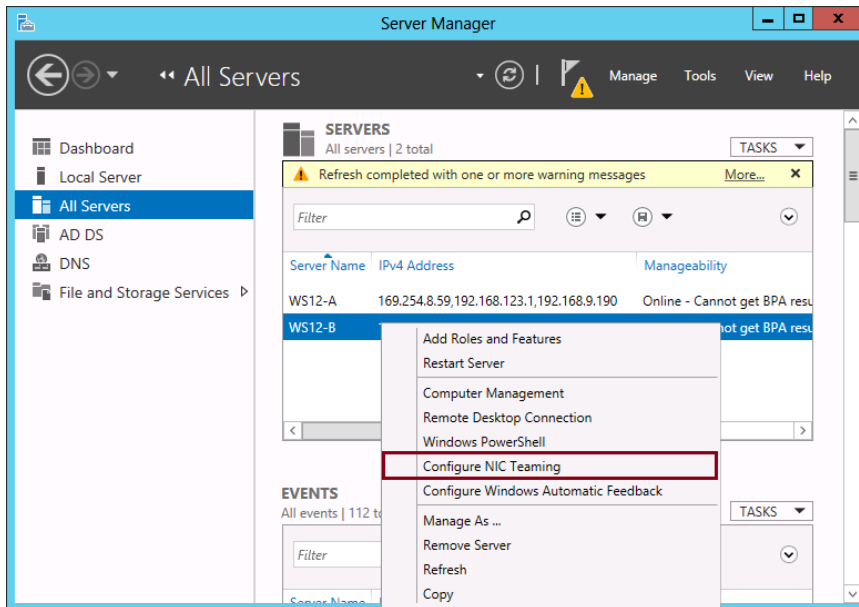


FIGURE 1-12 Configuring NIC teaming in Server Manager

In the NIC Teaming dialog box that opens, select the network adapters you want to team. Then right-click and select Add To New Team, as shown in Figure 1-13.

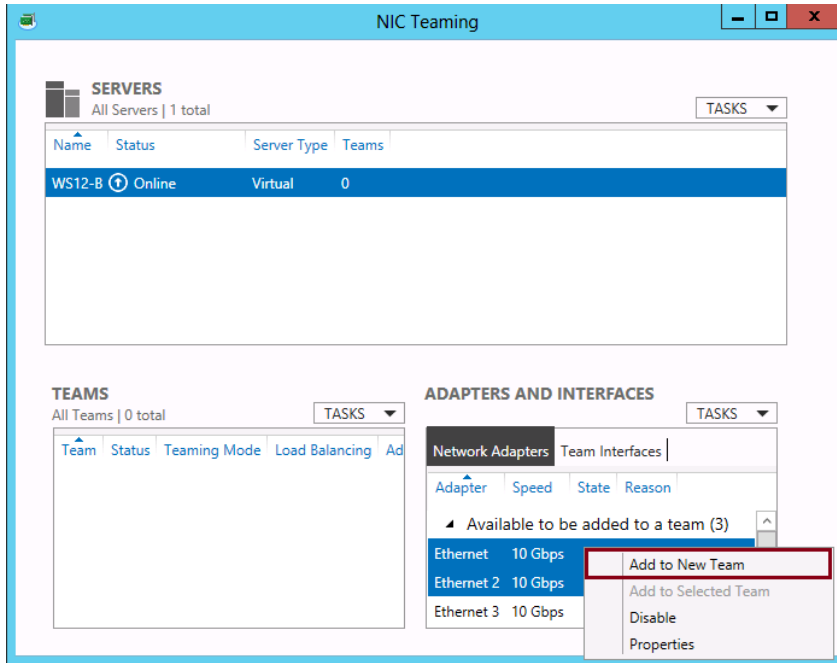


FIGURE 1-13 Adding network adapters to a new team

In the New Team dialog box, shown in expanded mode in Figure 1-14, you can configure the teaming mode and other settings as you want.

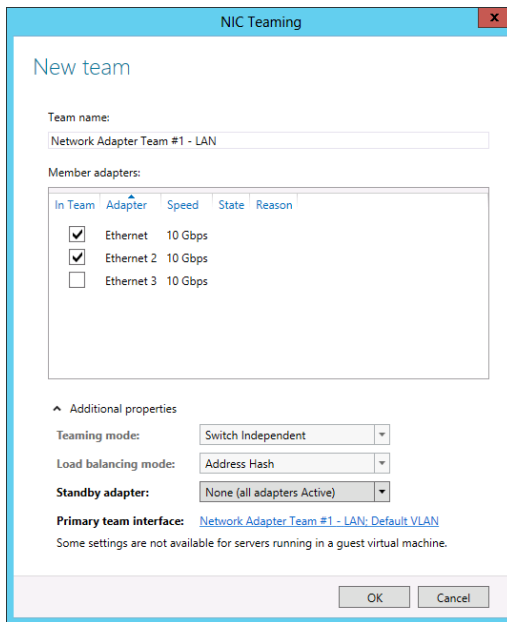


FIGURE 1-14 Configuring team properties

Clicking OK completes the process and, if successful, the new team will be displayed in both the Teams area and the Adapters And Interfaces area of the NIC Teaming dialog box, shown in Figure 1-15.

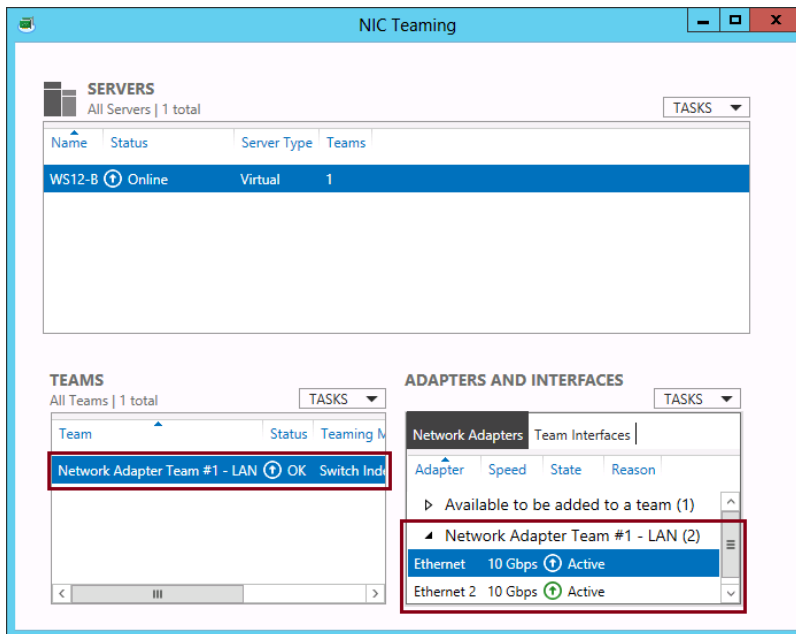


FIGURE 1-15 A newly configured network team

To configure and manage NIC teaming in Windows PowerShell, use cmdlets such as `New-NetLbfoTeam` to add a new team or `Get-NetLbfoTeam` to display the properties of a team. The cmdlets for managing NIC teaming are defined in the Windows PowerShell module named `NetLbfo`, and as Figure 1-16 shows, you can use the `Get-Command` cmdlet to display all the cmdlets defined in this module. You could then use the `Get-Help` cmdlet to learn the syntax for any of the functions displayed. For example, type **`Get-Help New-NetLbfoTeam`** to find out more about the `New-NetLbfoTeam` cmdlet.

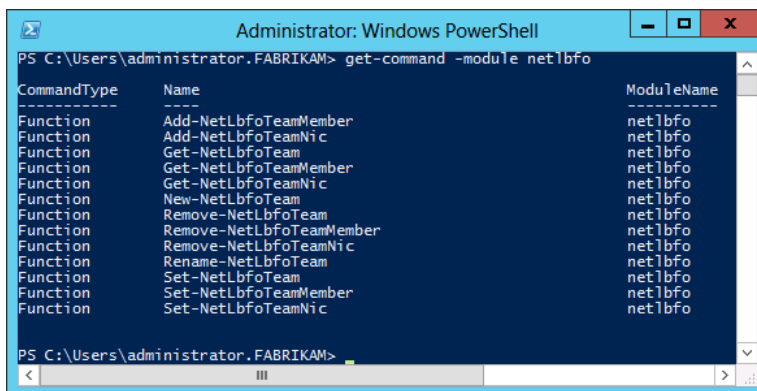


FIGURE 1-16 Cmdlets for NIC teaming



EXAM TIP

Remember that each network adapter team is assigned a single IP address.

MORE INFO For more information about NIC teaming in Windows Server 2012, see the NIC teaming overview at <http://technet.microsoft.com/en-us/library/hh831648.aspx>. For more in-depth information, search for the white paper titled "Windows Server 2012 NIC Teaming (LBFO) Deployment and Management" on <http://technet.microsoft.com>.

Objective summary

- The Dism.exe utility was introduced in Windows 7 and Windows Server 2008 R2. It allows you to service WIM files, VHD files, VHDX files, and online installations of Windows, including adding and removing features, packages, and drivers.
- New to Windows Server 2012 and Windows Server 2012 R2 is the ability to deploy roles and features to remote servers.

To perform this task in Windows PowerShell, use the following command:

```
Install-WindowsFeature -Name <feature_name> -ComputerName <computer_name> -Restart
```

To perform this task in Server Manager, you first need to add the remote server to the server pool. Then install the role or feature as you would to the local server.

- Windows PowerShell Desired State Configuration (DSC) is a new framework in Windows Server 2012 R2 that allows you to use Windows PowerShell scripting to ensure that selected servers are configured properly. To use DSC, you create a Configuration function in Windows PowerShell that defines a configuration for specified servers (nodes) and resources. You then call the function to generate an MOF file in a subdirectory. Finally, to ensure that the target server is configured according to your specification, you use the Start-DscConfiguration cmdlet to specify a path to the MOF file.
- In Windows Server 2012 and Windows Server 2012 R2 you can convert between a Server Core installation and a full (Server With A GUI) installation. To do so, you can begin from a full installation and then type the following command in Windows PowerShell:

```
Uninstall-WindowsFeature Server-Gui-Mgmt-Infra -restart
```

If you later want to reinstall the full graphical interface, type the following command in Windows PowerShell:

```
Install-WindowsFeature Server-Gui-Mgmt-Infra,Server-Gui-Shell -Restart
```

- NIC teaming is a feature introduced in Windows Server 2012 that allows you to group two or more NICs together to aggregate bandwidth and help ensure the availability of network connectivity. You can configure NIC teaming by using Server Manager or the `New-NetLbfoTeam` cmdlet.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the “Answers” section at the end of the chapter.

1. You work in the IT department for Contoso.com, which has approximately 200 employees. Your manager has created a new image named `Basic.wim` that will be used to deploy Windows Server 2012 R2. She has now asked you to modify the image with an index of 1 within this image file so that the IIS-WebServer feature is disabled. You move the `Basic.wim` file from network storage to your server, which is running Windows Server 2012 R2. Which of the following actions should you take next?
 - A. Use DISM with the `/Mount-Image` option.
 - B. Use DISM with the `/Disable-Feature` option.
 - C. Use the `Uninstall-WindowsFeature` cmdlet without the `-Remove` parameter.
 - D. Use the `Uninstall-WindowsFeature` cmdlet with the `-Remove` parameter.
2. You want to install Windows Server 2012 R2 and configure an interface that includes Server Manager but not Windows Explorer. What should you do? (Choose two.)
 - A. Choose the Server Core installation of Windows Server 2012 R2.
 - B. Choose the Server With A GUI installation of Windows Server 2012 R2.
 - C. Remove the Graphical Management Tools and Infrastructure feature.
 - D. Add the Graphical Management Tools and Infrastructure feature.
3. You have built a new server with network adapters from two different manufacturers. You want to use these two adapters to provide resiliency for the server’s network connectivity, so that in case one adapter fails, the other will continue to operate with the same configuration settings. What should you do?
 - A. Install the Network Load Balancing feature.
 - B. Install the Multipath I/O feature.
 - C. Use the `New-NetLbfoTeam` cmdlet.
 - D. Use the `Set-NetConnectionProfile` cmdlet.

4. You want to ensure that a server named Web1 has Internet Information Services installed. Web1 is currently running Windows Server 2012 R2. On Web1, you create the following function in Windows PowerShell ISE and then run the code.

```
Configuration MyConfig
{
    Node "Web1"
    {
        WindowsFeature IIS
        {
            Ensure = "Present"
            Name = "Web-Server"
        }
    }
}
```

You want to check the configuration of Web1 as soon as possible. What should you type next at the Windows PowerShell prompt?

- A. MyConfig
- B. MyConfig Web1
- C. Start-DscConfiguration .\MyConfig
- D. Start-DscConfiguration .\Web1

Objective 1.3: Configure local storage

For the 70-417 exam, this objective is likely to focus on Storage Spaces, an interesting feature introduced in Windows Server 2012 that adds SAN-like flexibility to your local storage. The topic of Storage Spaces can be broken down into primordial pools, new storage pools, and virtual disks.

This section covers the following topics:

- Creating and configuring storage pools
- Provisioning virtual disks
- Designing storage spaces

Introducing Storage Spaces

Storage Spaces is a feature introduced in Windows Server 2012 that provides for a single server the same storage flexibility provided by a storage area network (SAN)—but by using inexpensive, locally-attached disks. Installed by default, Storage Spaces allows you to create storage pools from which you can provision storage as needed.

Once you've created a storage pool using Storage Spaces, you can provision storage from the pool by creating virtual disks, also called *logical unit numbers (LUNs)*. A virtual disk behaves exactly like a physical disk except that it can span multiple physical disks within the storage pool.

Storage Spaces has the following requirements:

- Windows Server 2012 or Windows Server 2012 R2.
- One physical drive is required to create a storage pool; a minimum of two physical drives are needed to create a resilient mirror storage space.
- A minimum of three physical drives is required to create a storage space with resiliency through parity or three-way mirroring.
- Drives must be unpartitioned and unformatted.
- Drives must have at least 10 GB capacity.
- Drives can be attached either internally or externally (individually or in a just-a-bunch-of-disks [JBOD] enclosure). The following bus technologies are supported:
 - SATA (not possible to use in a failover cluster).
 - SCSI (not supported in a failover cluster).
 - Serial Attached SCSI (SAS) arrays that support SCSI Enclosure Services (SES).
 - USB (external drives for local storage only; not possible to use in a failover cluster or recommended for file servers).

Installing Storage Spaces

To install Storage Spaces, use the Add Roles And Features Wizard to add the File Server role service. This role service is found under File and iSCSI Services in the File and Storage Services role. You can also install the File Server role service by using Windows PowerShell as follows:

```
Install-WindowsFeature -Name FS-FileServer
```

NOTE Storage Services, another role service of the File and Storage Services role, is always installed by default on Windows Server 2012 and Windows Server 2012 R2 and provides general storage management functionality needed by other server roles.

Creating a storage pool

To create a storage pool, Storage Spaces requires a server to have at least one attached physical disk of at least 10 GB without any partitions or volumes. Any physical disks that meet these two criteria are automatically added to what is called the server's *primordial pool*. The primordial pool is the complete set of locally available disks from which a storage pool can be created. Figure 1-17 shows in Server Manager the primordial pools available to the server named WS12-A and WS12-B, respectively.



EXAM TIP

Physical disks that are initialized in Server Manager are automatically configured with the GUID Partition Table (GPT) partition style, not the Master Boot Record (MBR) partition style.

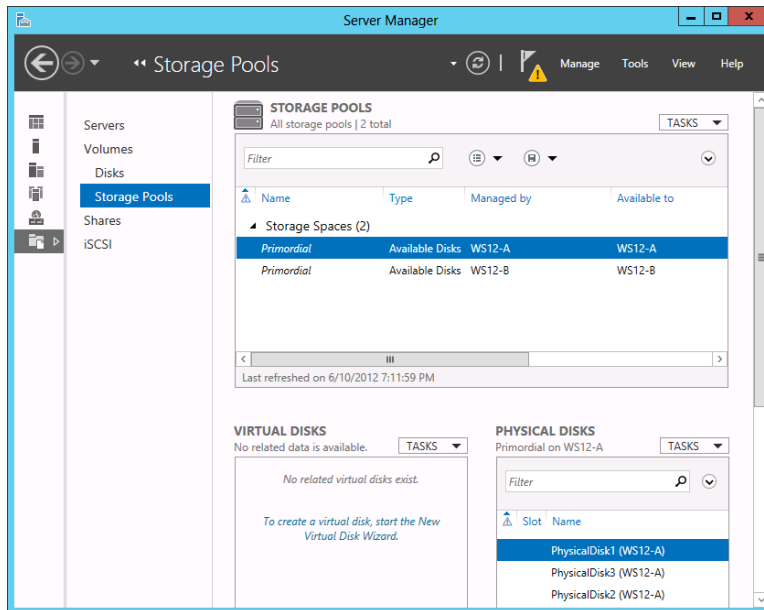


FIGURE 1-17 A primordial pool is composed of all unallocated physical disks larger than 10 GB available to a server

You can use Server Manager or Windows PowerShell to configure your storage pools from a primordial pool. To create a storage pool in Windows PowerShell, use the `New-StoragePool` cmdlet. To create a new storage pool using Server Manager, first make sure that you have navigated to File And Storage Services\Volumes\Storage Pools. Then select New Storage Pool from the Tasks menu in the Storage Pools area, as shown in Figure 1-18.

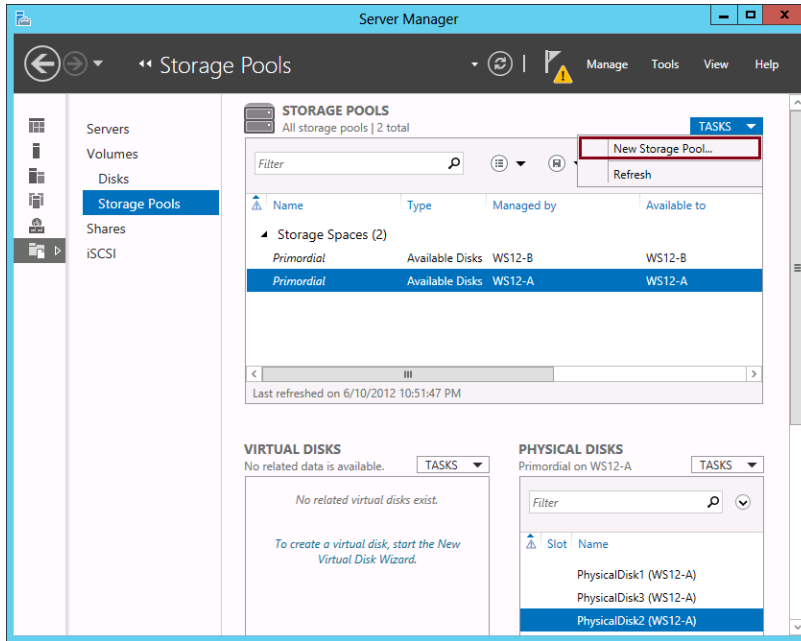


FIGURE 1-18 Creating a new storage pool

This step opens the New Storage Pool Wizard. After specifying a server (primordial pool) and name for your new pool, you can select which physical disks you want to include in your pool, as shown in Figure 1-19.

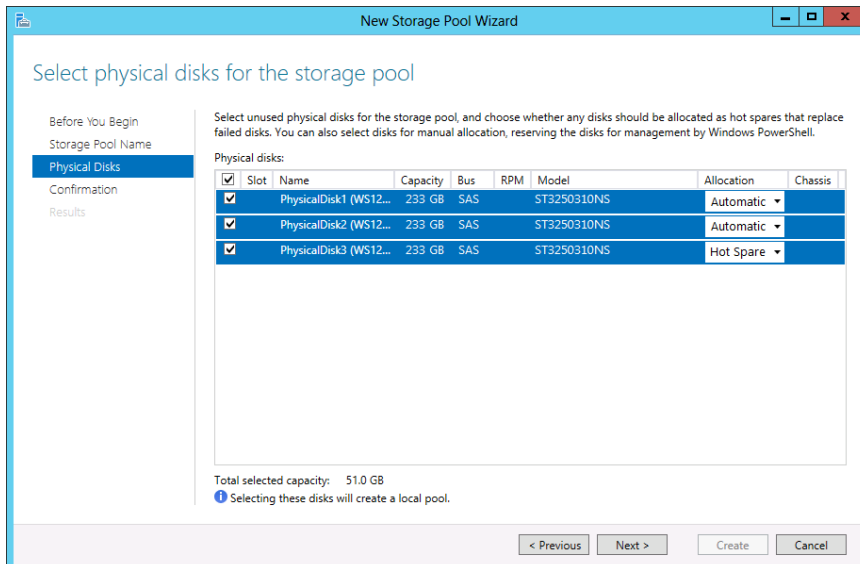


FIGURE 1-19 Selecting which physical disks to add to the storage pool



EXAM TIP

Remember that if you want the storage pool to support failover clusters, you have to use SAS storage arrays that support SES.

For each disk that you add to the pool, you can choose one of the following allocation types:

- **Automatic** This is the default setting. For this allocation type, the capacity on drives is set automatically.
- **Hot Spare** Physical disks added as hot spares to a pool act as reserves that are not available for provisioning in the creation of virtual disks. If a failure occurs on a drive in a pool that has an available hot spare, the spare will be brought online to replace the failed drive.



EXAM TIP

If you want to add a hot spare to your storage pool and plan to create a mirrored drive later, you need at least three physical disks in the storage pool: one for the hot spare and two to support the mirror.

Creating virtual disks

After a storage pool is created, you can use Server Manager to provision new virtual disks from this new available storage. These new virtual disks will appear as unallocated disks in Disk Management, from which you can then create volumes. Note that a virtual disk is the representation of virtualized storage and should not be confused with the VHD that is used in the context of Hyper-V or the iSCSI Target Server.

To create a virtual disk in Windows Powershell, use the `New-VirtualDisk` cmdlet. To create a virtual disk in Server Manager, complete the following steps:

1. In Server Manager, choose File And Storage Services and then Storage Pools.
2. Locate a storage pool (not a primordial pool) that you want to use to support the new virtual disk.
3. Right-click the storage pool and select New Virtual Disk to start the New Virtual Disk Wizard, as shown in Figure 1-20.
4. On the first pages of the wizard, verify that the correct server and storage pool are selected, and provide a name and description for the new virtual disk.
5. On the Select The Storage Layout page (see Figure 1-21), specify one of the following three data redundancy types for the virtual disk:
 - **Simple** A simple virtual disk provides data striping across physical disks, but no redundancy. Administrators should not host irreplaceable user data on a simple

space. A simple space maximizes capacity and throughput and therefore can be a good candidate for hosting temp files or easily re-created data at a reduced cost.

- **Parity** A parity virtual disk is similar to a hardware Redundant Array of Inexpensive Disks RAID5. Data, along with parity information, is striped across multiple physical disks. Parity enables Storage Spaces to continue to service read and write requests even when a drive has failed, and it provides this fault tolerance with efficient use of storage. A minimum of three physical disks is required for parity virtual disks. Note that a parity disk cannot be used in a failover cluster.
- **Mirror** A mirror virtual disk maintains either two or three copies of the data it hosts: two data copies for two-way mirror spaces or three data copies for three-way mirrors. All data writes are repeated on all physical disks to ensure that the copies are always current. Mirror spaces are attractive due to their greater data throughput and lower access latency compared to parity disks.

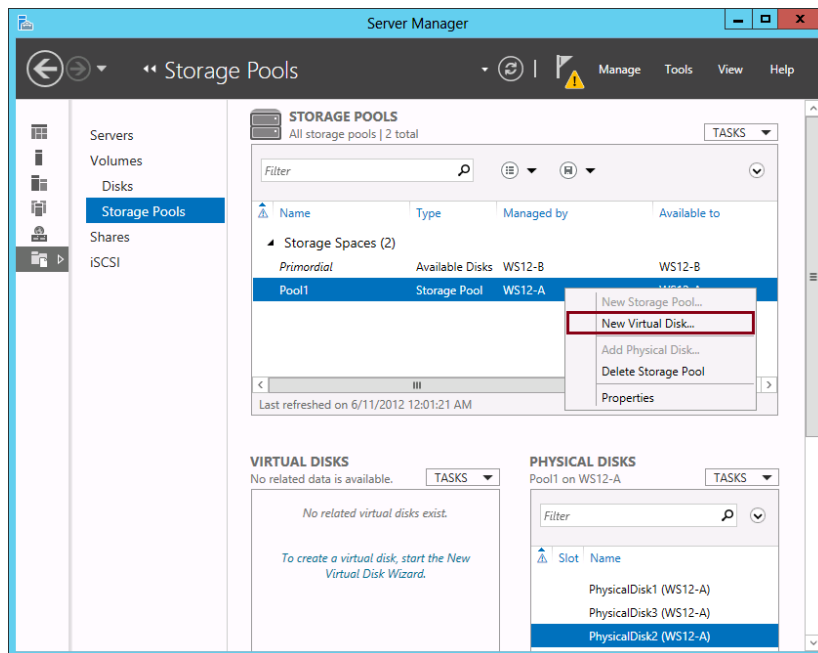


FIGURE 1-20 Creating a new virtual disk in a storage pool



EXAM TIP

Make sure you understand the advantages and disadvantages of simple, parity, and mirror spaces.

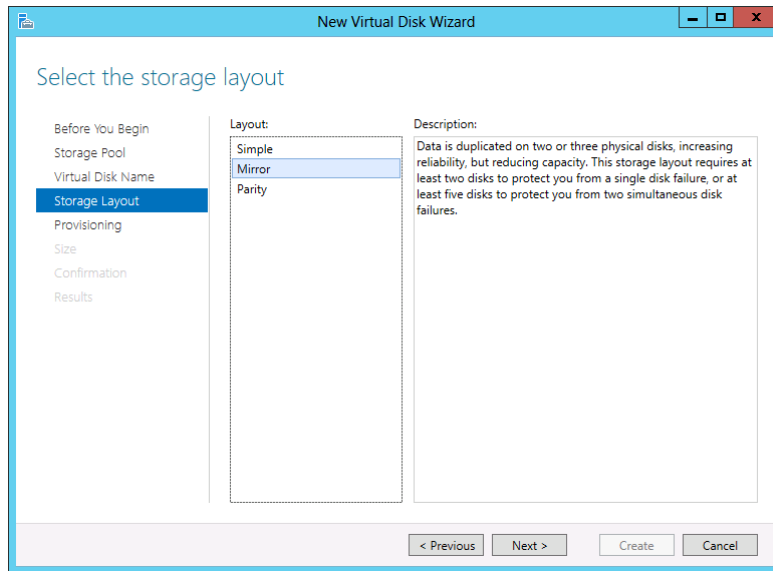


FIGURE 1-21 Selecting a storage layout

6. On the Specify The Provisioning Type page, choose one of the following provisioning types:
 - **Thin** Thin provisioning is a mechanism that allows storage capacity not to be allocated until datasets require the storage. You specify a maximum size for the virtual disk, and the capacity of the virtual disk grows as needed. Thin provisioning optimizes utilization of available storage, but it adds a few extra I/Os that can cause an occasional latency increase.
 - **Fixed** Fixed provisioned spaces allocate storage capacity upfront, at the time the space is created.



EXAM TIP

Expect to see a reference to thin provisioning on the 70-417 exam.

7. On the Specify The Size Of The Virtual Disk page, choose a size for the virtual disk.
8. Confirm all the selections and then click Create.

The new virtual disk appears in both Server Manager and Disk Management. The view in Server Manager is shown in Figure 1-22.




FIGURE 1-22 A new virtual disk created from a storage pool in Server Manager

Objective summary

- Storage Spaces was introduced in Windows Server 2012 and provides flexible provisioning of local storage to a server.
- All locally attached, unpartitioned physical disks with a capacity of at least 10 GB are automatically added to a server's primordial pool. A primordial pool is the complete set of locally available disks from which a storage pool can be created.
- Storage pools can be created from one or more physical disks. If you want to be able to create a mirrored virtual disk later from a storage pool, you need to add at least two physical disks to that storage pool. If you want to be able to create a virtual disk with parity later from a storage pool, you need to add at least three physical disks to that storage pool. On top of those requirements, you need to add one additional physical disk to a storage pool for each hot spare you want to be available to the storage.

- Thin provisioning is a new feature in Windows Server 2012 and Windows Server 2012 R2 that allows you to create drives that don't require all of their storage capacity to be allocated immediately. Thin provisioning optimizes available storage capacity for virtual disks.
- When you create new virtual disks from a storage pool, they appear in Disk Management as new, unallocated disks.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the "Answers" section at the end of the chapter.

1. You want to create a storage pool that can be used with a failover cluster. Which of the following disk types can you use?
 - A. Internal SCSI
 - B. Serial-Attached SCSI
 - C. Internal SATA
 - D. iSCSI
2. You want to create a storage pool that maximizes available storage capacity while including built-in fault tolerance and data resiliency. You also want to include a hot spare so that if a physical disk fails, another will be brought online to replace it. Which configuration should you choose for your storage pool? Assume you want to use the minimum number of physical disks possible.
 - A. Three physical disks and a mirror layout
 - B. Three physical disks and a parity layout
 - C. Four physical disks and a mirror layout
 - D. Four physical disks and a parity layout
3. You want to increase the size of a server's primordial pool. Which of the following disks can you use? (Choose all that apply.)
 - A. A 20 GB external USB drive
 - B. A 12 GB internal SCSI drive
 - C. An 8 GB SATA drive
 - D. A 5 GB Serial-Attached SCSI drive



Thought experiment

Installing and configuring servers at Fabrikam

You work as a network administrator in a company named Fabrikam.com. The Fabrikam Finance department requires a new server to support a web application that runs on Windows Server 2012 R2. Your manager asks you to help design and plan for the server. She specifies the following requirements:

- **Operating system installation and configuration**

The application requires a GUI to be installed. The server should not be limited to command-line administration only, but it can be managed through remote administration. Within these limitations, the attack surface of the server must be minimized, and performance must be optimized.

- **Network**

The network requires a single network connection with fault tolerance, so if one adapter fails, there will be no loss in connectivity, and the IP address will not change.

- **Storage**

You have an eight-bay disk array and eight 1 TB SATA disks available to attach to the new server. The disk array will be reserved for data storage. Your manager wants to use as few of these disks as possible while using the Storage Spaces feature to meet the following storage requirements:

- **Virtual Disk 1:** Triple mirrored with a capacity of 100 GB
- **Virtual Disk 2:** Parity disk with a capacity of 200 GB
- **One hot spare for each storage pool**

How do you answer the following design questions from your manager? You can find the answers to these questions in the “Answers” section.

1. How should you reconcile the requirement for a GUI during installation with the need to minimize the attack surface?
2. Is this server an ideal candidate for a Minimal Server Interface configuration? Why or why not?
3. What are two possible solutions to meet the needs for fault-tolerant network connectivity?
4. How many physical disks do you need to reserve for the application at a minimum? What is the minimum number of storage pools you need?

Answers

This section contains the answers to the Objective Reviews and the Thought Experiment.

Objective 1.1: Review

1. Correct answers: B, C

- A. Incorrect:** You cannot install Windows Server 2012 or Windows Server 2012 R2 on the existing server because those operating systems require a 64-bit CPU.
- B. Correct:** You need to replace the existing server with another that has a 64-bit CPU.
- C. Correct:** This command removes from disk the feature files of any uninstalled roles and features.
- D. Incorrect:** This command uninstalls every feature that is currently installed. It doesn't reduce the size of the installation image.

2. Correct answers: A, B

- A. Correct:** This command will delete from disk all the feature files for the Web Server role.
- B. Correct:** This command will remove from disk all the feature files for the Web Server role except for the manifest file.
- C. Incorrect:** You should not execute this command because it will remove the File Server role service from the server, and the server's only stated function is that of a file server.
- D. Incorrect:** You should not execute this command because it will remove the File Server role service from the server, and the server's only stated function is that of a file server.

3. Correct answers: A, D

- A. Correct:** When you specify a path to source files with the Install-WindowsFeature cmdlet, you can use the WIM: prefix at the beginning and specify at the end a particular index number to the image that contains the source files you want.
- B. Incorrect:** You need to use the WIM: prefix before the path and the source.
- C. Incorrect:** You need to specify an image index after the path.
- D. Correct:** You can specify a path to the \sources\sxs directory on the product DVD.

Objective 1.2: Review

1. **Correct answer:** A

- A. Correct:** You need to mount the image file before you can service it.
- B. Incorrect:** You need to mount the image before you can disable any feature in it.
- C. Incorrect:** You can use Uninstall-WindowsFeature on a VHD file, but not a WIM file.
- D. Incorrect:** You can use Uninstall-WindowsFeature on a VHD file, but not a WIM file.

2. **Correct answers:** A, D

- A. Correct:** The interface requirements describe the Minimal Server Interface. To configure this interface type, you can either start with a Server Core installation and add Graphical Management Tools And Infrastructure, or you can start with a Server With A GUI installation and remove Server Graphical Shell. Removing Server Graphical Shell is not provided as an answer choice, so you have to start with a Server Core.
- B. Incorrect:** Removing Server Graphical Shell is not provided as an answer choice, so you cannot start with a full installation of Windows Server 2012 R2.
- C. Incorrect:** Removing the Graphical Management Tools And Infrastructure feature would transform a full installation into a Server Core installation. The Server Core installation would not make Server Manager available.
- D. Correct:** The Graphical Management Tools And Infrastructure feature includes Server Manager and some other basic administrative tools, but it does not include Windows Explorer. Adding this feature to a Server Core installation would result in the desired configuration.

3. **Correct answer:** C

- A. Incorrect:** Network Load Balancing is used to configure many different servers to answer requests for a service at a single address.
- B. Incorrect:** Multipath I/O is used to provide multiple data paths to a storage device.
- C. Correct:** This cmdlet is used to create a new NIC team. NIC teaming is used to provide failure resiliency to physical network connections.
- D. Incorrect:** This cmdlet is used to set a profile to a network connection. It is not used to provide failure resiliency to physical network connections.

4. Correct answer: A

- A. Correct:** After you create the function, you need to invoke or call the function by typing its name. This step will generate an MOF file that provides configuration information that can be used with the Start-DscConfiguration cmdlet. The MOF file will be created in a subdirectory named after the new function.
- B. Incorrect:** You would specify the name of the server only if an appropriate parameter had been defined for the node in the function. The MyConfig function specifies a particular server by name.
- C. Incorrect:** You need to run this command after you call the function. Running this command will complete the process and ensure that IIS is installed on Web1.
- D. Incorrect:** Although you need to run the Start-DscConfiguration cmdlet, the path specified here is incorrect. You would need to specify the .\MyConfig path, not the .\Web1 path.

Objective 1.3: Review

1. Correct answer: B

- A. Incorrect:** SCSI disks cannot be used in failover clusters.
- B. Correct:** Serial-Attached SCSI (SAS) disks can be used to create a storage pool that can support a failover cluster.
- C. Incorrect:** SATA disks cannot be used in failover clusters.
- D. Incorrect:** iSCSI disks cannot be used to create a storage pool that can support a failover cluster.

2. Correct answer: D

- A. Incorrect:** You want a parity layout so that the storage capacity is maximized.
- B. Incorrect:** You need four physical disks: three to support the parity layout and a fourth for the hot spare.
- C. Incorrect:** You want a parity layout so that the storage capacity is maximized.
- D. Correct:** You want a parity layout so that the storage capacity is maximized, and you need four physical disks: three to support the parity layout and a fourth for the hot spare.

3. Correct answers: A, B

- A. Correct:** You can use external USB drives in storage pools as long as they are at least 10 GB in size.
- B. Correct:** You can use SCSI drives in storage pools as long as they are at least 10 GB in size.
- C. Incorrect:** You can use SATA drives in storage pools, but they need to be at least 10 GB in size.
- D. Incorrect:** You can use Serial-Attached SCSI drives in storage pools, but they need to be at least 10 GB in size.

Thought experiment

- 1.** You can first install Windows Server 2012 R2 with a full GUI and then install the application. After you install the application, you can remove the GUI features.
- 2.** This is not an ideal candidate for a Minimal Server Interface configuration because even though its administration should not be restricted to the command line, it can be managed through a GUI on remote computers. The requirements of a minimal attack surface and optimal performance suggest that a Server Core installation is a better fit for this scenario.
- 3.** Two possible solutions to meeting the requirements for network fault tolerance are using the built-in NIC teaming feature of Windows Server 2012 R2 or using NIC teaming provided by an independent hardware vendor of network adapters.
- 4.** You need four disks to meet the requirements of this configuration. You can create both virtual disks out of one storage pool. Both of these virtual disks require three physical disks, but they can be provisioned out of the same pool. In addition, the hot spare requires a fourth, separate physical disk to be assigned to the storage pool.

This page intentionally left blank

Configure server roles and features

Exam 70-417 distills the three Windows Server 2012 MCSA certification track exams (70-410, 70-411, and 70-412) down to a single exam. The raw material from these three original MCSA exams includes a total of 18 major content areas or “domains” within Windows Server 2012. These 18 content areas are then further broken down into 62 specific objectives corresponding to job tasks.

That’s a lot of source material, but just a fraction—about one-third—of this original content is specified for 70-417. Specifically, just 22 objectives within 14 content areas made it into the official “Skills Measured” list for 70-417. These 14 official content areas are represented as the chapter names in this book.

As a result, most of the content areas in 70-417 appear only in a partial form compared to the original exams from which they were taken. This chapter is an example of that. The Configure Server Roles and Features domain is taken from the 70-410 exam, but only one of the three original objectives remains.

The questions you see from the Configure Server Roles and Features domain on the 70-417 exam will fall *almost* completely within Objective 2.1, within the topic of remote management. However, it’s a good idea to review the Microsoft website for the original objectives associated with this content area on the 70-410 exam. Don’t be surprised if on 70-417 you see a question relating to any of the original topics.

Objectives in this chapter:

- Objective 2.1: Configure servers for remote management

Objective 2.1: Configure servers for remote management

Windows Server 2012 and Windows Server 2012 R2 are much better suited to administering remote servers on the network than any of their predecessors. It’s not just that these operating systems offer new capabilities in remote management, which they do. It’s also that behind the scenes, existing technologies have been revised to simplify remote management.

As you study this section for the exam, above all, do not rely on what you learned for Windows Server 2008 and Windows Server 2008 R2. Some features might look the same, but they have in fact changed—such as the inbound rules you might use to enable various types of remote management, or the function of a particular command, or the name of the relevant Group Policy settings.

This section covers the following topics:

- Managing multiple servers with Server Manager
- Configuring various server types for remote management
- Configuring Group Policy to enforce remote management settings

Managing multiple servers with Server Manager

The new Server Manager in Windows Server 2012 and Windows Server 2012 R2 reveals big changes—both cosmetic and functional—and some of these changes are relevant for the 70-417 exam. The most significant of these new features is that you can now use Server Manager to manage multiple servers, as shown in Figure 2-1. One way you can use Server Manager to manage multiple servers is through the All Servers option in the navigation pane. You can also use the Create Server Group option on the Manage menu to create server groups, which are custom pages in Server Manager that allow you to manage subsets of your servers such as all your DNS servers. Finally, you can also multi-select servers on Server Manager pages, allowing you to perform some management tasks simultaneously on all the servers you selected.

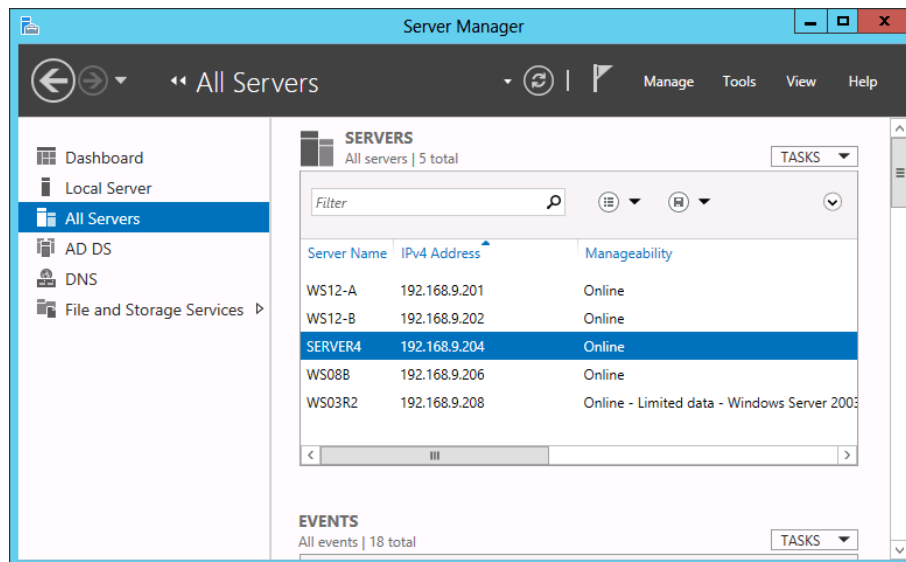


FIGURE 2-1 Server Manager is a multiserver management tool in Windows Server 2012 and Windows Server 2012 R2

MORE INFO See Chapter 1, “Install and Configure Servers,” for information about how to add servers to the All Servers page in Server Manager.

Adding non-domain-joined servers to Server Manager

It’s easy to add and manage servers through Server Manager when all of the servers are in the same Active Directory domain. These steps were shown in Chapter 1, “Install and Configure Servers,” in the section “Objective 1.2: Configure Servers.” Adding servers in the same domain to Server Manager is so simple, in fact, that it’s almost certainly too simple to serve as the basis for an exam question.

Managing non-domain joined servers through Server Manager is, however, more complicated and requires a level of expertise that is more suitable for the exam. Typically, workgroup environments are not covered in MCSE exams, but “manage non-domain joined servers” has officially been added as a new task in the updated objectives for Windows Server 2012 R2. For these reasons, you will likely see a question on the 70-417 exam that tests your knowledge of this topic.

Simply adding workgroup servers to Server Manager is not difficult. You can easily add a workgroup server to Server Manager by using the Add Servers option on the Manage menu and then specifying it by IP address or DNS name as long as the target server is accessible on the network. However, the remote non-domain-joined server will appear in Server Manager only with an error message at first; no management operations will be possible immediately. These complications with managing remote workgroup servers in Server Manager have to do with the following four issues:

NOTE Successful name resolution is an assumed prerequisite for remote management. The managing server and target server must be able to resolve each other’s computer names through DNS.

1. Remote computers aren’t automatically trusted for remote management on a managing server (as they are in a shared domain environment).

This is the only issue that is always present when the managing server running Server Manager and the server being targeted in Server Manager are not members of the same Active Directory domain.

To fix this problem, you need to configure the server running Server Manager to add the target server to its TrustedHosts list for remote management. To add the first server to the list, type the following at an elevated Windows PowerShell prompt:

```
Set-Item wsman:\localhost\Client\TrustedHosts TargetServer
```

where TargetServer is the name of the target server you want to manage through Server Manager.

To add additional servers to the list, add the `-Concatenate` and `-Force` options. For example, to add the server named `Server2` to a `TrustedHosts` list that already includes one or more entries, type the following at an elevated Windows PowerShell prompt:

```
Set-Item wsman:\localhost\Client\TrustedHosts Server2 -Concatenate -Force
```

If you neglect to use the `-Concatenate` and `-Force` options when the `TrustedHosts` list contains pre-existing entries, the pre-existing entries will be deleted. To see the existing `TrustedHosts` lists for remote management, type the following at an elevated command prompt:

```
winrm get winrm/config/client
```

In the output, read the values next to `TrustedHosts`. If no computer name entries are present, you don't need to use the `-Concatenate` and `-Force` options.

2. The correct administrator credentials aren't automatically available for the remote server because there is no one administrator account common to both the managing server and the target server.

This issue is present whenever the user name and password of the currently logged-in administrator on the server running Server Manager do not exactly match the credentials (user name and password) of an administrator on the target server.

To fix this problem, right-click the server in Server Manager and select `Manage As`. Then specify the credentials of an administrator on the remote server.

3. User Account Control by default allows only the local built-in Administrator account on a target workgroup computer to run elevated processes remotely through Server Manager. By default, no other administrator accounts can perform management actions through Server Manager.

This issue is present on remote workgroup servers only. It is not present on remote servers joined to any domain.

To fix this problem, on the remote workgroup server, create a `REG_DWORD` registry entry named `LocalAccountTokenFilterPolicy` in `HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System` and then set its value to `1`. You can do this by typing the following command at an elevated Windows PowerShell prompt:

```
New-ItemProperty -Name LocalAccountTokenFilterPolicy -path HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System -propertyType DWord -value 1
```

4. The firewall rule allowing remote management for the Public network profile allows remote management only from computers located on a computer's local subnet by default.

This issue is present on remote servers only when both of the following are true: The active network profile is `Public` (not `Private` or `Domain`) *and* when the server running Server Manager is on a different subnet.

To fix this problem, in Windows Firewall with Advanced Security, open the firewall rule named Windows Remote Management (HTTP-In) *for the Public profile*. On the Scope tab, in the Remote IP Address list, type the IP address of the server running Server Manager.



EXAM TIP

Remember all four of these issues and their fixes for the exam.

Managing remote management tasks through All Servers

After you add remote servers to Server Manager, you can manage those particular servers from your local Server Manager console as long as they are enabled for remote management. Figure 2-2 shows a menu of management tasks you can perform on remote servers listed in the All Servers – Servers section in Server Manager. (These servers could naturally be added to a custom server group you create as well.) These tasks include adding roles and features, opening Computer Management (to review event logs, for example), opening a Windows PowerShell prompt on the remote server, and configuring NIC teaming.

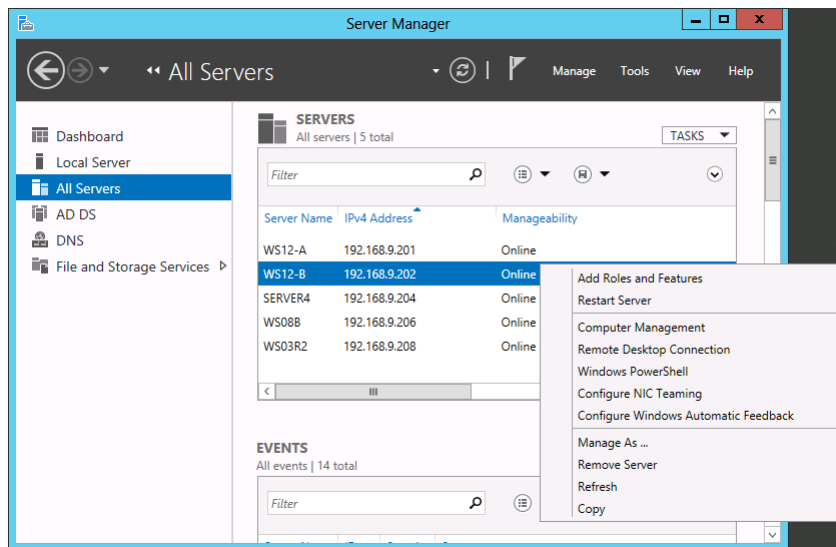


FIGURE 2-2 Remote management tasks in Server Manager

If you right-click a server that is a domain controller, you can access a much larger set of administrative options, including running many diagnostic tools, as shown in Figure 2-3.




FIGURE 2-3 Remote management options on a domain controller



EXAM TIP

None of the utilities shown in the shortcut menu in Figure 2-3 is new to Windows Server 2012 or Windows Server 2012 R2, but you are still likely to see one or more of them on the 70-417 exam. Make sure you review them.

On the All Servers page in Server Manager, you can use the following sections to perform everyday maintenance on the status of your servers:

- **Events** Use this section to check for errors and warnings on your servers without having to open a console on the remote machine.
- **Services** Use this section to check for stopped services.
- **Best Practices Analyzer** Use this section to compare the server configuration to a Best Practices standard.
- **Performance** This section provides CPU and memory usage data of a server over time. To start CPU and memory performance monitoring, you need to right-click a server and select Start Performance Counters, as shown in Figure 2-4.

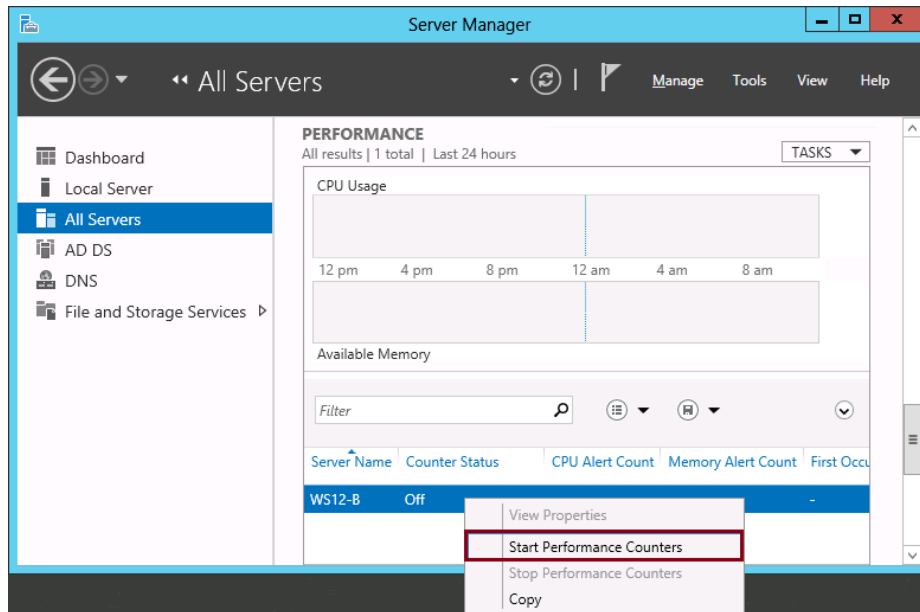


FIGURE 2-4 Monitoring CPU and memory performance

- **Roles And Features** Use this section to verify the installed roles and features on your servers. You can also use this section to remove roles and features.

Remote management in Windows Server 2012 and Windows Server 2012 R2 is enabled and configured by default. You can start remotely managing servers out of the box, as long as they are running Windows Server 2012 or Windows Server 2012 R2, they are in the same domain, and you have the proper administrative privileges. This ease of remote management is a welcome change from both Windows Server 2008 and Windows Server 2008 R2, in which you needed to configure a server to manage it remotely with admin tools.

Unfortunately, the fact that remote management is easy in Windows Server 2012 and Windows Server 2012 R2 doesn't mean that questions about this topic on the exam will be just as easy. It's the exam writers' jobs to develop questions which require a reasonable amount of expertise to solve. For remote management, these questions could likely include scenarios in which you need to re-enable Windows Server 2012 R2 for remote management at the command line, enforce remote management settings through Group Policy, or create proper firewall settings on down-level servers, such as those running Windows Server 2008.

Understanding remote management types: DCOM and WinRM

A brief review of remote management technologies in Windows networks is helpful for exam preparation. Generally speaking, it's important to remember that remote management in Windows networks is an umbrella term. Different remote management tools use different underlying technologies and require different preconfiguration steps.

For both local and remote management, Windows Management Instrumentation (WMI) provides an administrative interface to scripts, programs, and built-in Windows tools. For the purposes of remote management, WMI-based queries and configuration commands are typically passed through either of two protocols: *Distributed Component Object Model (DCOM)* or *Windows Remote Management (WinRM)*. DCOM is an older, proprietary technology for software component communication across networks; WinRM is the Microsoft implementation of an independent standard called WS-Management Protocol.

WMI OVER DCOM

Traditional console-based tools such as Microsoft Management Console (MMC) snap-ins and Computer Management rely on WMI over DCOM when used for remote management. When used remotely, DCOM tools require only that you open certain ports on the firewall of the server you want to manage. If you don't open those ports, you get a message like the one shown in Figure 2-5. This particular error message is helpful in that it informs you exactly which predefined inbound rules you need to enable by using either the Windows Firewall With Advanced Security tool or the `Enable-NetFirewallRule` cmdlet on the remote server:

- COM+ Network Access (DCOM-In)
- All rules in the Remote Event Log Management group

These two sets of rules allow you to connect to most MMC consoles in Windows Server 2012 and Windows Server 2012 R2. Other inbound rules you might need to create are Remote Volume Management (to use Disk Management remotely) and Windows Firewall Remote Management (to use Windows Firewall With Advanced Security remotely).

If you see a question about these DCOM-based remote management tools on the 70-417 exam, it's unlikely to mention DCOM by name. Instead, the question will probably mention Computer Management or the name of another MMC console or snap-in.



EXAM TIP

Remember that if you need to remotely manage a computer running Windows Server 2012 or Windows Server 2012 R2 by using Computer Management, you should enable certain firewall rules either by using Windows Firewall With Advanced Security or the `Enable-NetFirewallRule` cmdlet. (Enabling the remote management property is not sufficient.) This statement is true for both Server Core installations and Server With A GUI installations.

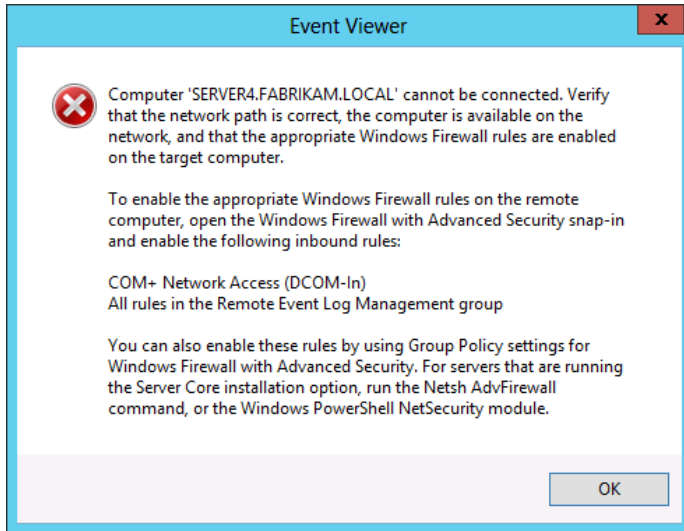


FIGURE 2-5 If you get this message, you've likely tried to use an MMC console for remote management without opening the proper ports

WMI OVER WINRM

The second type of protocol for accessing WMI is WinRM. WinRM isn't new, but within Windows Server there's been a movement toward WinRM-based tools and away from DCOM since Windows Server 2008. As you might remember, WinRM is a Windows service. The most notable tools that use WinRM for remote management are Windows PowerShell, WinRS (Windows Remote Shell), and Server Manager in Windows Server 2012 and Windows Server 2012 R2.



EXAM TIP

Even though the WinRS command is not new, you should remember it for the exam. Use WinRS with the `/r` switch to specify the target computer on which you want to run another command. For example, type `wins /r:myserver ipconfig` to run `Ipconfig` on a server named `Myserver`.

From an exam standpoint, the implications of WinRM being a service is that when a WinRM tool fails, the underlying cause could be that the WinRM service has stopped. (Note that WinRM by default starts automatically in Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, and Windows Server 2012 R2.)

There are also a couple of other important points to remember about WinRM. WinRM tools are firewall-friendly in that they communicate over a single port: either 5985 over HTTP or 5986 over HTTPS. (Yes, you read that correctly: The usual ports 80 and 443 are avoided for security reasons.) Besides requiring a single port to be open for communication, WinRM also requires a WinRM listener to be enabled on the server you want to manage remotely. As

you might remember, both the listener and the port can be configured at once by executing the simple command **winrm quickconfig** at an elevated command prompt on the server you want to manage. (On servers running Windows Server 2012 or Windows Server 2012 R2, however, this step is normally not required because WinRM is enabled by default.)

Re-enabling Windows Server 2012 and Windows Server 2012 R2 for remote management through Server Manager

Remote management is governed in the GUI by the Remote Management property in Server Manager. However, this property enables only WinRM-based remote management, not DCOM-based administration. As a result, by default you can open a Windows PowerShell prompt on a remote server running Windows Server 2012 or Windows Server 2012 R2, or restart a remote server running Windows Server 2012 or Windows Server 2012 R2, because both of these options rely on WinRM. However, you receive an error if you attempt to open Computer Management without opening the needed ports on the remote server, even though this option appears on the shortcut menu of a server that has been added to Server Manager (as shown in Figure 2-2 and Figure 2-3.).

If you discover that you aren't able to use Server Manager to remotely manage a server in the same domain running Windows Server 2012 or Windows Server 2012 R2, it's possible that remote management has been disabled manually. If the server is running either a Server With A GUI installation or Minimal Server Interface, you can re-enable this functionality in the Server Manager interface on that remote server. To do so, perform the following steps:

1. In Server Manager, in the Properties area of the Local Server page, click the hyperlink for the Remote Management property, as shown in Figure 2-6.

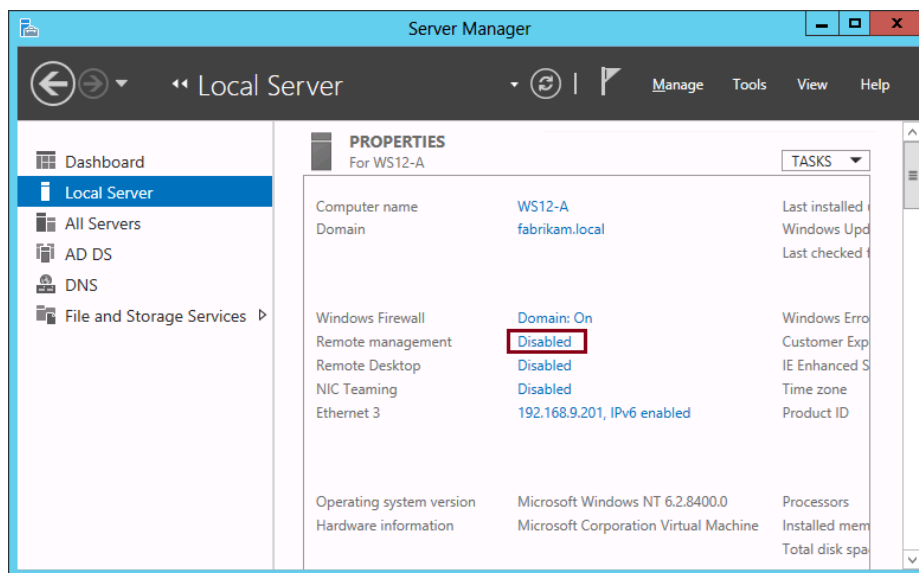


FIGURE 2-6 Re-enabling remote management for Server Manager

2. In the dialog box that opens, select Enable Remote Management Of This Server From Other Computers and then click OK.

Although you need to know how to perform this quick procedure, it might be a bit too straightforward to appear in an exam question. It's easier to imagine a question in which you need to know how to re-enable remote management at the command prompt:

```
Configure-SMRemoting.exe -Enable
```

If you want to disable remote Server Manager management, type the following:

```
Configure-SMRemoting.exe -Disable
```

NOTE You can also use `Configure-SMRemoting -Get` to view the current remote management setting on the server.

Note also that before Windows Server 2012, remote Server Manager management required many open ports, and `Configure-SMRemoting` was a Windows PowerShell script that opened all of those ports. In Windows Server 2012 and Windows Server 2012 R2, remote Server Manager management now relies only on WinRM for most features (such as deploying roles, restarting, and Windows PowerShell) and DCOM for some additional features (such as Computer Management). Consistent with this more efficient remote management method, `Configure-SMRemoting` configures only WinRM in Windows Server 2012 and Windows Server 2012 R2, and it's now the equivalent of the command `Winrm Quickconfig`. As with enabling remote management in the interface or with `Winrm Quickconfig`, if you enable remote management with `Configure-SMRemoting`, you still need to enable the DCOM ports manually at a later point if you want more complete remote management functionality by using MMC consoles.

ENABLING REMOTE MANAGEMENT ON SERVER CORE WITH SCONFIG

`Sconfig` is a text-based configuration tool that is available in the Server Core version of Windows Server. `Sconfig` first appeared in Windows Server 2008 R2, so if you received your MCSA in the first release of Windows 2008, you might have missed this handy utility.

Using `Sconfig` is easy. Just type **Sconfig** at the command prompt in Server Core and you get a menu of self-explanatory configuration options, one of which (choice #4) is to Configure Remote Management, as shown in Figure 2-7.

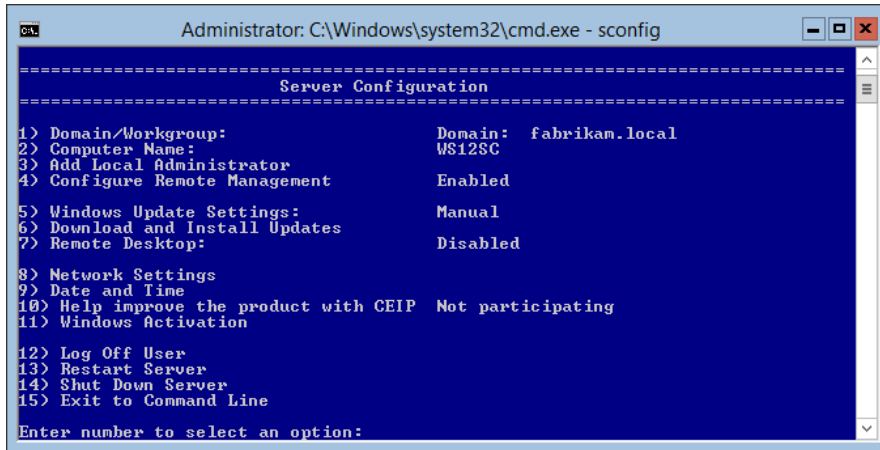


FIGURE 2-7 The Sconfig tool makes it easy to enable remote management in Server Core

Again, as with Winrm Quickconfig and Configure-SMRemoting, enabling remote management in Sconfig configures only WinRM-based remote management.

CONFIGURING REMOTE MANAGEMENT OF EARLIER VERSIONS OF WINDOWS SERVER

Server Manager can be used to remotely administer computers running older versions of Windows Server. For Windows Server 2008 and Windows Server 2008 R2, Server Manager can be used to perform many of the tasks available in Server Manager, but it can't be used to add or remove roles. To manage remote servers that are running Windows Server 2008 or Windows Server 2008 R2, or to manage a remote server running Windows Server 2012 from a server running Windows Server 2012 R2, you must first install the following updates, in the order shown.

1. .NET Framework 4 (if the management computer is running the first release of Windows Server 2012) or .NET Framework 4.5 (if the management computer is running Windows Server 2012 R2).
2. Windows Management Framework 3.0 (if the management computer is running the first release of Windows Server 2012) or Windows Management Framework 4.0 (if the management computer is running Windows Server 2012 R2).

The Windows Management Framework download package updates WMI providers on the target servers. The updated WMI providers let Server Manager collect information about roles and features that are installed on the managed servers.

3. The performance update associated with Knowledge Base (KB) article 2682011 (or a superseding update), which allows Server Manager to collect performance data from Windows Server 2008 and Windows Server 2008 R2. (You don't need to install this update on servers running Windows Server 2012 R2.)

Installing these updates makes these operating systems compatible with Server Manager in 2012. To configure the servers for remote management, run the Winrm Quickconfig

command and (optionally) create the inbound firewall rules needed to support MMC traffic. As an alternative to running the Winrm Quickconfig command, you also can perform the following steps:

1. Open an elevated Windows PowerShell prompt.
2. Type **Set-ExecutionPolicy RemoteSigned**.
3. Type **Configure-SMRemoting.ps1 -force -enable**.

NOTE You cannot use Server Manager to manage a server running Windows Server 2012 R2 from a server running the first release of Windows Server 2012. Also note that for servers running Windows Server 2003 SP2, Server Manager can indicate only whether a server is online or offline.

Using Group Policy to enable remote management

The most efficient way to configure remote management on multiple servers is to use Group Policy. Through Group Policy you can achieve two things: Create WinRM listeners on IP address ranges of your choice and create inbound firewall rules allowing WinRM and DCOM traffic. These steps are described in the following procedure:

1. In a Group Policy Object (GPO) Editor, navigate to Computer Configuration\Administrative Templates\Windows Components\Windows Remote Management and then select WinRM Service. This location within a GPO is shown in Figure 2-8.

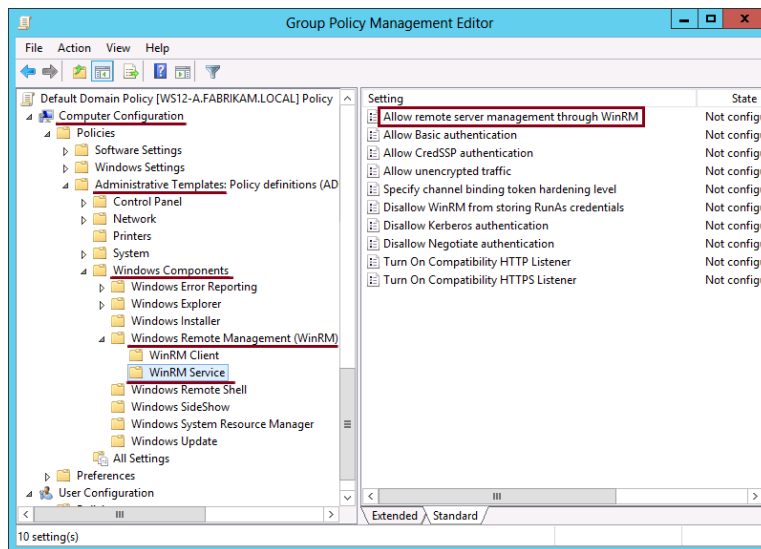


FIGURE 2-8 Configuring WinRM listeners through Group Policy

2. In the details pane, double-click Allow Remote Server Management Through WinRM.
3. In the dialog box that opens, select Enabled (shown in Figure 2-9).

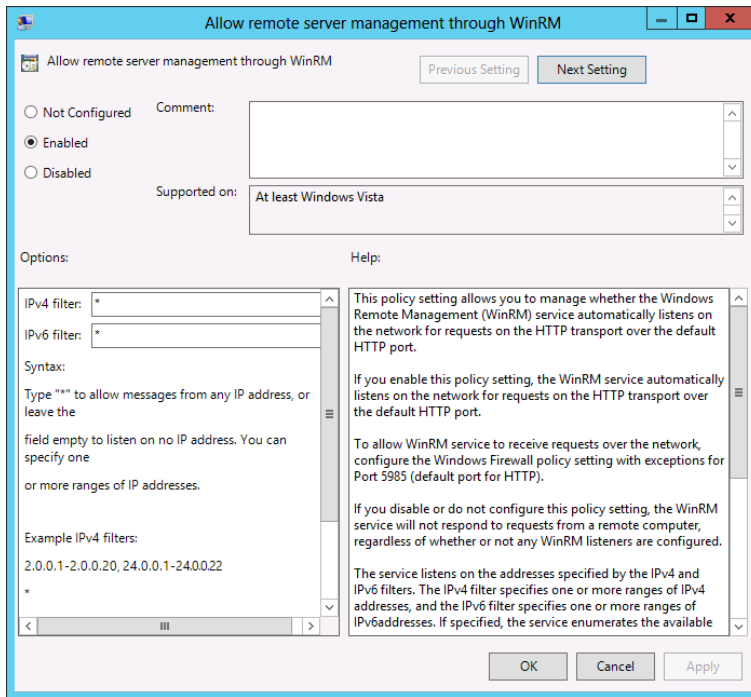


FIGURE 2-9 Configuring WinRM in Group Policy

4. In the IPv4 Filter text box and the IPv6 Filter text box, type the IP address on which you want to allow remote management through WinRM. If you want to enable remote management on all IP addresses, type *.
5. Click OK.
6. In the GPO console tree, navigate to Computer Configuration\Windows Settings\Security Settings\Windows Firewall With Advanced Security\Windows Firewall With Advanced Security.
7. Right-click Inbound Rules and then click New Rule.
8. In the New Inbound Rule Wizard, on the Rule Type page, select Predefined.
9. On the Predefined drop-down menu, select Remote Event Log Management. Click Next.
10. On the Predefined Rules page, click Next to accept the new rules.
11. On the Action page, leave Allow The Connection as the default selection and then click Finish.

12. Repeat steps 7 through 11 to create new inbound rules for the following additional predefined rule types:
- Windows Remote Management
 - COM+ Network Access
 - Remote Volume Management
 - Windows Firewall Remote Management

A GPO configured with these inbound firewall rules is shown in Figure 2-10.

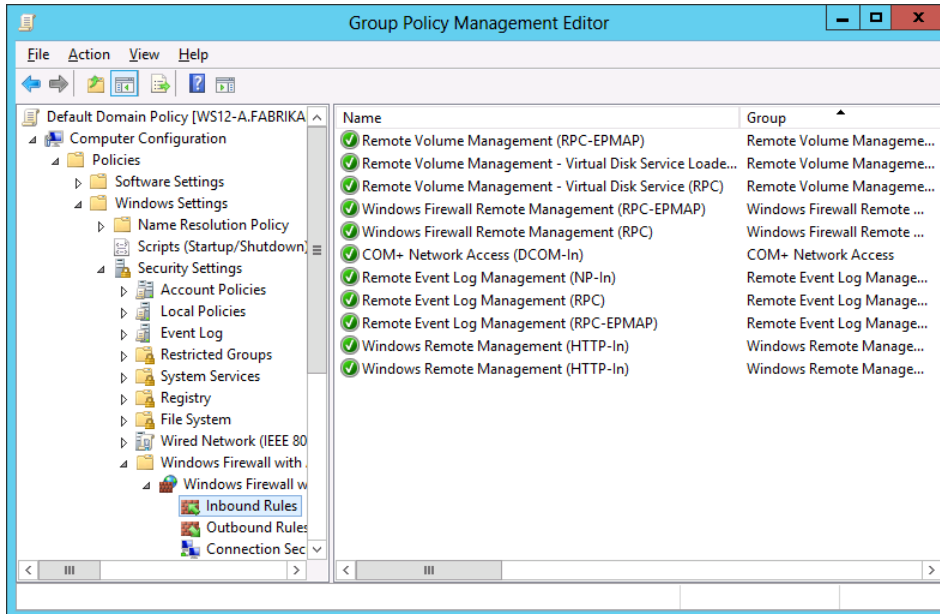


FIGURE 2-10 Firewall rules for remote management

Using Remote Server Administration Tools for Windows 8 and Windows 8.1

To support remote server management from client computers, you can download and install Remote Server Administration Tools for Windows 8 or Remote Server Administration Tools for Windows 8.1. Remote Server Administration Tools includes Server Manager, MMC snap-ins, consoles, Windows PowerShell modules, and some command-line tools for managing roles and features that run on Windows Server 2012 and Windows Server 2012 R2.



EXAM TIP

In its original form on the 70-410 exam, the Configure Server Roles and Features domain includes objectives named “Configure File and Share Access” and “Configure Print and Document Services.” Even though these objectives aren’t officially indicated for the 70-417 exam, you might see a question about these topics that relates to older features which have not changed in Windows Server 2012 or Windows Server 2012 R2. If your knowledge about these topics is rusty, you should review them. For example, you should be able to answer the following questions: How do you configure access-based enumeration (ABE)? What is the Easy Print driver used for? When might you need to configure two printers for the same print device? What is the Volume Shadow Copy Service (VSS) used for?

Objective summary

- Remote management based on the WinRM service is enabled and configured by default in Windows Server 2012 and Windows Server 2012 R2. This is a change from earlier versions of Windows Server.
- Server Manager in Windows Server 2012 and Windows Server 2012 R2 allows you to manage multiple servers, including deploying roles to remove servers and opening a remote Windows PowerShell session. You can use Server Manager to manage remote servers without any additional configuration if the server running Server Manager and the target server are members of the same Active Directory domain. If the servers are not members of the same domain, you will need to add the remote server to the list of trusted hosts for WinRM on the server running Server Manager. Potentially, you will also have to perform other configuration steps, such as creating a registry key to override UAC restrictions. These other steps depend on factors discussed in this chapter.
- MMC consoles rely on DCOM as opposed to WinRM, so you need to enable different inbound firewall rules to use them for remote management. The number of inbound rules you need to enable has been greatly reduced in Windows Server 2012 and Windows Server 2012 R2, compared to previous versions.
- Servers running pre-Windows Server 2012 versions of Windows Server can be managed remotely in Server Manager. To take full advantage of the administrative tasks and information available, you need to update these remote servers with .NET Framework 4 and Windows Management Framework 3.0 (if the management server is running Windows Server 2012) and with .NET Framework 4.5 and Windows Management Framework 4.0 (if the management server is running Windows Server 2012 R2).
- The best way to configure multiple servers for remote management is to use Group Policy. In Group Policy, you can configure WinRM and create all the inbound firewall rules you need to support your remote management.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the “Answers” section at the end of the chapter.

1. A server named SC2012 is running a Server Core installation of Windows Server 2012 R2. You want to manage SC2012 remotely by using Server Manager.

Which of the following will *not* help you achieve your goal?

 - A. Configure-SMRemoting.exe
 - B. Sconfig.exe
 - C. Winrm Quickconfig
 - D. Repadmin.exe
2. Your company network includes 25 servers running either Windows Server 2012 or Windows Server 2008 R2 in a single domain. The servers running Windows Server 2008 R2 have been updated with Windows Management Framework 3.0. You now want to configure all of these servers for remote management through Windows PowerShell. What should you do? (Choose all that apply.)

 - A. Enable the Allow Remote Server Management Through WinRM setting in Group Policy.
 - B. Enable the Allow Remote Shell Access setting in Group Policy.
 - C. Configure an inbound firewall policy rule for COM+ Remote Administration.
 - D. Configure an inbound firewall policy rule for Windows Remote Management.
3. All of your servers are running Windows Server 2012 R2 and are enabled for remote management. You want to remotely manage three of these servers from a server named Admin01. However, when you attempt to manage remote servers through Computer Management on Admin01, you receive error messages.

You create a GPO named Remote Computer Management Firewall Rules and link it to the domain. You now want to use the GPO to create predefined firewall rules in Windows Firewall With Advanced Security. You want to create only the predefined firewall rules that open the ports necessary for remote management through Computer Management. Which of the following predefined inbound rules should you enable? (Choose all that apply.)

 - A. COM+ Network Access
 - B. Remote Event Log Management
 - C. Remote Volume Management
 - D. WMI

4. You work from a management server named Mgmt1 that is joined to the Corp.contoso.com domain. You want to use Server Manager on Mgmt1 to manage a standalone server named Web1.contoso.com while you are logged in with your normal domain account. Note the following information about Web1:
- Web1 is located on the same subnet as Mgmt1.
 - Web1 is not joined to any Active Directory domain.
 - The name of the local administrator account on Web1 is named Web1Admin.
 - Web1 has a single Ethernet connection and its active network profile is Public.
- Which of the following steps are required to allow you to successfully manage Web1 through Server Manager? (Choose all that apply.)
- A. On Web1, open the Windows Remote Management (HTTP-In) firewall rule for the Public profile and add the IP address of Mgmt1 to the list of remote IP addresses on the Scope tab.
 - B. Use the Manage As option in Server Manager to provide credentials of an administrator on Web1.
 - C. Set the value of the LocalAccountTokenFilterPolicy property to 1.
 - D. Add Web1 to the list of TrustedHosts for WinRM on Mgmt1.



Thought experiment

Managing servers remotely at Fabrikam

You are a network administrator at Fabrikam.com, whose network includes 20 servers and 250 clients, all of which belong to the Fabrikam.com domain. Ten servers are running Windows Server 2012 and ten servers are running Windows Server 2008 R2. All clients are running Windows 8.

The IT department is instituting a policy that removes IT personnel from the server room for most day-to-day administration. Administration of all servers from now on will normally be conducted remotely through Server Manager and various MMC consoles. Currently, administration is conducted locally in the server room or through a Remote Desktop connection. The remote management settings on all servers remain at their original defaults.

You can find the answers to these questions in the “Answers” section.

- 1.** All of your servers are located in the server room. Which tool should you use to administer servers remotely from computers running Windows 8?
- 2.** Which inbound rules do you need to create or enable on the servers running Windows Server 2012 to enable remote management through Server Manager?
- 3.** You want to be able to remotely manage the servers running Windows Server 2012 using Computer Management, Disk Management, and Windows Firewall With Advanced Security. Which inbound rules should you enable in Group Policy?
- 4.** You run the Winrm Quickconfig command on your servers running Windows Server 2008 R2. However, you find that you cannot manage these servers remotely by using Server Manager. In addition, some MMC administration tools don't work as they do on the servers running Windows Server 2012. How should you fix this problem?

Answers

This section contains the answers to the Objective Review and the Thought Experiment.

Objective 2.1: Review

1. Correct answer: D

- A. Incorrect:** This command, when used with the `-enable` parameter, enables remote management on Windows Server 2012 and Windows Server 2012 R2.
- B. Incorrect:** This command opens a utility in Server Core that allows you to enable remote management of the local server.
- C. Incorrect:** This command enables remote management on Windows Server 2012 and Windows Server 2012 R2.
- D. Correct:** This tool helps administrators diagnose Active Directory replication problems between domain controllers. It doesn't help you enable remote management.

2. Correct answers: A, D

- A. Correct:** Remote management through Windows PowerShell relies on the WinRM service. You can use this policy setting to configure WinRM listeners on your servers.
- B. Incorrect:** This policy setting does not affect Windows PowerShell.
- C. Incorrect:** This firewall rule does not open any of the ports needed by Windows PowerShell.
- D. Correct:** This firewall rule opens the port needed for WinRM-based communication and is required for Windows PowerShell remoting.

3. Correct answers: A, B, C

- A. Correct:** This predefined rule allows you to connect to a remote computer through Computer Management and use a few system tools, such as Shared Folders and Local Users And Groups.
- B. Correct:** This predefined rule group allows you to manage computers remotely in Computer Management through the Event Viewer system tool and the Task Scheduler system tool.
- C. Correct:** This predefined rule group allows you to use Disk Management remotely in Computer Management.
- D. Incorrect:** This predefined rule doesn't enable you to use any tools in Computer Management.

4. Correct answers: B, C, D

- A. Incorrect:** This step is necessary only when the active network profile on the target server is Public *and* the server running Server Manager is found on a subnet that is different from the target server's subnet. In this case, the Public profile is active on Web1, but Mgmt1 and Web1 are located on the same subnet.
- B. Correct:** The question states that you want to use your normal domain account on Mgmt1, so to manage Web1 remotely, you will need to specify the credentials of an administrator on that machine.
- C. Correct:** To overcome restrictions related to User Account Control in workgroup environments, you need to create this REG_DWORD entry in HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion on the target server and set its value to 1.
- D. Correct:** Server Manager relies on WinRM. To use WinRM to manage a remote server outside of a domain environment, you need to add the target computer to the list of TrustedHosts on the managing computer. In this case, you need to add Web1 to the list of TrustedHosts on Mgmt1.

Thought experiment

1. You should use Remote Server Administration Tools for Windows 8.
2. You don't need to create or enable any inbound rules. Servers running Windows Server 2012 have the required rules enabled by default.
3. You should enable COM+ Network Access (DCOM-In), Remote Event Log Management, Remote Volume Management, and Windows Firewall Remote Management.
4. You should install .NET Framework 4, Windows Management Framework 3.0, and the update associated with KB article 2682011.

This page intentionally left blank

Configure Hyper-V

If you were to name the one server role that has seen the most changes between Windows Server 2008 R2 and Windows Server 2012, you'd have to say Hyper-V. To begin with, Hyper-V now has its own Windows PowerShell module, so the role is completely manageable at the Windows PowerShell prompt. Beyond this new command-line manageability, there are new improvements in memory configuration, storage, Resource Metering, security, extensibility, and other areas (such as fault tolerance) that aren't even covered in this domain.

In Windows Server 2012 R2, the changes from the first release of Windows Server 2012 are more incremental but still significant. One highly visible change is that "snapshots" are now called "checkpoints," so don't be surprised when you see that word on the exam. Other important new features in Windows Server 2012 R2 Hyper-V include enhanced session mode, which improves the usability of VMs, and generation 2 VMs, which offer improved installation and startup times.

The good news is that none of these new features is particularly difficult to understand, at least at the level they will be tested on for the 70-417 exam. Your studying efforts in this area should therefore pay off well.

Objectives in this chapter:

- Objective 3.1: Create and configure virtual machine settings
- Objective 3.2: Create and configure virtual machine storage
- Objective 3.3: Create and configure virtual networks

Objective 3.1: Create and configure virtual machine settings

Of the features mentioned by Microsoft in the description of this exam objective, three are at least partially new to Windows Server 2012 (Dynamic Memory, Smart Paging, and Resource Metering) and two are new to Windows Server 2012 R2 (generation 2 virtual machines and enhanced session mode). A sixth topic, RemoteFX, was introduced in Windows Server 2008 R2 but has only recently been added to the objectives.

Beyond learning about these six topics, though, you should also know that all new and old settings in Hyper-V can now be configured at the Windows PowerShell prompt. From

that perspective, every virtual machine setting is new; any configuration option could be covered on the exam. You should be sure, therefore, to supplement your study of the new features in Windows Server 2012 and Windows Server 2012 R2 Hyper-V with a review of the new cmdlets related to virtual machine (VM) configuration in the Hyper-V module.

This section covers the following topics:

- Hyper-V module in Windows PowerShell
- Generation 1 and generation 2 virtual machines
- Enhanced session mode
- Dynamic Memory
- Smart Paging
- Resource Metering
- Non-uniform memory access (NUMA) topology
- RemoteFX

Hyper-V Module in Windows PowerShell

As you've already learned, Windows PowerShell in Windows Server 2012 and Windows Server 2012 R2 includes a new module called Hyper-V that provides a command-line administration interface for almost all VM settings. It's uncertain how many cmdlets will appear on the 70-417 exam, and there are too many of them (more than 150) to document here.

Instead, you can use `Get-Command` to review the names of these cmdlets so that you can at least recognize the most important ones. You can sort the output by the cmdlet nouns to make it easier to understand. (The noun portion of a cmdlet represents the object that is configured.)

For example, to see a list of all cmdlets in the module and group them by cmdlet noun, type the following:

```
Get-Command -Module Hyper-V | Sort Noun,Verb
```

If you want to see cmdlets that contain the string `*VM*` (and are likely to relate specifically to VM management and configuration), type the following:

```
Get-Command *VM* | Sort Noun,Verb
```

To further filter your results, you can use the wildcard character twice or more, as in the following example:

```
Get-Command *VM*adapter* | Sort Noun,Verb
```

You can then use `Update-Help` and `Get-Help`, optionally with the `-Examples` or `-Full` option, to get the latest documentation about any particular cmdlet that interests you.

Generation 1 and generation 2 virtual machines

Beginning with Windows Server 2012 R2, the New Virtual Machine Wizard now includes a Specify Generation page, shown in Figure 3-1.

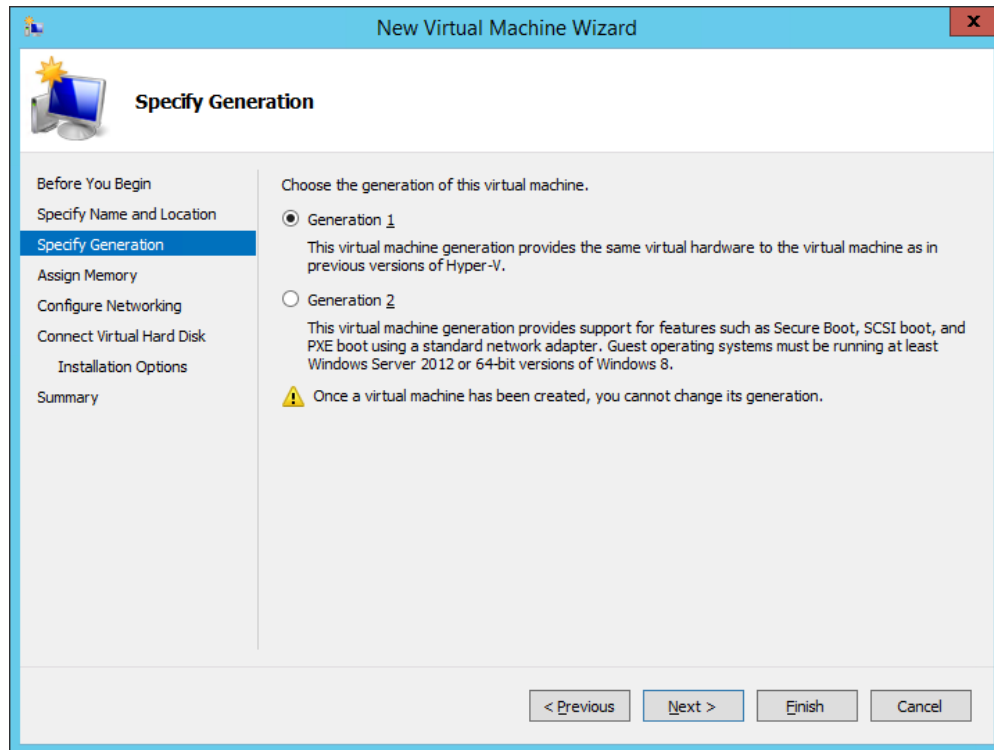


FIGURE 3-1 Choosing the generation of a new VM

The concept of a VM generation is new to Windows Server 2012 R2, and you'll likely see one or more questions about this topic on the 70-417 exam. Generation 1 VMs, as they are now called, are the familiar VMs that have existed in all versions of Hyper-V since it was first introduced. Generation 2 VMs are a new option available only when the host system is running Windows Server 2012 R2.

Here are the changes in generation 2 VMs that you need to understand:

- **Removal of legacy emulated hardware devices** Generation 1 VMs emulate a set of legacy hardware devices, including two IDE controllers, two COM ports, and a floppy disk drive. These emulated devices have been removed from generation 2 VMs. The advantage of removing support for emulated devices is faster boot times (by about 20 percent) and faster installations (by about 50 percent).

Figure 3-2 and Figure 3-3 show the difference in the number of default hardware devices in generation 1 and 2 virtual machines. Each figure displays the complete list of hardware devices in a default VM of each generation.

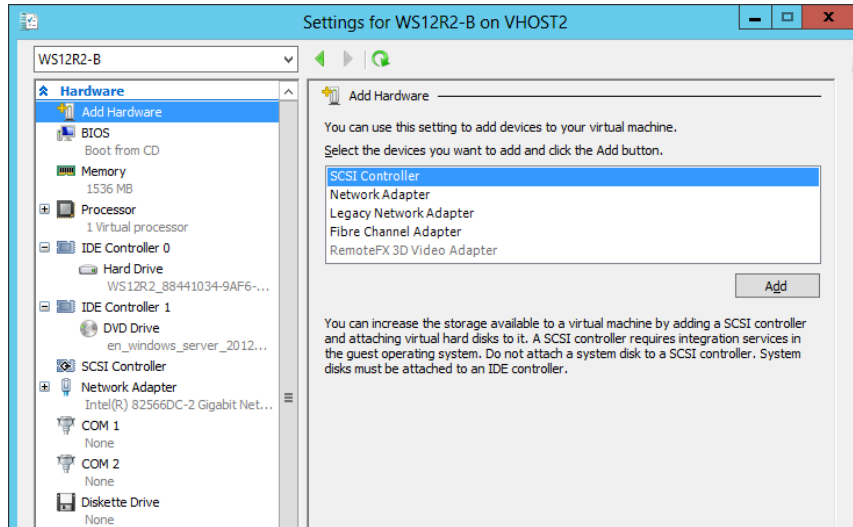


FIGURE 3-2 Default hardware devices for a generation 1 VM

The generation 2 VM lacks both IDE controllers, COM ports, and the diskette drive.




FIGURE 3-3 Default hardware devices for a generation 2 VM

- UEFI vs. BIOS** Generation 1 VMs use Basic Input Output System (BIOS) as a firmware interface to initiate the boot process and begin loading the operating system. Generation 2 VMs use Unified Extensible Firmware Interface (UEFI) for the same purpose. The main advantage of UEFI is that it allows Secure Boot, an option that you can enable on the Firmware page of a VM's settings, as shown in Figure 3-4. Secure Boot ensures that no malicious code is installed beneath the operating system and that the UEFI has not been altered from an approved version.




FIGURE 3-4 Enabling Secure Boot on a generation 2 VM

- **SCSI boot** In generation 1 VMs, you can boot from IDE devices only. In generation 2 VMs, there are no IDE controllers, and you can boot from SCSI devices, including ISO files.
- **PXE boot-compatible network adapters** In generation 1 VMs, only legacy network adapters are PXE-boot compatible. In generation 2 VMs, there are no more legacy network adapters, and normal network adapters are PXE-boot compatible.
- **Limited operating system support** Generation 2 VMs support only the following guest operating systems:
 - Windows 8 (64-bit)
 - Windows Server 2012
 - Windows 8.1 (64-bit)
 - Windows Server 2012 R2
- **No RemoteFX support** Generation 2 VMs do not support RemoteFX.
- **No Physical CD/DVD support** Generation 2 VMs do not support physical CDs or DVDs. You must use ISO files instead.
- **No VHD support** Generation 2 VMs do not support VHD files. You must use VHDX files. (The VM's operating system must also be installed originally on a VHDX file, not on a VHDX that has been converted from VHD.)
- **No virtual floppy disk support** Generation 2 VMs do not support VFD files.



EXAM TIP

To perform a network boot on a VM from a PXE-compatible adapter, select BIOS in a generation 1 VM's settings, and Firmware in a generation 2 VM's settings. Then adjust the Startup Order (for BIOS) or the Boot Order (for Firmware) so that the desired adapter is listed at the top.

Enhanced session mode

Enhanced session mode is a new feature in Windows Server 2012 R2 that improves the connectivity between a guest operating system and the host operating system. In short, enhanced session mode provides the VM connection window with most of the benefits of a Remote Desktop connection window, including the ability to copy and paste between the host operating system desktop and the guest VM.

Enhanced session mode isn't enabled by default, and it's available only when the guest is running Windows 8.1 or Windows Server 2012 R2. To enable enhanced session mode, you have to enable two options in Hyper-V Settings of the host computer: First, as shown in Figure 3-5, navigate to Enhanced Session Mode Policy in the Server menu on the left and then select Allow Enhanced Session Mode in the right pane. Second, as shown in Figure 3-6, navigate to Enhanced Session Mode in the User menu on the left and then select Use Enhanced Session Mode.

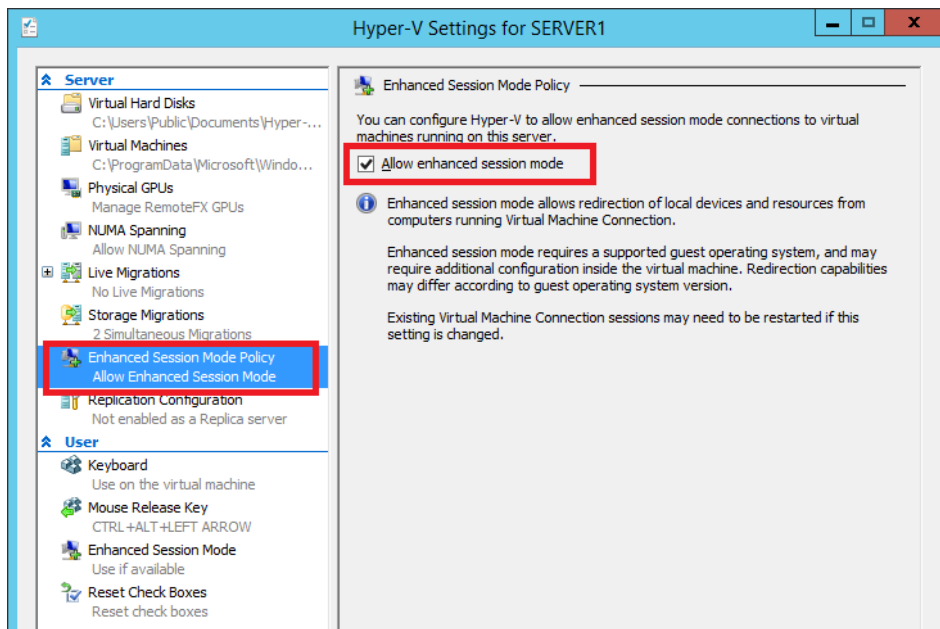


FIGURE 3-5 Step one in enabling enhanced session mode

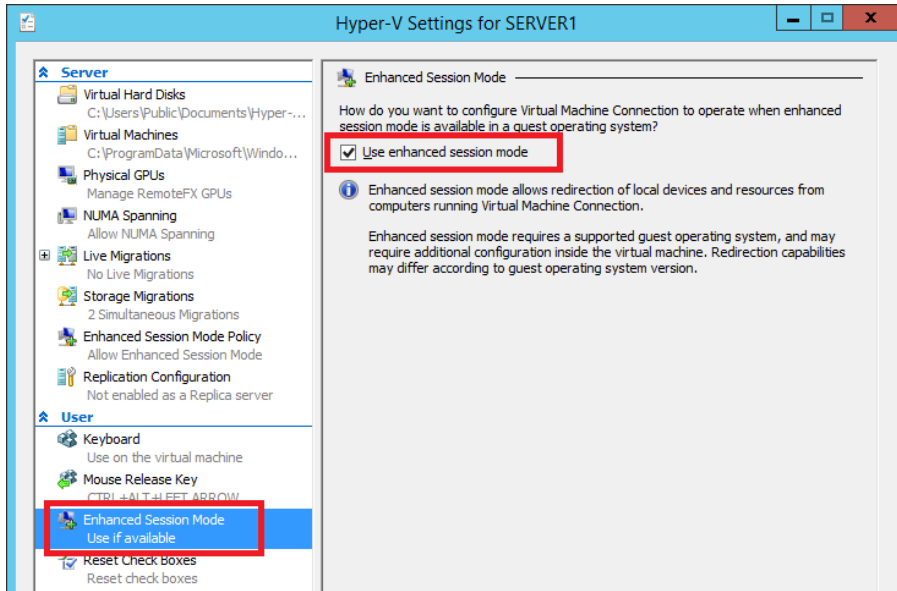


FIGURE 3-6 Step two in enabling enhanced session mode

After you enable these two options, your next VM connection to a Hyper-V guest running Windows 8.1 or Windows Server 2012 R2 will open the window shown in Figure 3-7.

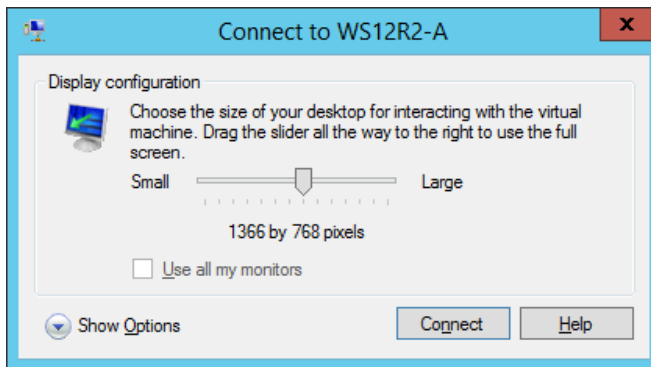


FIGURE 3-7 Enhanced session mode options

Clicking the down arrow next to Show Options reveals a Local Resources tab, which is similar to the tab of the same name available in a Remote Desktop Connection window. By default, remote audio on the VM is configured to be played on the host operating system, and the local printers and Clipboard data are shared through the VM connection. (The Clipboard allows copying and pasting between the host and guest machines.) Smart card readers in the host operating system are also redirected to the VM by default.

You can also use the Local Resources tab to configure the following resources in the VM connection window:

- **Drives** Makes local drives on the host operating system available from within the VM. Drives can be selected individually
- **Other Supported Plug and Play (PnP) Devices** Redirects PnP devices on the host, such as USB drives to the VM

Dynamic Memory

Dynamic Memory was introduced in Windows Server 2008 R2 Service Pack 1. Just one new configuration setting (Minimum RAM) has been added in Windows Server 2012 and Windows Server 2012 R2, but you should be prepared to be tested on any or all of the feature's settings on the 70-417 exam.

If you haven't had the chance to learn about this feature, remember the following point: Dynamic Memory pools the available RAM on a Hyper-V host for all running VMs for which Dynamic Memory is enabled. Using this pool, Dynamic Memory automatically modifies on the fly the amount of RAM assigned to each running VM as the need increases or decreases. The biggest benefit of Dynamic Memory is that it allows you to use your RAM resources in a highly efficient manner, dramatically increasing the number of VMs you can run on that Hyper-V host. (Marketing materials talk about the benefit Dynamic Memory offers in "improving consolidation ratios" on your virtualization servers. It's good to know that phrase because you might find it on the exam.)

The second most important concept you need to remember about Dynamic Memory is that starting a VM often requires more memory than does running the VM after it starts, and dynamic RAM assignment in Windows Server 2012 naturally mirrors these changing needs. If, for example, you have 6 GB of RAM on a server and try to start 10 VMs at once, you might get an error message regardless of whether Dynamic Memory is enabled. However, only if Dynamic Memory is enabled might you be able to get them all up and running if you start them one at a time. The prototypical example that illustrates low memory usage after startup is with virtual desktop infrastructure (VDI), where you might have a pool of unused virtual machines available in case several people happen to need a desktop all at once. (If you see a scenario on the exam about VDI and desktop pools, expect Dynamic Memory to play a part in the solution somehow.)

Now let's take a look at Dynamic Memory settings. They appear where you'd expect, which is in the Memory section of a VM's settings in Hyper-V Manager, as shown in Figure 3-8. You also can enable and configure Dynamic Memory with Windows PowerShell by using the Set-VM cmdlet, which can be used to configure the various properties of a VM. Note that you can enable or disable Dynamic Memory only when the VM is in a stopped state. (Dynamic Memory does *not* mean you can manually adjust RAM settings while a VM is running.)

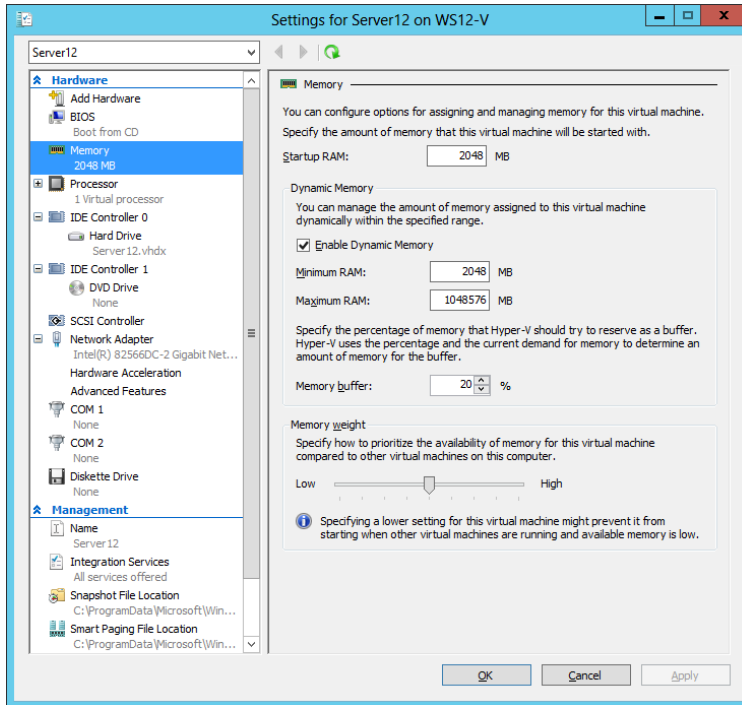


FIGURE 3-8 Configuring Dynamic Memory settings

The settings here affect *how* memory is assigned dynamically. You need to understand the implications of these settings on behavior and performance.

- **Startup RAM** This value existed before Windows Server 2012, but it used to have a slightly different meaning. Before Windows Server 2012, the Startup RAM was both the amount of RAM used at startup and the minimum amount of RAM ever assigned to the VM.

In Windows Server 2012 and Windows Server 2012 R2, the Startup RAM setting is now only the amount of RAM assigned at startup and is no longer the minimum RAM. If a running VM uses less RAM after startup, some of that RAM can now be reclaimed by other running VMs.

Here is another important point to remember about Startup RAM: The more RAM you assign to a VM when it starts up, the faster it will be able to start up (of course). But don't forget the flip side: If you set this level too high, you might temporarily (during startup) deprive other VMs of the RAM they need to perform at an acceptable level.

- **Minimum RAM** This is the only new setting that has been added in Windows Server 2012 and Windows Server 2012 R2, so make sure you understand it. If you have enabled Dynamic Memory on a VM running on a Windows Server 2012 or Windows Server 2012 R2 host, by default this value is the same as the Startup RAM value. However, you can lower Minimum RAM to allow the amount of RAM allocated to the VM to decrease after startup.

Why would you want to manually lower the Minimum RAM level? One reason is that by allowing unused physical memory of a running VM to be reclaimed, you can make sure that physical memory is available to other VMs that might need it. On the other hand, by keeping the value higher, you can ensure that enough is available to the same VM when it restarts.

- **Maximum RAM** This is the maximum amount of memory that can ever be dynamically assigned to the VM. There's always a balancing act here. If you don't set this value high enough, the VM's performance could suffer. However, for a RAM-intensive workload, setting this value too high could deprive other VMs of needed RAM.
- **Memory Buffer** This is the preferred amount of extra RAM (defined as a percentage) that is assigned to the system beyond what is determined to be needed to run the active workload at any given point. The default is set at 20 percent. You don't normally have to change this setting, but if memory usage spikes intermittently on a VM, you might want to increase this percentage to help ensure that enough RAM is available when needed.
- **Memory Weight** This parameter determines how available memory on the host is allocated among the different VMs running on the host. If you want to prioritize the performance and memory allocation of a given VM relative to other VMs, you would raise the memory weight setting on that VM.

NOTE To review and configure Dynamic Memory settings in Windows PowerShell, use `Get-VMMemory` and `Set-VMMemory`.

Smart Paging

What if, with Dynamic Memory enabled, you have just enough RAM to start your VMs but not enough to restart a particular VM once they are all up and running? Maybe, for example, you used the last 256 MB of available RAM to start a VM running Microsoft Windows XP, and now you can't restart a VM running Windows 8, which requires 512 MB of RAM to start. To prevent this kind of scenario from happening, Hyper-V in Windows Server 2012 and Windows Server 2012 R2 introduces a new feature called Smart Paging. *Smart Paging* allows a VM that's being restarted to use disk resources temporarily on the host as a source for any additional memory needed to restart a VM. Then, once the VM has started successfully and its memory requirements decrease, Smart Paging releases the disk resources. The downside of Smart Paging, as you probably have guessed, is that performance is compromised. VMs restart, but slowly, with Smart Paging.

To minimize the performance impact of Smart Paging, Hyper-V uses it only when all of the following are true:

- The VM is being restarted.
- There is no available physical memory.
- No memory can be reclaimed from other VMs running on the host.

Smart Paging is *not* used in the following cases:

- A VM is being started from an “off state” (instead of a restart).
- Oversubscribing memory for a running VM is required.
- A VM is failing over in Hyper-V clusters.

Smart Paging is a new feature that is specifically mentioned in the objectives for the 70-417 exam, so don't be surprised if it appears as an important element in a test question. With this in mind, be aware that the only configuration option for Smart Paging relates to the storage location for the Smart Paging file. Why would the location of the Smart Paging file ever matter in a test question? Well, if the disk on which the Smart Paging file is stored nears its capacity, there might not be enough disk space to allow a VM to restart. If this happens, the way to solve the problem would be to move the Smart Paging file to a disk with more space. (That's assuming you can't add more RAM to the host server, of course.)



EXAM TIP

If you create a checkpoint (formerly called a snapshot) of a live VM, the checkpoint will save the running memory. If you create a checkpoint of a stopped VM, no data in RAM needs to be saved. Therefore, if you want to reduce the size of your checkpoint file, shut down the VM before creating the checkpoint.

Resource Metering

Resource Metering is a new feature of Windows Server 2012 and Windows Server 2012 R2 that is designed to make it easy to build tools that measure VM usage of CPU, memory, disk space, and network. This feature was primarily designed for hosting VMs for a customer. In such a scenario, you need to know how much of your computing resources are used so that you can charge the customer accordingly.

You can use Resource Metering in Windows Server 2012 and Windows Server 2012 R2 to collect and report on historical resource usage of the following seven metrics:

- Average CPU usage by a VM
- Average physical memory usage by a VM
- Minimum physical memory usage by a VM
- Maximum physical memory usage by a VM
- Maximum amount of disk space allocated to a VM
- Total incoming network traffic for a virtual network adapter
- Total outgoing network traffic for a virtual network adapter

You can view this functionality in Windows PowerShell even though it is intended to be used primarily with additional tools.

To enable Resource Metering on a VM, use the `Enable-VMResourceMetering` cmdlet on the host server. For example, to enable Resource Metering on a VM named `VSrv1`, type the following at a Windows PowerShell prompt:

```
Enable-VMResourceMetering -VMName VSrv1
```

At this point, the Resource Metering counters start running. To view all Resource Metering statistics on the VM since you ran the last command, use the Measure-VM cmdlet. For example, type the following to display the Resource Metering data on VSrv1 for all seven metrics:

```
Measure-VM -VMName VSrv1
```

Alternatively, you could save the usage statistics into a report with this command:

```
$UtilizationReport = Get-VM VSrv1 | Measure-VM
```

You could then display the contents of the report at a later time with the following command:

```
Write-Output $UtilizationReport
```

To reset the counters to start counting usage again from zero, you use the following command:

```
Reset-VMResourceMetering -VMName VSrv1
```

To stop the counters from running on VSrv1, type the following:

```
Disable-VMResourceMetering -VMName VSrv1
```

These metrics can be collected even when the VMs are moved between hosts using live migration or when their storage is moved using storage migration.

For the 70-417 exam, what's most important to remember about Resource Metering is that it allows you to measure CPU, memory, disk, and network usage on a particular VM. You should also know the general steps required to configure Resource Metering, but you won't have to know the specific syntax used in Windows PowerShell cmdlets.

NOTE If you want to measure Internet traffic as opposed to network traffic in general, you can use network metering port access control lists (ACLs), which are described later in this chapter.



EXAM TIP

Remember that Resource Metering doesn't let you measure *current* resource usage. You can, however, use Task Manager to view current CPU and memory usage for individual VMs. To do so, open a Virtual Machine Connection to each VM, and then view the Processes tab in Task Manager. Each VM will appear as a separate instance of Virtual Machine Connection along with the current CPU and memory usage for that VM.

You can also use counters in Performance Monitor on the host server to track VM resource usage over time. For example, to measure CPU usage in one or more particular VMs, use the Hyper-V Hypervisor Virtual Processor counter set. To measure Dynamic RAM usage in on or more particular VMs, use the Hyper-V Dynamic Memory VM counter set.

MORE INFO For an overview of Resource Metering in Windows Server 2012, see the topic “Hyper-V Resource Metering Overview” in the TechNet Library at <http://technet.microsoft.com/en-us/library/hh831661.aspx>. Also search for the specific Windows PowerShell cmdlets on <http://technet.microsoft.com>.

Non-uniform memory access (NUMA)

Non-uniform memory access (NUMA) is a new configuration node beneath the Processor node in a VM’s settings. NUMA is a technology that improves system scalability by optimizing memory and memory bus usage in multi-processor systems. In Windows Server 2012 and Windows Server 2012 R2, VMs are NUMA-aware, which means that multi-processor VMs can access memory resources in a more optimal and scalable way. Generally speaking, you don’t need to change the default settings in the NUMA topology configuration area because they are automatically configured correctly based on the host server’s hardware. On rare occasions, however, it might be necessary to modify these settings if you have moved a VM between two physical hosts with different NUMA topologies. Configuring these settings is beyond the scope of the 70-417 exam, but you should know that the Use Hardware Topology button resets NUMA settings to the default settings.

EXAM TIP

Be sure to review VM settings that have not changed since Windows Server 2008. For example, you should know that Integration Services enable VM features such as time synchronization, host-backup awareness, and system shutdown awareness. Also review VM settings such as Resource Control, which allows you to prioritize CPU resources for certain VMs.

RemoteFX

RemoteFX is a set of technologies that improves video rendering, graphics, and overall user experience over the RDP protocol. RemoteFX can work only if a RemoteFX-compatible graphics processing unit (GPU) is available on the remote server to which clients are connecting over RDP.

RemoteFX can be used with Hyper-V. In this case, clients connect to remote VMs over RDP. All VMs on a physical host can share the GPU of that host, and each VM is configured with a virtual GPU (vGPU) that points to the physical GPU.

Here are the requirements for running RemoteFX with Hyper-V:

- Windows Server 2008 R2 SP1 or later
- DX11 vGPU with WDDM v1.2 driver
- SLAT-capable processor
- Remote Desktop Virtualization Host component of the Remote Desktop Services role must be installed (to enable RemoteFX vGPU)

- GPU or GPUs must be enabled for use with RemoteFX in Hyper-V Settings
- VMs must have the “RemoteFX 3D Video Adapter” hardware component added
- VMs must be generation 1

MORE INFO For more information about configuring RemoteFX in Windows Server 2012, search for “RemoteFX vGPU Setup and Configuration Guide for Windows Server 2012” on <http://technet.microsoft.com>.

Objective summary

- In Windows Server 2012 and Windows Server 2012 R2, almost all VM settings can be configured in Windows PowerShell.
- Windows Server 2012 R2 introduces the option to create generation 2 virtual machines. Generation 2 virtual machines drop support for legacy hardware devices, but they boot faster and perform operating system installations faster. Generation 2 VMs also allow the option for Secure Boot, which ensures that no malicious software is installed beneath the operating system.
- Enhanced session mode is a new feature in Windows Server 2012 R2. It provides a VM connection with many of the benefits of a Remote Desktop connection, including the ability to share features with the host operating system such as printers, Clipboard data, and drives.
- Dynamic Memory pools all the memory available on a host server for all VMs hosted on that server. Because computers tend to use more memory when they are starting than when they are running, Dynamic Memory allows you to use available RAM much more efficiently.
- Important Dynamic Memory settings include Startup RAM, Minimum RAM, and Maximum RAM.
- Smart Paging allows VMs to use virtual (paged) memory to complete a restart operation when insufficient physical memory is available.
- With the Resource Metering feature in Windows Server 2012 and Windows Server 2012 R2, you can use the `Enable-VMResourceMetering` cmdlet to start metering the CPU, memory, disk, and network usage of a VM. To display usage statistics, use the `Measure-VM` cmdlet. To reset usage counters to zero, use `Reset-VMResourceMetering`. To disable Resource Metering, use `Disable-VMResourceMetering`.
- RemoteFX improves graphics over RDP and can be used with Hyper-V to improve Remote Desktop connections to individual VMs. The physical host requires a compatible GPU, and you have to select this GPU in Hyper-V Settings and enable it for RemoteFX. You then need to add a RemoteFX 3D Video Adapter in each chosen VM. RemoteFX is not compatible with generation 2 VMs.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the “Answers” section at the end of the chapter.

1. A server named HYPV1 is running Windows Server 2012 R2 and has been configured with the Hyper-V role. HYPV1 has 20 GB of RAM and is hosting 12 VMs. All VMs are running Windows Server 2012 R2 and have Dynamic Memory enabled.

One of the VMs hosted on HYPV1 is named VM1. VM1 is hosting a web application. VM1 averages five concurrent user connections to this web application and its performance is considered acceptable by users.

VM1 has the following memory settings:

- Startup Memory: 1024 MB
- Minimum Memory: 384 MB
- Maximum Memory: 4096 MB
- Memory Weight: Medium

You need to perform a scheduled restart of VM1 once per week. You have noticed during past scheduled restarts of VM1 that restarts have sometimes occurred only with the help of Smart Paging and have required several minutes to complete. You want to minimize downtime associated with restarting VM1 and reduce the likelihood that any restart operation will rely on Smart Paging. What should you do?

- A.** Increase the Startup Memory on VM1.
 - B.** Increase the Minimum Memory setting on VM1.
 - C.** Decrease the Maximum Memory on other VMs.
 - D.** Change the Memory Weight setting on VM1 to High.
2. A server named HYPV2 is running Windows Server 2012 R2 and has been configured with the Hyper-V role. HYPV2 has 16 GB of RAM and is hosting 10 VMs. All VMs are running Windows Server 2012 R2 and have Dynamic Memory enabled.
One of the VMs on HYPV2 is named VM2. VM2 hosts a little-used application that is used for testing only and is not used for any other purposes. You attempt to restart VM2 but receive an error message indicating that there is insufficient memory to perform the operation.
You want to restart VM2 successfully. What should you do? (Choose all that apply.)
 - A.** Increase the Startup Memory setting on VM2.
 - B.** Decrease the Maximum Memory on other VMs.
 - C.** Increase the Memory Buffer % setting on VM2.
 - D.** Move the Smart Paging file to a disk with more space.

3. A server named HYPV3 is running Windows Server 2012 R2 and has been configured with the Hyper-V role. HYPV3 hosts a VM named VM3. You have been measuring the CPU, memory, network, and disk space usage of VM3 for the past 24 hours. You would now like to display the collected usage data at the Windows PowerShell prompt. Which of the following commands should you type at an elevated Windows PowerShell prompt?
- A. `Enable-VMResourceMetering -VMName VM3`
 - B. `Disable-VMResourceMetering -VMName VM3`
 - C. `Measure-VM -VMName VM3`
 - D. `$UtilizationReport = Get-VM VSrv1 | Measure-VM`

Objective 3.2: Create and configure virtual machine storage

There are three topics in this objective that are most likely to be tested: VHDX, virtual Fibre Channel, and storage Quality of Service (QoS). Of these three, VHDX is the one feature you're pretty much guaranteed to see on the 70-417 exam. Fortunately, though, all three topics are easy to understand.

This section covers the following topics:

- New VHDX disk format
- Virtual Fibre Channel adapter
- Storage Quality of Service (QoS)

New VHDX disk format

Virtual hard disk (VHD) files have a size limit of 2 TB, which can prevent you from virtualizing some workloads such as extra-large databases. To fix this problem, Windows Server 2012 and Windows Server 2012 R2 introduce a new VHDX file format, which has a 64 TB limit.

Size is the biggest advantage of the VHDX, so if it appears in a test question, it will most likely be in the context of a scenario in which you need to support files that are larger than 2 TB. What is the disadvantage of VHDX? Backward compatibility. If you need to migrate storage to servers running Windows Server 2008 R2 or earlier, use VHD. Also note that the larger

size of VHDX applies only to non-boot volumes. VHDX boot disks are also limited to 2 TB because of limitations found in the legacy AMI BIOS used in Hyper-V virtual machines.

Remember that VHDX is the default selection for a new VHD file, as shown in Figure 3-9, but you can opt to create a VHD just as easily.

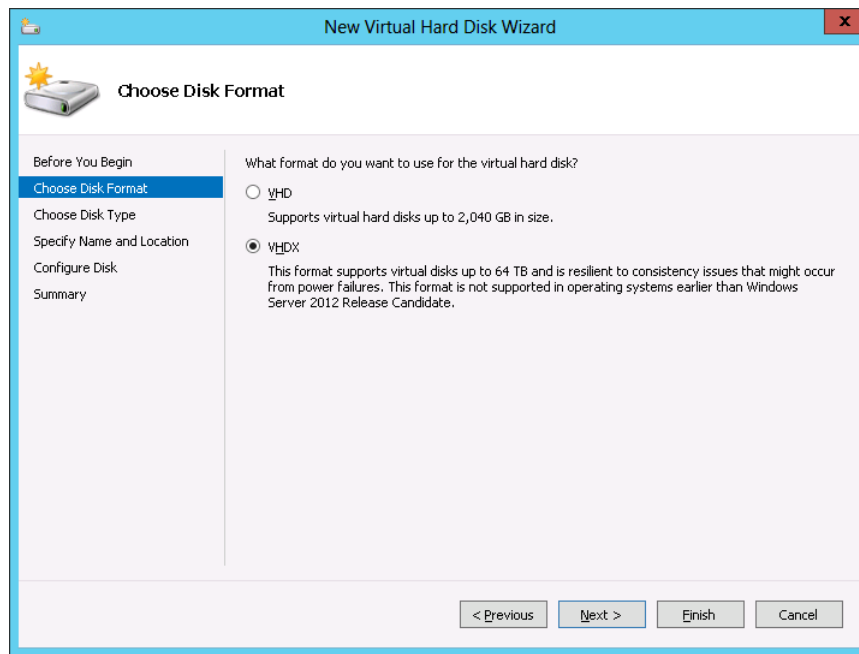


FIGURE 3-9 Creating a new VHDX

You can also convert a disk from a VHDX to a VHD and vice versa, as long as the disk isn't bigger than 2 TB. To do so, just select the virtual disk in the VM settings and click Edit, as shown in Figure 3-10.



EXAM TIP

Aside from Hyper-V Manager, you can also use Computer Management or the New-VHD cmdlet to create a new VHD or VHDX. (Note that New-VirtualDisk is different: That cmdlet is used to create a new virtual disk in a specific storage pool.) To convert a virtual hard disk between the VHD and VHDX formats in Windows PowerShell, use the Convert-VHD cmdlet.

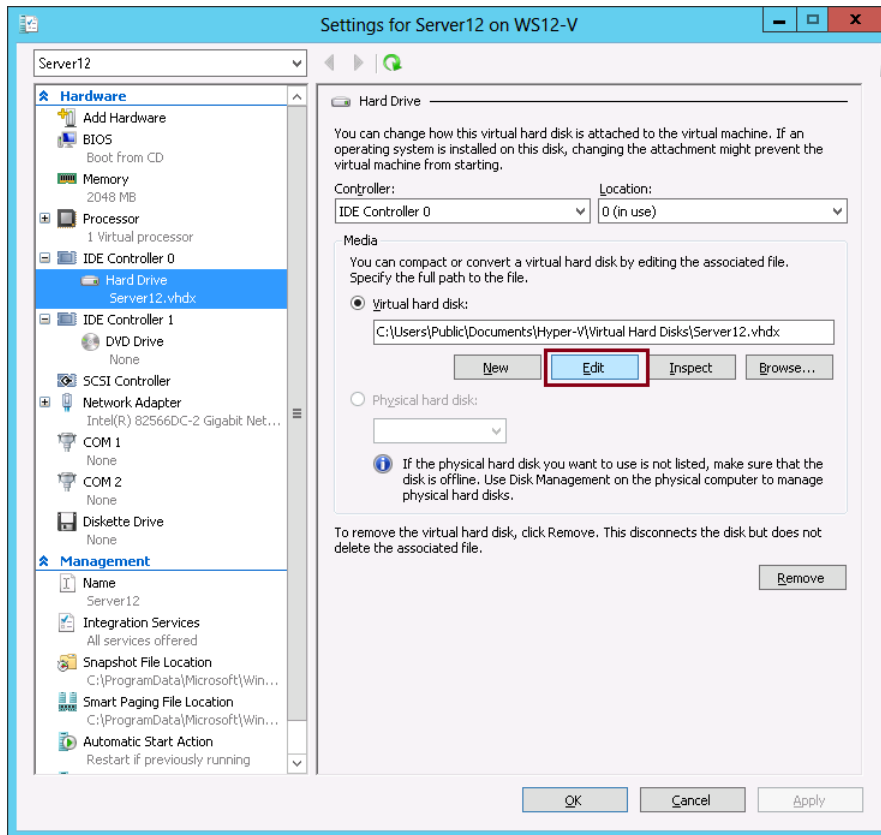


FIGURE 3-10 Converting a virtual hard disk



EXAM TIP

Remember the purpose of the Physical Hard Disk option shown in Figure 3-10. This option is often called a “pass-through disk” and has been available since Windows Server 2008. With a pass-through disk, you add a physical disk (as opposed to a VHD or VHDX) to a VM. As stated in the description of the feature in Figure 3-10, you need to take a physical disk offline before you can attach it to a VM as a pass-through disk.

Then, in the Edit Virtual Hard Disk Wizard, choose the Convert option, shown in Figure 3-11.

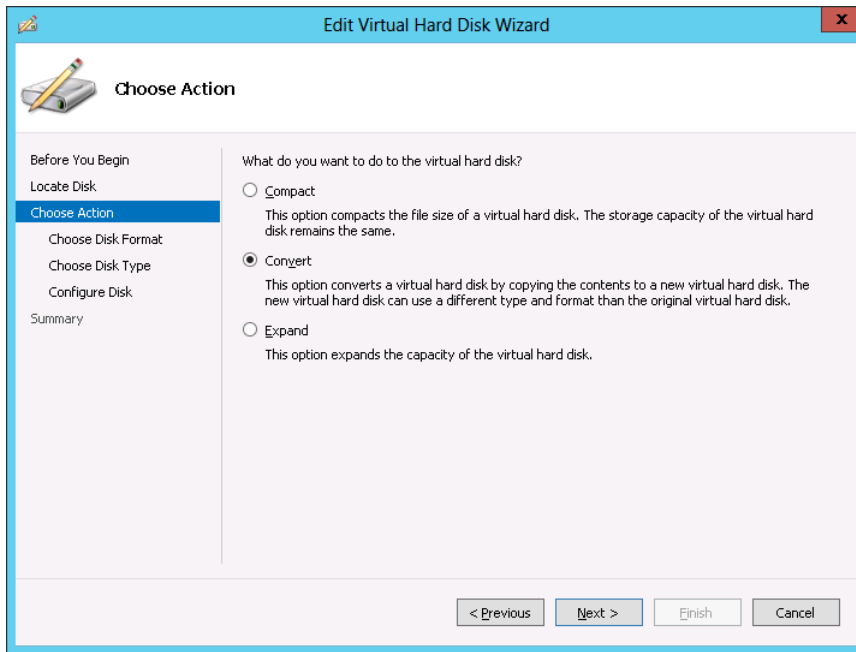


FIGURE 3-11 Converting a virtual hard disk to VHD or VHDX

NOTE To convert a VHD to a VHDX file in Windows PowerShell, use the Convert-VHD cmdlet.

Although size is the biggest advantage of a VHDX, it isn't the only advantage. VHDX files also offer the following benefits:

- Improved resiliency from power failure, thanks to a new disk log.
- Support for new low-cost storage options thanks to 4 KB sector disks.
- Better performance thanks to large block sizes.
- Support for user-defined file metadata. You could use metadata, for example, to include information about the service pack level of the guest operating system on the VM.

Any of these advantages could appear as requirements in a scenario question, so be sure to remember them.

MORE INFO For more information about the new VHDX format in Windows Server 2012, see the article titled "Hyper-V Virtual Hard Disk Format Overview" in the TechNet Library at <http://technet.microsoft.com/en-us/library/hh831446.aspx>.

Virtual Fibre Channel adapter

Before Windows Server 2012, you could provision storage from a Fibre Channel storage area network (SAN) and then use that storage in a guest VM. However, you had to prepare everything in the host operating system so that the source of the storage was transparent to the guest.

What's new in Windows Server 2012 and later is that you can create a Fibre Channel adapter for your VM and then provision storage from your Fibre Channel SAN from within the guest operating system. This might be useful, for example, if you want to migrate to a virtual environment application that is already connected to specific logical unit numbers (LUNs) in your Fibre Channel SAN. Another advantage of the Fibre Channel adapter is that it allows you to cluster guest operating systems to provide high availability for VMs.

To configure virtual Fibre Channel, first use the Virtual SAN Manager option in the Actions pane of Hyper-V Manager to create a new virtual Fibre Channel SAN. Virtual Fibre Channel SANs are connected to one or more physical host bus adapters (HBAs). Then add a new Fibre Channel adapter to the VM. To add a new Fibre Channel adapter to a VM, first open the settings of the VM and select Add Hardware from the menu on the left. Lastly, select Fibre Channel Adapter and click Add, as shown in Figure 3-12.

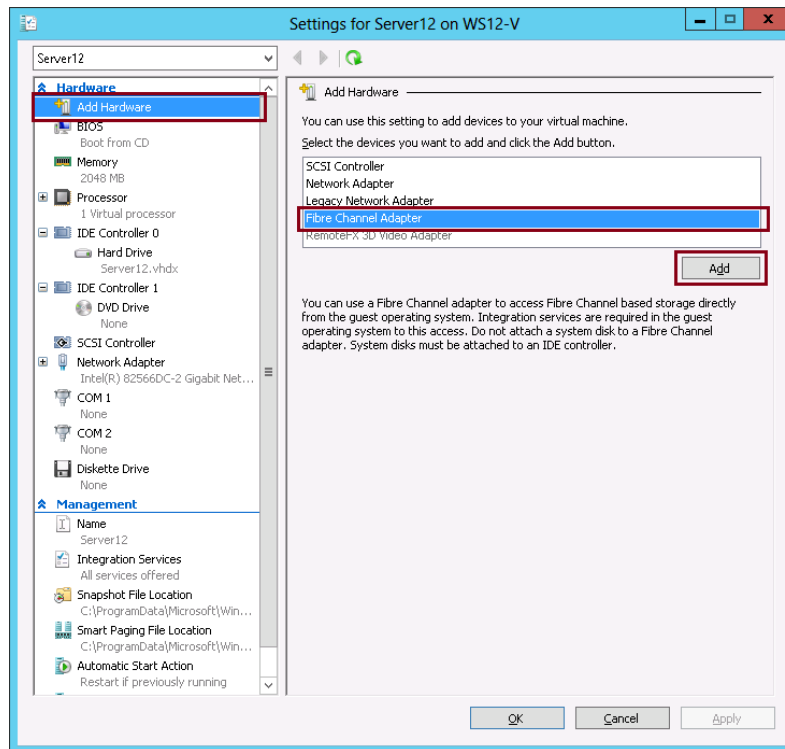


FIGURE 3-12 Adding a new virtual Fibre Channel adapter



EXAM TIP

To configure a VM to connect to a Fibre Channel SAN, first create a virtual Fibre Channel SAN that connects to one or more physical HBAs.

You configure virtual Fibre Channel adapter settings by specifying a virtual SAN. Port addresses are supplied automatically, but you can edit them by clicking Edit Addresses. The port addresses include hexadecimal values representing the World Wide Node Name (WWNN) and World Wide Port Name (WWPN), as shown in Figure 3-13.

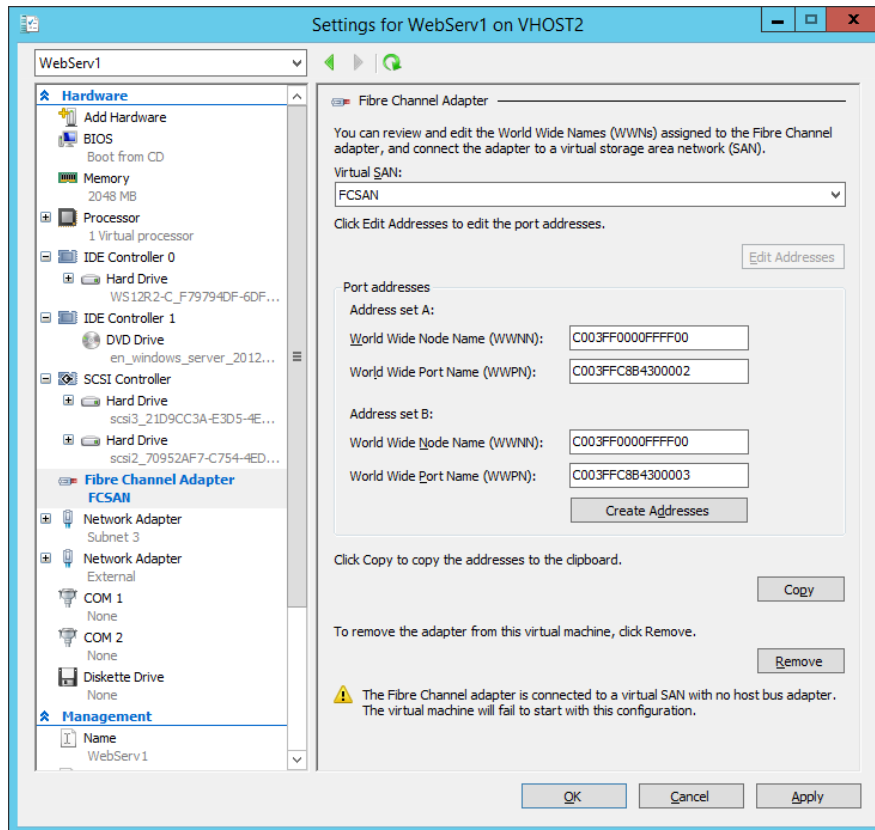


FIGURE 3-13 Configuring Fibre Channel settings

The Fibre Channel adapter in Hyper-V includes a few limitations you need to be aware of. First, the drivers for your HBAs must support virtual Fibre Channel. Second, you can't use virtual Fibre Channel to connect to boot media for your VMs. Finally, you can't use the Fibre Channel adapter with just any guest operating system. The guest has to be running Windows Server 2008 or later.

NOTE The cmdlets used for creating, configuring, and managing virtual Fibre Channel adapters are Add-VMFibreChannelHba, Set-VMFibreChannelHba, Remove- VMFibreChannelHba, and Get-VMFibreChannelHba.

MORE INFO For more information about the Fibre Channel adapter in Hyper-V, see the topic “Hyper-V Virtual Fibre Channel Overview,” at <http://technet.microsoft.com/en-us/library/hh831413.aspx>.

Storage Quality of Service (QoS)

Storage Quality of Service (QoS) is a new feature in Windows Server 2012 R2 that allows you to define a minimum and maximum level of I/O throughput for a virtual disk in Hyper-V. The throughput is defined as an input/output per second (IOPS) value, where each IO is considered to be 8 KB of data. The IOPS limits you set apply only to an individual disk, not to a VM in general.

To configure storage QoS, open the settings of a VM, expand the desired virtual disk in the Hardware menu on the left and then select Advanced Features. In the Advanced Features configuration area on the right, click Enable Quality Of Service Management, and then define a minimum and maximum level for the IOPS. You may leave one value set to zero to accept the system defaults. (Note that the minimum setting does not ensure that this minimum IOPS will be met. The minimum value merely defines a threshold that will trigger an event-based notification.)

Figure 3-14 shows the configuration settings for storage QoS.

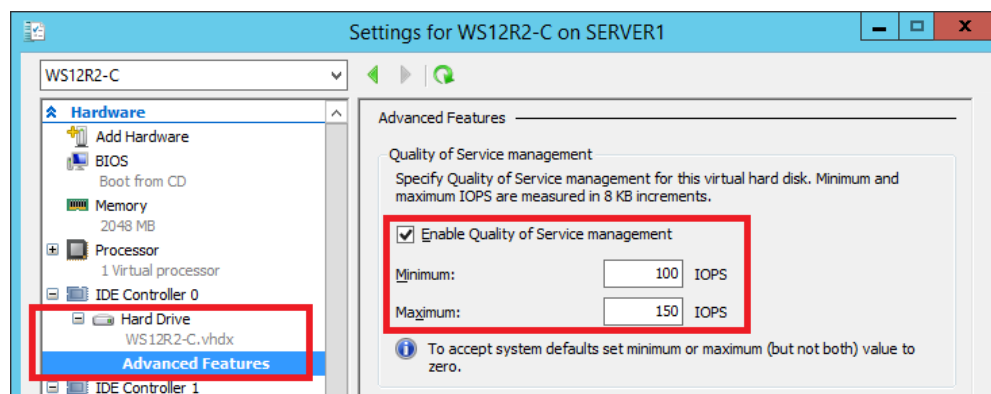


FIGURE 3-14 Configuring storage QoS settings

Why would you need to define storage QoS? One potential application would be to honor different levels of service-level agreements (SLAs) for different clients or for different areas of your organization. For example, your organization might provide three levels of service to clients corresponding to three different ranges of IOPS.

MORE INFO For more information storage QoS, see “Storage Quality of Service for Hyper-V” at <http://technet.microsoft.com/en-us/library/dn282281.aspx>.

Objective summary

- Windows Server 2012 and Windows Server 2012 R2 introduce VHDX files, which have a 64 TB size limit. (VHD files have a 2 TB limit.) Other advantages of the VHDX file format are improved resiliency from power failures, user-defined metadata, and better performance.
- You can convert a VHD to a VHDX and vice versa.
- Hyper-V in Windows Server 2012 and Windows Server 2012 R2 allows you to create virtual Fibre Channel adapters for virtual machines. If you have a Fibre Channel SAN and compatible HBA drivers, you can then provision SAN storage from within a guest VM.
- Storage Quality of Service (QoS) is a new feature in Windows Server 2012 R2 that allows you to define an acceptable range of IOPS for a selected virtual disk in Hyper-V. Each IO is defined as 8 KB.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the “Answers” section at the end of the chapter.

1. You have a VHD that is stored on a server running Windows Server 2012 R2. The VHD is 1.5 TB in size and stores a rapidly growing database file that is 1.0 TB. You want to provide at least 4 TB of space for the database file. What should you do?
 - A. Use the Edit Virtual Hard Disk Wizard and choose the Convert option.
 - B. Use the Edit Virtual Hard Disk Wizard and choose the Expand option.
 - C. Move the contents of the VHD to a new dynamically expanding disk.
 - D. Move the contents of the VHD to a new differencing disk.

2. You work as a network administrator for Fabrikam.com. Fabrikam.com has a server room that includes 20 servers, 10 of which are virtualized on a server named HYPV4 running Windows Server 2012 R2.
- The Fabrikam.com office recently experienced a power outage. After the power outage, the universal power supply connected to HYPV4 did not gracefully shut down HYPV4 or its 10 hosted virtual servers. Some VHDs were corrupted, which required you to restore the VHDs from backup, resulting in a loss of data.
- You want to help ensure that future power outages do not corrupt data on your virtualized servers. What should you do?
- A. Configure NIC teaming for every VM.
 - B. Convert the VHDs on your VMs to VHDX files.
 - C. Create Fibre Channel adapters for each VM and move the VHDs to shared storage.
 - D. Enable data deduplication on HYPV4.
3. You work as a network administrator for Fabrikam.com. One of your servers, named HYPV5, is running Windows Server 2012 R2 and has been configured with the Hyper-V role. HYPV5 hosts five VMs running Windows Server 2008 R2.
- You want to attach new VHDs to the VMs hosted on HYPV5 to increase storage space to these VMs. Until now, the VMs have relied on locally attached storage on HYPV5 to store VHDs attached to the VMs. However, adequate storage space is no longer available on HYPV5 for any new VHDs.
- Your network includes a Fibre Channel SAN, from which HYPV5 can already provision storage. You want to provision new storage from the Fibre Channel SAN and use it for the new VMs, and you want to achieve this with the least amount of administrative effort. What should you do? (Choose all that apply.)
- A. Upgrade the VM operating systems to Windows Server 2012 R2.
 - B. From within the host operating system, provision new storage from the SAN.
 - C. From within the guest operating system, provision new storage from the SAN.
 - D. Convert the VHD files to VHDX files.

Objective 3.3: Create and configure virtual networks

This objective covers the bulk of the new features in Windows Server 2012 and Windows Server 2012 R2 Hyper-V, but it's unclear which of these many features will actually appear on the 70-417 exam. Some, such as virtual switch extensions, are difficult to write questions about for an exam on Windows Server, as opposed to System Center Virtual Machine Manager or Windows development. Others are almost too easy: They can't be set up in a question without giving the answer away, as is the case with bandwidth management, DHCP guard, and router advertisement guard. Still others, such as port ACLs, are constrained by a relative

lack of documentation compared to other features. SR-IOV stands out as a feature for which questions suitable to this exam can be written without too much difficulty, but even it is not currently mentioned by name as a topic in the objective description provided by Microsoft.

As a result, it's difficult to predict what questions you will see on the exam for this objective, so you can only learn the salient points about each of these features and expect to be surprised by any question you might see on the exam.

This section covers the following topics:

- Virtual switch extensions
- Network isolation
- Single-root I/O virtualization (SR-IOV)
- Bandwidth management
- Advanced features for virtual network adapters

Virtual switch extensions

The “virtual networks” that appeared in the Windows Server 2008 and Windows Server 2008 R2 interface have been replaced in Windows Server 2012 and Windows Server 2012 R2 by elements called virtual switches. From an administration point of view, virtual networks appear simply to have been renamed. Network adapters now connect to virtual switches instead of virtual networks, and just like the old virtual networks, virtual switches can be external, internal, or private.

But there is more to virtual switches than meets the eye at first glance. One of the key innovations in Windows Server 2012 and Windows Server 2012 R2 Hyper-V is that the functionality of these new virtual switches can be expanded through extensions provided by Microsoft or independent software vendors. You add these new extensions as you would install any new software.

Windows Server 2012 and Windows Server 2012 R2 allow allows for the following kinds of virtual switch extensions:

- Capturing extensions, which can capture packets to monitor network traffic but cannot modify or drop packets
- Filtering extensions, which are like capturing extensions but also can inspect and drop packets
- Forwarding extensions, which allow you to modify packet routing and enable integration with your physical network infrastructure

Once installed, extensions are made available to all switches but are enabled and disabled on a per-switch basis. To manage installed extensions for a virtual switch, from the Actions pane in Hyper-V Manager, select Virtual Switch Manager, as shown in Figure 3-15.




FIGURE 3-15 Opening the new Virtual Switch Manager

Then, in the Virtual Switch Manager dialog box that opens, expand the desired switch and select Extensions, as shown in Figure 3-16. In the Switch Extensions box, you can enable, disable, and rearrange the order of installed extensions.




FIGURE 3-16 Managing virtual switch extensions

By default, each switch has two extensions: Microsoft NDIS Capture, which is disabled, and Microsoft Windows Filtering Platform, which is enabled.

You can also use PowerShell to create, delete, and configure extensible switches on Hyper-V hosts. Use the `Get-VMSwitchExtension` cmdlet to display details concerning the extensions installed on a specific switch. To see the full list of cmdlets available to manage virtual switches in general, type **Get-Command *VMSwitch*** at a Windows PowerShell prompt.

If any questions about virtual switch extensions appear on the 70-417 exam, they will most likely involve fictional or hypothetical extensions. One possible scenario could involve two extensions that you need to enable but that don't work well together. If such a problem were to occur and you wanted the functionality of both extensions, you could create two separate virtual switches, with one of these extensions enabled on each. Then you could connect chosen VMs to the appropriate switch, as needed.

NOTE To manage virtual switch extensions in Windows PowerShell, you can use `Enable-VMSwitchExtension`, `Disable-VMSwitchExtension`, and `Get-VMSwitchExtension` cmdlets. For a complete list, use `Get-Command` to search for the string `*VMSwitchExtension*`.

Network isolation

You can isolate VMs from unwanted network traffic by using the `Add-VMNetworkAdapterAcl` cmdlet in Windows PowerShell. The feature is sometimes called port ACLs in Microsoft documentation, but on the 70-417 exam, it's possible you will see this feature mentioned only by its associated cmdlets.

Each port ACL is like a firewall rule that allows or denies traffic associated with a Media Access Control (MAC) or IP address. If you configure the port ACL on a Hyper-V host running Windows Server 2012 or later, it remains in effect even if you move the VM to another host server.

For example, to deny both inbound and outbound traffic between the remote address 192.168.9.111 and the VM named Server12, type the following at an elevated Windows PowerShell prompt on the Hyper-V host:

```
Add-VMNetworkAdapterAcl-VMName Server12 -RemoteIPAddress 192.168.9.111 -Direction Both  
-Action Deny
```

You can then review the effects of this last action by using the `Get-VMNetworkAdapterACL` cmdlet. The specific command for this example and its associated output would be as follows:

```
Get-VMNetworkAdapterACL -VMName Server12
VMName: Server12
VMId: eefb383d-5070-4a74-a16b-3e46a5d2b90c
AdapterName: Network Adapter
AdapterId: Microsoft:EEFB383D-5070-4A74-A16B-3E46A5D2B90C\C3F8188F-EF58-480E-A00F-36F55F6CDA52
```

| Direction | Address | Action |
|-----------|----------------------|--------|
| ----- | ----- | ----- |
| Inbound | Remote 192.168.9.111 | Deny |
| Outbound | Remote 192.168.9.111 | Deny |

To remove the port ACL and the associated traffic restriction, use the `Remove-VMNetworkAdapterACL` cmdlet. For instance, following our example, you would type the following:

```
Remove-VMNetworkAdapterACL -VMName Server12 -RemoteIPAddress 192.168.9.111 -Direction Both -Action Deny
```

Resource Metering through port ACLs

You can use the same `Add-VMNetworkAdapterAcl` cmdlet to meter traffic to or from a specific address. To achieve this, use the `Meter` action instead of `Allow` or `Deny`, as in the following example:

```
Add-VMNetworkAdapterAcl -VMName Server12 -RemoteIPAddress 192.168.9.111 -Direction Both -Action Meter
```

You would then use the `Get-VMNetworkAdapterACL` cmdlet to view the metered usage. The following shows the command used with the same example and the associated output:

```
Get-VMNetworkAdapterACL -VMName Server12
VMName: Server12
VMId: eefb383d-5070-4a74-a16b-3e46a5d2b90c
AdapterName: Network Adapter
AdapterId: Microsoft:EEFB383D-5070-4A74-A16B-3E46A5D2B90C\C3F8188F-EF58-480E-A00F-36F55F6CDA52
```

| Direction | Address | Action |
|-----------|----------------------|------------------|
| ----- | ----- | ----- |
| Inbound | Remote 192.168.9.111 | Meter (1 Mbytes) |
| Outbound | Remote 192.168.9.111 | Meter (0 Mbytes) |

Metering usage through port ACLs might seem like an obscure feature, but don't be surprised if it shows up on an exam question. In a way, it's actually a showcase feature of Windows Server 2012 and Windows Server 2012 R2 because it allows virtual hosting providers to meter Internet usage (traffic to the default gateway) specifically as opposed to network usage in general. Like the Resource Metering feature, this base functionality is intended to be leveraged through scripts and programs.

Single-root I/O virtualization (SR-IOV)

Single-root I/O virtualization (SR-IOV) is an extension to the PCI Express (PCIe) standard that can improve network performance. SR-IOV support in Hyper-V is new to Windows Server 2012 and Windows Server 2012 R2. In Hyper-V, SR-IOV enables network traffic to bypass the software switch layer of the Hyper-V virtualization stack and reduce I/O overhead. If you assign only SR-IOV-enabled virtual network adapters and switches to a VM, the network performance of the VM can be nearly as good as that of a physical machine. In addition, the processing overhead on the host is reduced.

To enable SR-IOV, you first need to create a new virtual switch. (You cannot enable SR-IOV on any existing switch, such as the default virtual switch.) In Hyper-V Manager, from the Actions pane, select Virtual Switch Manager. In the Virtual Switch Manager window that opens, choose the option to create a new external virtual switch. Then, in the Virtual Switch Properties pane, in the Connection Type area (shown in Figure 3-17), select the Enable Single-Root I/O Virtualization (SR-IOV) check box. Supply a Name and any Notes for the new virtual switch and then click OK.

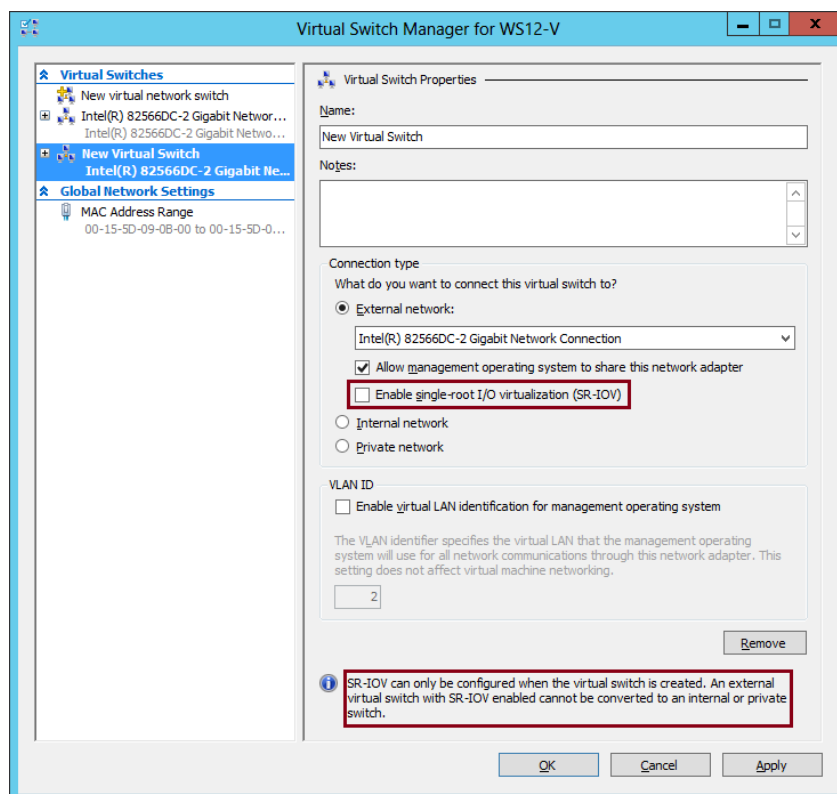


FIGURE 3-17 Enabling SR-IOV on a new virtual switch

NOTE To create a new switch enabled for SR-IOV in Windows PowerShell, use the `New-VMSwitch` cmdlet with the `-EnableIOV $True` parameter.

After you create an SR-IOV-enabled virtual switch, open the settings of the VM for which you want to enable the adapter for SR-IOV and connect the network adapter to the new virtual switch you have just created. Then expand the Network Adapter settings in the Hardware pane, select Hardware Acceleration, and select the Enable SR-IOV check box, shown in Figure 3-18.

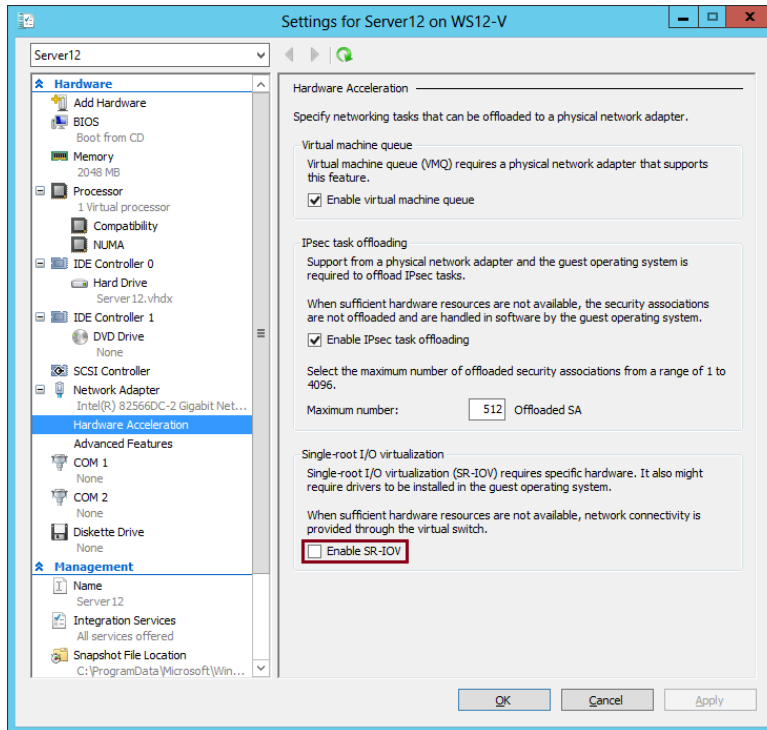


FIGURE 3-18 Enabling SR-IOV on a virtual network adapter

Finally, depending on your hardware configuration, you might need to install drivers within the guest operating system to fully enable SR-IOV. You can check the status of SR-IOV by clicking the Networking tab for a particular VM in Hyper-V Manager. If SR-IOV is active, this information is displayed as shown in Figure 3-19.

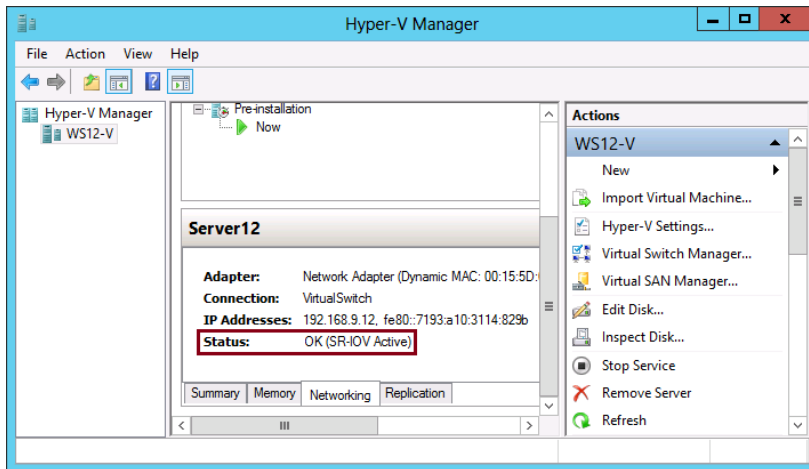


FIGURE 3-19 A status message indicating the SR-IOV is active



EXAM TIP

Remember that single-root I/O virtualization improves network performance on a VM by allowing a VM's network traffic to bypass virtual switches.

MORE INFO For more information about SR-IOV in Hyper-V, search for “Everything you wanted to know about SR-IOV in Hyper-V. Part 1” on <http://blogs.technet.com> or visit <http://blogs.technet.com/b/jhoward/archive/2012/03/12/everything-you-wanted-to-know-about-sr-iov-in-hyper-v-part-1.aspx>.

Bandwidth management

Bandwidth management is a new feature in Windows Server 2012 and Windows Server 2012 R2 Hyper-V that lets you set both a minimum and maximum Mbps of throughput for any virtual network adapter. In Windows Server 2008 R2, you could configure a maximum bandwidth but not a minimum. Now you can configure both a minimum and maximum for each virtual network adapter.

You enable and configure bandwidth management on a virtual network adapter in the settings of a VM, as shown in Figure 3-20. For either the Minimum Bandwidth setting or the Maximum Bandwidth setting, configuring a value of 0 leaves that setting unrestricted.

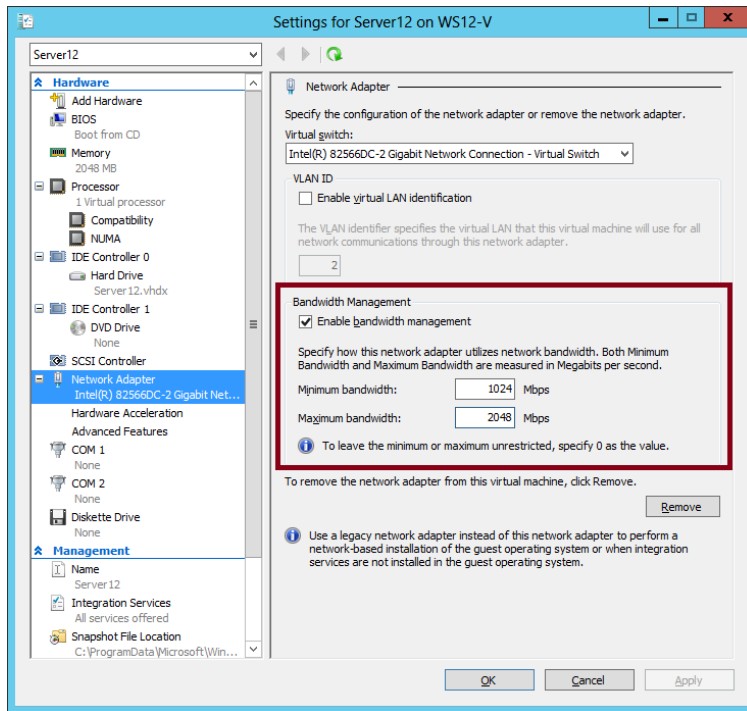


FIGURE 3-20 Enabling bandwidth management

You can also use the `Set-VMNetworkAdapter` to configure minimum and maximum bandwidth on a virtual network adapter. As an alternative to specifying a value for Mbps, you can use this cmdlet to specify a relative bandwidth weight between 0 and 100 relative to other virtual network adapters. To ensure that all virtual network adapters are ensured an equal minimum or maximum bandwidth, you can assign the same bandwidth weight to all adapters. For example, by specifying a bandwidth weight of 1 to all network adapters on servers named `Srv1`, `Srv2`, and `Srv3`, the following command ensures that the same minimum bandwidth is assigned to those network adapters:

```
Get-VMNetworkAdapter -VMName Srv1,Srv2,Srv3 | Set-VMNetworkAdapter
-MinimumBandwidthWeight 1
```



EXAM TIP

Bandwidth management is not available as an option on legacy network adapters. Bandwidth management is available only on standard network adapters in Hyper-V.

MORE INFO For more information about bandwidth management (also called Quality-of-Service for Hyper-V) in Windows Server 2012, visit <http://technet.microsoft.com/en-US/library/hh831511>.

Advanced features for virtual network adapters

A number of new features can be enabled for virtual network adapters in Hyper-V. These options appear when you select Advanced Features after you expand a Network Adapter in the Hardware menu, as shown in Figure 3-21. The new features in this area are defined next.

- **DHCP Guard** Helps safeguard against Dynamic Host Configuration Protocol (DHCP) man-in-the-middle attacks by dropping DHCP server messages from unauthorized VMs pretending to be DHCP servers.
- **Router Guard** Helps safeguard against unauthorized routers by dropping router advertisement and redirection messages from unauthorized VMs pretending to be routers.
- **Port Mirroring** Enables monitoring of a VM's network traffic by forwarding copies of destination or source packets to another VM being used for monitoring purposes.
- **NIC Teaming** In Windows Server 2012 and Windows Server 2012 R2, the NIC teaming feature can be configured for virtual network adapters as well as for physical network adapters.

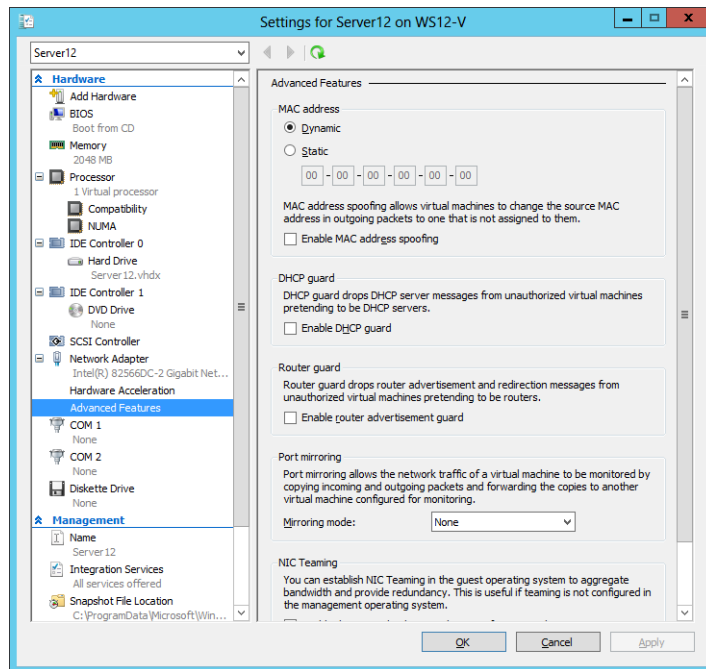


FIGURE 3-21 Configuring advanced features for a virtual network adapter



EXAM TIP

You need to remember the names and functions of these four features for the 70-417 exam.

NOTE To configure settings for a virtual network adapter (including those for SR-IOV, bandwidth management, DHCP guard, router advertisement guard, port mirroring, and NIC teaming), use the `Set-VMNetworkAdapter` cmdlet. Use `Get-Help` to learn about the specific syntax used to configure each feature.

Objective summary

- The functionality of virtual networks in previous versions of Windows Server has been replaced by virtual switches in Windows Server 2012 and Windows Server 2012 R2. Virtual switch features can be enhanced or expanded through extensions, which can be managed in the Hyper-V Manager interface.
- Port ACLs are like firewall rules that allow or deny traffic to a VM based on MAC or IP address. You can also use a port ACL to meter traffic between a VM and a specific address.
- SR-IOV is a way to optimize network performance between a Hyper-V guest and a physical network. To configure SR-IOV, you must create a new virtual switch enabled for SR-IOV, connect a VM's network adapter to that switch, and then enable SR-IOV on the adapter. You might also have to install drivers within the guest operating system.
- Windows Server 2012 and Windows Server 2012 R2 include many new configurable options for network adapters, such as bandwidth management, DHCP guard, router advertisement guard, port mirroring, and NIC teaming.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the "Answers" section at the end of the chapter.

1. You work as a network administrator for Fabrikam.com. Fabrikam.com includes its own cloud infrastructure, which is used to provide virtual hosting services to external customers. Customer servers are hosted as VMs on your servers running Windows Server 2012 R2.

You want to block all traffic to and from the customer virtual servers except for communication with the default gateway.

Which of the following cmdlets should you use on the host servers to isolate the guest VMs?

- A. `Add-VMNetworkAdapterACL`
- B. `Set-VMNetworkAdapterVLAN`
- C. `Set-VMSwitchExtensionPortFeature`
- D. `New-NetFirewallRule`

- 2.** You install the Hyper-V role on a server running Windows Server 2012 R2 and then create a new VM. You now want to optimize network performance for the VM by enabling SR-IOV. What should you do? (Choose all that apply.)
- A.** Create a new private switch.
 - B.** Enable SR-IOV on the virtual switch.
 - C.** Create a new external switch.
 - D.** Enable SR-IOV on the virtual network adapter.
- 3.** You want to maximize security on a VM and help prevent man-in-the-middle attacks. Which of the following settings will help achieve this goal? (Choose all that apply.)
- A.** Enable MAC Spoofing
 - B.** DHCP Guard
 - C.** Router Guard
 - D.** Port Mirroring



Thought experiment

Configuring Hyper-V at Fabrikam

In this thought experiment, apply what you've learned about this objective. You can find answers to these questions in the "Answers" section.

You work as a network administrator for Fabrikam.com, a hosting provider that uses a private cloud infrastructure to provide virtual hosting services to external customers.

Your cloud infrastructure is composed of 20 physical servers running Windows Server 2012 R2 with the Hyper-V role installed. Customer servers are hosted as VMs on these physical servers. Each physical server is equipped with 64 GB of RAM. Shared storage is provided by a Fibre Channel SAN.

Your goals are to use your physical resources as efficiently as possible and to provide a high level of security and performance for customers.

- 1.** You are working with an in-house developer to create a tool that measures CPU, disk, and Internet usage for each customer VM. The developer wants to know how to access this raw information in Windows Server 2012 R2 so that he can build a tool around it. Which method should you show the developer to retrieve the desired usage information?
- 2.** A customer has a database application hosted in your cloud. The application is running in a VM that is running Windows Server 2008 R2 and SQL Server 2008. The database is stored on a VHD drive (stored on the host server) whose size is fixed at 2 TB, but it will soon outgrow the space available. How can you provide more storage space for the database application in a way that minimizes the effort required for all stakeholders?
- 3.** Your IT department has purchased two virtual switch extensions from independent software vendors. The first switch extension is a filtering extension that enables customers to search incoming packets for specific strings or patterns that are useful for security or market research. The second switch extension is a forwarding extension that forwards all incoming traffic received on a switch to any chosen IP address.

You want to be able to use these extensions to provide customers with the ability to search packets on the wire without significantly degrading network performance for services hosted on the customer VM. How can you achieve this goal?

Answers

This section contains the answers to the Objective Reviews and the Thought Experiment.

Objective 3.1: Review

1. Correct answer: B

- A. Incorrect:** Increasing the Startup Memory value will only increase the likelihood that Smart Paging will be used during startup.
- B. Correct:** Increasing the Minimum Memory setting will help ensure that more physical memory remains allocated to VM1 when a restart begins.
- C. Incorrect:** This isn't the best option because it could deprive other important applications of needed RAM.
- D. Incorrect:** This setting would prioritize memory allocation to VM1 when needed. It wouldn't ensure that more memory is allocated to VM1 at the time of a restart operation.

2. Correct answer: D

- A. Incorrect:** Increasing the Startup Memory setting would decrease the possibility that VM2 will be able to start successfully because it will require HYPV2 to find more RAM to allocate to the startup operation. In addition, the fact that Smart Paging is not helping VM2 start indicates most likely that the drive that stores the Smart Paging file has run out of space.
- B. Incorrect:** Decreasing the Maximum Memory on other VMs would have an unpredictable effect on the availability of RAM for VM2 during a restart operation. In addition, the other running VMs might host high-priority applications that need the memory. Finally, the fact that Smart Paging is not helping VM2 start most likely indicates that the drive that stores the Smart Paging file has run out of space.
- C. Incorrect:** Increasing the Memory Buffer % setting would allocate more RAM to VM2 while it is running and would likely make some more memory available at the time of a restart. However, VM2 hosts only low-priority applications that are rarely used. Allocating RAM to VM2 while it is running would deprive other VMs of the RAM they might need to support good performance in higher priority applications. In addition, the fact that Smart Paging is not helping VM2 start most likely indicates that the drive that stores the Smart Paging file has run out of space.
- D. Correct:** Both the host server and the guest VM are running Windows Server 2012 R2, which supports Smart Paging. If insufficient RAM is available for a restart operation, the Smart Paging feature will normally rely on disk storage as virtual memory to help perform the restart. If a guest VM cannot restart in this scenario, it is most likely because not enough free space is available on the disk that currently stores the Smart Paging file.

3. Correct answer: C

- A. Incorrect:** This command would enable Resource Metering on the VM. However, according to the question, Resource Metering is already enabled.
- B. Incorrect:** This command would stop the metering of resources on VM3 but would not display any usage statistics.
- C. Correct:** This command would display usage statistics on VM3 because Resource Metering was enabled or reset.
- D. Incorrect:** This command would save the resource data into a stored variable, not display it on the screen.

Objective 3.2: Review

1. Correct answer: A

- A. Correct:** VHDs have a size limit of 2 TB. The Convert option allows you to change the disk type to a VHDX, which has a size limit of 64 TB.
- B. Incorrect:** VHDs have a size limit of 2 TB; you need a VHD file that is larger than 4 TB. Choosing the Expand option would allow you to expand the size of the VHD from 1.5 TB to 2.0 TB.
- C. Incorrect:** Creating a dynamically expanding VHD would not allow you to move beyond the 2 TB limit for VHD files. You need to convert the disk to a VHDX file.
- D. Incorrect:** Creating a differencing VHD would not allow you to move beyond the 2 TB limit for VHD files. You need to convert the disk to a VHDX file.

2. Correct answer: B

- A. Incorrect:** NIC teaming will help ensure against network outages, but it will not help ensure against data corruption after a power failure.
- B. Correct:** VHDX files—unlike VHD files—contain a log that helps these virtual disks avoid corruption resulting from a power outage.
- C. Incorrect:** Moving the VHDs to shared storage will not make them more resilient to power outages.
- D. Incorrect:** Data deduplication allows data to be stored more efficiently, but it doesn't help prevent corruption from power outages.

3. Correct answer: B

- A. Incorrect:** You don't need to upgrade. You can currently provide new storage for the VMs simply by provisioning new storage for the host server. You would need to upgrade to Windows Server 2012 or later only if you needed to provision storage directly from the guest operating system.
- B. Correct:** You can provision storage from the SAN in the host operating system running Windows Server 2012 R2. Then you can configure new volumes on the host server and then store new VHDs for the VMs on those new volumes.
- C. Incorrect:** You don't need to provision new storage from the SAN from the guest operating system. To do this would require you to upgrade the guest operating systems to Windows Server 2012 or later. You would then need to create and configure virtual Fibre Channel ports. This set of actions would not allow you to achieve your goal with the least amount of administrative effort.
- D. Incorrect:** Converting the VHD files to VHDX files would require you to upgrade the guest operating systems to Windows Server 2012 or later. In addition, converting to VHDX would not help you attach more available storage to your VMs.

Objective 3.3: Review

1. Correct answer: A

- A. Correct:** You can use `Add-VMNetworkAdapterAcl` to create a port ACL and allow or deny traffic between a VM and any specified addresses.
- B. Incorrect:** This cmdlet allows you to associate a VLAN ID with a network adapter. It does not isolate network traffic in a way that would be useful in this specific scenario.
- C. Incorrect:** This cmdlet allows you to configure a feature on a virtual network adapter. It doesn't allow you to restrict network traffic in a way that would be helpful in this scenario.
- D. Incorrect:** This cmdlet allows you to restrict traffic between any address and the host server, not the guest VMs.

2. Correct answers: B, C, D

- A. Incorrect:** You can enable SR-IOV only on an external switch.
- B. Correct:** You need to enable SR-IOV on a new external virtual switch.
- C. Correct:** You can enable SR-IOV only on a new switch. The switch must be external.
- D. Correct:** You need to enable SR-IOV on the virtual network adapter connected to the new virtual switch.

3. Correct answers: B, C

- A. Incorrect:** MAC spoofing enables you to choose a MAC address manually. It doesn't prevent man-in-the-middle attacks.
- B. Correct:** DHCP guard prevents man-in-the-middle attacks from unauthorized VMs pretending to be legitimate DHCP servers.
- C. Correct:** Router guard prevents man-in-the-middle attacks from unauthorized VMs pretending to be legitimate routers.
- D. Incorrect:** Port mirroring is used to forward traffic to a remote VM. It is not used to prevent man-in-the-middle attacks.

Thought experiment

- 1.** To measure CPU and disk usage, use the `Enable-VMResourceMetering`, `Measure-VM`, and `Reset-VMResourceMetering` cmdlets. To measure Internet usage, create a port ACL that measures traffic specifically between a VM and the default gateway by using the `Add-VMNetworkAdapterAcl` cmdlet with the `-Meter` action.
- 2.** Back up the VHD. Convert the VHD to a VHDX. Expand the new VHDX to a desired size up to 64 TB. (Only the host needs to be running Windows Server 2012 or later to support VHDX files. You don't need to upgrade the guest operating system to Windows Server 2012 or later.)
- 3.** Enable only the forwarding extension on the virtual switch currently used by the services hosted on the VM. Create a second virtual switch that enables only the filtering extension.

Install and administer Active Directory

The important feature changes that have appeared since Microsoft Windows Server 2008 R2 in the Install and administer Active Directory domain all fall within a single objective, “Install domain controllers.” Within this small area, however, the change could hardly be more significant: The very tool used in pre-Windows Server 2012 versions of Windows Server to install a domain controller, Dcpromo.exe, has been deprecated (which in Microsoft lingo means “officially set on a path to obsolescence”). More specifically, the use of Dcpromo is highly restricted in Windows Server 2012 and Windows Server 2012 R2. You can use it for domain controller promotion *only with an answer file*. (Dcpromo also retains some specialized uses, such as the force-removal of Active Directory Domain Services with the /ForceRemoval parameter.)

What takes the place of Dcpromo in Windows Server 2012 and Windows Server 2012 R2? A new Active Directory Domain Services Configuration Wizard and a new set of Windows PowerShell cmdlets.

You need to understand these new installation tools well for the 70-417 exam.

Objectives in this chapter:

- Objective 4.1: Install domain controllers

Objective 4.1: Install domain controllers

In Windows Server 2008 and Windows Server 2008 R2, you had the option of installing the Active Directory Domain Services server role before promoting the server to a domain controller. In Windows Server 2012 and Windows Server 2012 R2, that step is now mandatory, and as a result, installing a domain controller is now a two-step process.

This section covers the following topics:

- Installing domain controllers by using GUI
- Installing domain controllers by using Windows PowerShell
- Install From Media option without defragmentation
- Deploying a domain controller in Windows Azure

Installing domain controllers in the GUI

The first step in deploying a domain controller is to add the Active Directory Domain Services server role. If you perform this first step by using the Add Roles And Features Wizard, the second step in deploying a domain controller is easy: You can just choose the option to promote the server to a domain controller on the final page of this wizard, as shown in Figure 4-1.

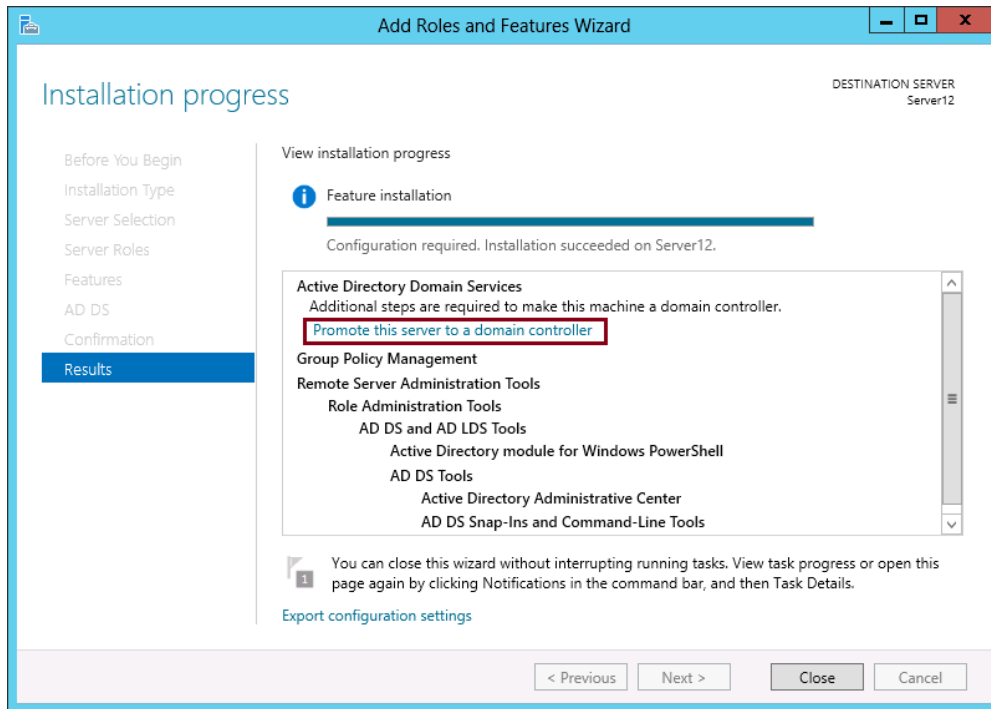


FIGURE 4-1 Installing a domain controller after installing the Active Directory Domain Services server role



EXAM TIP

Make sure the RID Master is online and available whenever you add a new domain controller to an existing domain. If you are installing the first domain controller in a new domain in an existing forest, then the Domain Naming Master must be available.

If you prefer to promote the server later, you can do so using Server Manager. In Server Manager, expand the notifications menu and choose the option to promote the server to a domain controller, as shown in Figure 4-2. (Note that this option appears only if you have added the Active Directory Domain Services server role.)

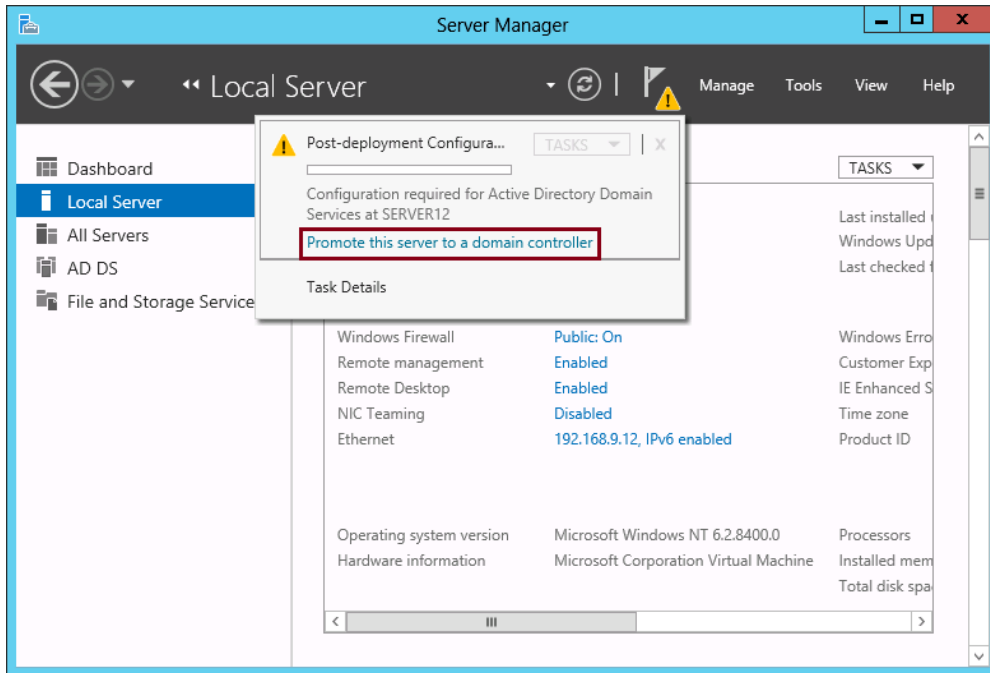


FIGURE 4-2 Installing a domain controller in Server Manager

Whether you select the option immediately at the end of the Add Roles And Features Wizard or later on the notification menu in Server Manager, the result is the same: The Active Directory Domain Services Configuration Wizard is started. This wizard is very similar to the old Active Directory Domain Services Installation Wizard in Windows Server 2008 and Windows Server 2008 R2, even though the code behind it has been rewritten completely from scratch. The first page of the Active Directory Domain Services Configuration wizard is shown in Figure 4-3.




FIGURE 4-3 The Active Directory Domain Services Configuration Wizard is new to Windows Server 2012 and Windows Server 2012 R2

The options that appear in the wizard are so similar to those that appeared in the corresponding wizard in Windows Server 2008 and Windows Server 2008 R2 that it is not necessary to review them all. However, the Review Options page reveals an interesting change, shown in Figure 4-4. Remember how in Windows Server 2008 and Windows Server 2008 R2, you could export the settings you had selected in the wizard to an *answer file* to be used with Dcpromo? In Windows Server 2012 and Windows Server 2012 R2, that option is gone and is replaced by a new option to export the settings you have selected to a *Windows PowerShell script*.

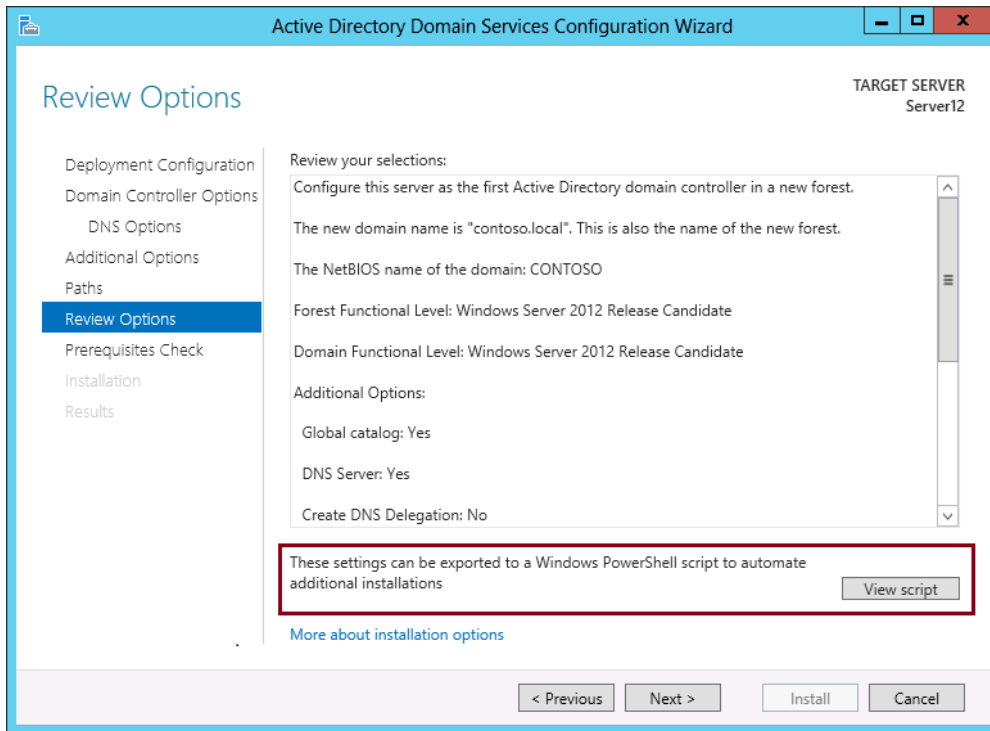


FIGURE 4-4 Exporting domain controller installation settings to a Windows PowerShell script

The following code shows the contents of the exported Windows PowerShell script that results when you choose to add a new domain controller to an existing domain (fabrikam.local):

```
#
# Windows PowerShell script for AD DS Deployment
#
Import-Module ADDSDeployment
Install-ADSDomainController `
-NoGlobalCatalog:$false `
-CreateDnsDelegation:$false `
-Credential (Get-Credential) `
-CriticalReplicationOnly:$false `
-DatabasePath "C:\Windows\NTDS" `
-DomainName "fabrikam.local" `
-InstallDns:$true `
-LogPath "C:\Windows\NTDS" `
-NoRebootOnCompletion:$false `
-SiteName "Default-First-Site-Name" `
-SysvolPath "C:\Windows\SYSVOL" `
-Force:$true
```

Here's another example of the contents of the exported Windows PowerShell script. This version results when you choose to add a first domain controller to a new forest (contoso.local):

```
#
# Windows PowerShell script for AD DS Deployment
#
Import-Module ADDSDeployment
Install-ADDSForest `
-CreateDnsDelegation:$false `
-DatabasePath "C:\Windows\NTDS" `
-DomainMode "Win2012R2" `
-DomainName "contoso.local" `
-DomainNetbiosName "CONTOSO" `
-ForestMode "Win2012R2" `
-InstallDns:$true `
-LogPath "C:\Windows\NTDS" `
-NoRebootOnCompletion:$false `
-SysvolPath "C:\Windows\SYSVOL" `
-Force:$true
```



EXAM TIP

Remember that if you save the Windows PowerShell script, you should save it as a .ps1 file.

MORE INFO For a detailed walkthrough of the Active Directory Domain Services Configuration Wizard, search for "Introducing the first Windows Server 2012 Domain Controller (Part 1 of 2)" or visit <http://blogs.technet.com/b/askpfeplat/archive/2012/09/03/introducing-the-first-windows-server-2012-domain-controller.aspx>.

Adprep runs automatically

Windows Server 2012 and Windows Server 2012 R2 introduce another new feature related to the Active Directory Domain Services Configuration Wizard: This wizard runs Adprep /domainprep and Adprep /forestprep automatically as needed. Before Windows Server 2012, if upgrading the schema was necessary to install a new domain controller, you needed to run Adprep beforehand.

In Windows Server 2012 and Windows Server 2012 R2, you still have the option of running Adprep to upgrade the schema, but if you haven't done so before completing the Active Directory Domain Services Configuration Wizard, the schema is upgraded automatically without prompting you first.

Of course, in a production environment, you wouldn't want to make such a significant change without testing, preparation, and planning. Extending the schema ahead of time and in a controlled manner is preferred as a way to minimize risk.

Installing domain controllers with Windows PowerShell

The ability to install a domain controller by using Windows PowerShell is new to Windows Server 2012 and Windows Server 2012 R2. As with using the GUI to perform the same task, promoting a domain controller is a two-step process when you use Windows PowerShell.

First, you need to add the Active Directory Domain Services server role by typing the following at an elevated Windows PowerShell prompt:

```
Install-WindowsFeature -Name AD-Domain-Services -IncludeManagementTools
```

After the role installation is complete, you can get information about the available cmdlets in the ADDSDeployment module. To do so, type the following:

```
Get-Command -Module ADDSDeployment
```

MORE INFO For a listing of the cmdlets in the ActiveDirectory module (as opposed to the ADDSDeployment module), visit <http://technet.microsoft.com/en-us/library/hh852274.aspx> or type the following at a Windows PowerShell prompt:

```
Get-Command -Module ActiveDirectory
```

Table 4-1 shows the 10 cmdlets available in the ADDSDeployment module.

TABLE 4-1 The cmdlets in the ADDSDeployment module

| Cmdlet | Description |
|---|--|
| Add-ADDSTReadOnlyDomainControllerAccount | Creates a Read-Only Domain Controller (RODC) account that can be used to install an RODC in Active Directory |
| Install-ADDSDomain | Installs a new Active Directory domain configuration |
| Install-ADDSDomainController | Installs a domain controller in Active Directory |
| Install-ADDSTForest | Installs a new Active Directory forest configuration |
| Test-ADDSDomainControllerInstallation | Runs the prerequisites (only) for installing a domain controller in Active Directory |
| Test-ADDSDomainControllerUninstallation | Runs the prerequisites (only) for uninstalling a domain controller in Active Directory |
| Test-ADDSDomainInstallation | Runs the prerequisites (only) for installing a new Active Directory domain configuration |
| Test-ADDSTForestInstallation | Runs the prerequisites (only) for installing a new forest in Active Directory |
| Test-ADDSTReadOnlyDomainControllerAccountCreation | Runs the prerequisites (only) for adding an RODC account |
| Uninstall-ADDSDomainController | Uninstalls a domain controller in Active Directory |



EXAM TIP

You need to understand the function of all of these cmdlets for the 70-417 exam.

Installing the first domain controller in a new forest

To install a domain controller in a new forest, use the Test-ADDSForestInstallation and Install-ADDSForest cmdlets. (The Test-ADDSForestInstallation cmdlet is optional.)

TEST-ADDSFORESTINSTALLATION

Use Test-ADDSForestInstallation to verify that your environment meets the prerequisites to install the first domain controller in the new forest with the parameters specified. These same prerequisite tests are run if you use the Install-ADDSForest cmdlet.

For example, the following command runs the prerequisite tests for installing a new forest named corp.contoso.com. Because it doesn't specify a password with the -SafeModeAdministratorPassword parameter, the user will be prompted to supply a Directory Services Restore Mode (DSRM) password.

```
Test-ADDSForestInstallation -DomainName "corp.contoso.com"
```

The following command provides a more complex example. Here, the prerequisite tests are run for installing a new forest with the following specifications:

- Create a DNS delegation in the parent contoso.com domain (-CreateDNSDelegation).
- Set the domain functional level to Windows Server 2008 (-DomainMode Win2008).
- Set the forest functional level to Windows Server 2008 R2 (-ForestMode Win2008R2).
- Install the Active Directory database and SYSVOL on the D drive (-DatabasePath "D:\NTDS" -SysvolPath "D:\SYSVOL").
- Install the log files on the E drive (-LogPath "E:\Logs").
- Have the server automatically restart after Active Directory Domain Services installation is complete (No -NoRebootOnCompletion parameter).
- Prompt the user to provide and confirm the DSRM password (No -SafeModeAdministratorPassword parameter).

```
Test-ADDSForestInstallation -DomainName corp.contoso.com -CreateDNSDelegation  
-DomainMode Win2008 -ForestMode Win2008R2 -DatabasePath "D:\NTDS" -SysvolPath  
"D:\SYSVOL" -LogPath "E:\Logs"
```

MORE INFO For more information about the Test-ADDSForestInstallation cmdlet, visit <http://technet.microsoft.com/en-us/library/hh974717>.

INSTALL-ADDSFOREST

After you have tested the new forest creation with the Test-ADDSDomainControllerInstallation cmdlet, you are ready to use the Install-ADDSDomainController cmdlet to install the domain controller and create the new forest.

For example, the following command will create a new forest with the name corp.contoso.com, install a Domain Name System (DNS) server on the new local domain controller, and prompt the user to provide and confirm a DSRM password:

```
Install-ADDSDomainController -DomainName "corp.contoso.com" -InstallDNS
```

Note that in this next, more complex example, all of the same parameters used with the second Test-ADDSDomainControllerInstallation example are used again. Running the earlier test gives you assurance that the following command will work:

```
Install-ADDSDomainController -DomainName corp.contoso.com -CreateDNSDelegation -DomainMode Win2008 -ForestMode Win2008R2 -DatabasePath "D:\NTDS" -SysvolPath "D:\SYSVOL" -LogPath "E:\Logs"
```

MORE INFO For more information about the Install-ADDSDomainController cmdlet, visit <http://technet.microsoft.com/en-us/library/hh974720>.



EXAM TIP

Remember that the Netlogon.dns file lists all of the SRV records registered in DNS when a domain controller is installed.

Installing an additional domain controller in an existing domain

To install a domain controller in an existing domain, use the Test-ADDSDomainControllerInstallation and Install-ADDSDomainController cmdlets. (The Test-ADDSDomainControllerInstallation cmdlet is optional.)

TEST-ADDSDOMAINCONTROLLERINSTALLATION

This cmdlet runs prerequisite checks that verify that you can use the Install-ADDSDomainController cmdlet to install a domain controller in Active Directory.

For example, the following command runs prerequisite checks to verify the possibility of installing a domain controller in the existing corp.contoso.com domain (using domain administrator credentials). The domain controller will include a DNS server. The user will be prompted to enter and confirm the DSRM password.

```
Test-ADDSDomainControllerInstallation -InstallDNS -Credential (Get-Credential CORP\Administrator) -DomainName "corp.contoso.com"
```


MORE INFO For more information about the Test-ADDSDomainControllerInstallation cmdlet, visit <http://technet.microsoft.com/en-us/library/hh974725>.

INSTALL-ADDSDOMAINCONTROLLER

If the test completes successfully, you can use the Install-ADDSDomainController cmdlet with the same parameters to install the new domain controller in the existing corp.contoso.com domain:

```
Install-ADDSDomainController -InstallDns -Credential (Get-Credential CORP\Administrator)
-DomainName "corp.contoso.com"
```



EXAM TIP

For the 70-417 exam, remember that you first need to use Install-WindowsFeature to install the AD DS binaries. To then promote a server domain controller, use Install-ADDSTForest, Install-ADDSDomain, or Install-ADDSDomainController. Finally, to demote a domain controller, use Uninstall-ADDSDomainController. (For a force removal, you can use Uninstall-ADDSDomainController -ForceRemoval or even Dcpromo.exe /ForceRemoval.)

MORE INFO For more information about the Install-ADDSDomainController cmdlet, visit <http://technet.microsoft.com/en-us/library/hh974723>.

Installing the first domain controller in a new domain in an existing forest

To install the first domain controller in a new domain in an existing forest, use the Test-ADDSDomainInstallation and Install-ADDSDomain cmdlets. (The Test-ADDSDomainInstallation cmdlet is optional.)

TEST-ADDSDOMAININSTALLATION

This cmdlet runs the prerequisite checks that would be performed if you were to use the Install-ADDSDomain cmdlet to install a new domain controller in a new domain in an existing forest.

The following example runs prerequisite checks that verify the possibility of creating a new child domain named child.corp.contoso.com by using credentials of CORP\EnterpriseAdmin1. Because it doesn't use the -SafeModeAdministratorPassword parameter, the user will be prompted to provide and confirm the DSRM password to complete the checks. The following example also verifies the possibility of accomplishing the following:

- Installing a DNS server (-InstallDNS)
- Creating a DNS delegation in the corp.contoso.com domain (-CreateDNSDelegation)

- Setting the domain functional level to Windows Server 2003 (-DomainMode Win2003)
- Making the domain controller a global catalog server (no -NoGlobalCatalog parameter).
- In a site named Houston (-SiteName Houston)
- Using DC1.corp.contoso.com as the replication source domain controller (-ReplicationSourceDC DC1.corp.contoso.com)
- Installing the Active Directory database and SYSVOL on the D drive (-DatabasePath "D:\NTDS" -SYSVOLPath "D:\SYSVOL")
- Installing the log files on the E drive (-LogPath "E:\Logs")
- Avoiding an automatic restart after the domain installation is complete (-NoRebootOnCompletion)

```
Test-ADDSDomainInstallation -Credential (Get-Credential CORP\EnterpriseAdmin1)
-NewDomainName child -ParentDomainName corp.contoso.com -InstallDNS -CreateDNSDelegation
-DomainMode Win2003 -ReplicationSourceDC DC1.corp.contoso.com -SiteName Houston
-DatabasePath "D:\NTDS" -SYSVOLPath "D:\SYSVOL" -LogPath "E:\Logs" -NoRebootOnCompletion
```

MORE INFO For more information about the Test-ADDSDomainInstallation cmdlet, visit <http://technet.microsoft.com/en-us/library/hh974715.aspx>.

INSTALL-ADDSDOMAIN

After you have tested the possibility of creating a new domain with the Test-ADDSDomainInstallation cmdlet, you are ready to use the Install-ADDSDomain to install the first domain controller in the new domain in an existing forest.

The following cmdlet creates the domain with the configuration verified in the test:

```
Install-ADDSDomain -Credential (Get-Credential CORP\EnterpriseAdmin1) -NewDomainName
child -ParentDomainName corp.contoso.com -InstallDNS -CreateDNSDelegation -DomainMode
Win2003 -ReplicationSourceDC DC1.corp.contoso.com -SiteName Houston -DatabasePath
"D:\NTDS" -SYSVOLPath "D:\SYSVOL" -LogPath "E:\Logs" -NoRebootOnCompletion
```

MORE INFO For more information about the Install-ADDSDomain cmdlet, visit <http://technet.microsoft.com/en-us/library/hh974722>.

Adding an RODC account

Use the Test-ADDSDomainReadOnlyDomainControllerAccountCreation and Add-ADDSDomainReadOnlyDomainControllerAccount cmdlets to pre-create a computer account for an RODC. (The Test-ADDSDomainReadOnlyDomainControllerAccountCreation is optional.) You can optionally use the -DelegatedAdministratorAccountName parameter to give a non-administrator (or a "delegated administrator") the rights and permissions required to install

the RODC. Once you have added the RODC account, you or the delegated administrator can use the `Install-ADDSDomainController` cmdlet with the `-ReadOnlyReplica` switch parameter to install an RODC.

For example, the following command adds a new RODC account to the `corp.contoso.com` domain using the North America site as the source site for the replication source domain controller, while delegating the rights and permissions required to install the RODC to a user named `User1`.

```
Add-ADDSDomainControllerAccount -DomainControllerAccountName RODC1 -DomainName corp.contoso.com -SiteName NorthAmerica -DelegatedAdministratorAccountName corp.contoso.com\User1
```



EXAM TIP

Remember that pre-creating an RODC computer account is a way to let you give permission to a non-administrator to install an RODC. If you don't specify the delegated administrator when the RODC computer account is created, you can specify a user or group account on the **Managed By** tab of the RODC computer account in Active Directory Users and Computers. In general, it's recommended that you create a test RODC computer account and review all of its properties tabs before you take the 70-417 exam. For example, you should know that you can specify replication partners for an RODC through the NTDS Settings available on the **General** tab.

MORE INFO For more info on the `Test-ADDSDomainControllerAccountCreation` cmdlet and `Add-ADDSDomainControllerAccount` cmdlet, visit <http://technet.microsoft.com/en-us/library/hh974721> and <http://technet.microsoft.com/en-us/library/hh974718>.

Uninstalling a domain controller

Use the `Test-ADDSDomainControllerUninstallation` and `Uninstall-ADDSDomainController` cmdlets to uninstall a domain controller. Unlike the previous cmdlets, these cmdlets can be used without any parameters. If you do so, you will be prompted to supply a local Administrator password.

MORE INFO For more information about the `Test-ADDSDomainControllerUninstallation` and `Uninstall-ADDSDomainController` cmdlets, visit <http://technet.microsoft.com/en-us/library/hh974716> and <http://technet.microsoft.com/en-us/library/hh974714>.

Install from Media (IFM) option without defragmentation

Windows Server has included an Install from Media (IFM) option for deploying domain controllers since Windows Server 2003. With this option, Active Directory Domain Services data is stored on a local drive, on removable media such as a DVD, or on a network shared folder. Using IFM allows you to avoid replicating all directory data over the network when you install the new domain controller.

The recommended method for creating Active Directory Domain Services installation media is to use the Ntdsutil.exe tool that is available when the Active Directory Domain Services server role is installed. Ntdsutil includes an IFM subcommand menu that creates the files necessary to install Active Directory Domain Services by using the IFM option.

Windows Server 2012 and Windows Server 2012 R2 introduce two additional options to this IFM menu. These two options allow you to create IFM stores without first performing an offline defrag of the exported NTDS.dit database file. An offline defrag is performed by default, an operation that can be time-consuming. When disk space is not a premium and you do not need to compact the Active Directory database, these two options save time creating the IFM.

Table 4-2 describes the two menu items.

TABLE 4-2 Creating IFM media without defragmentation

| Menu Item | Description |
|--------------------------------|---|
| Create Full NoDefrag %s | Create IFM media without defragmenting for a full Active Directory domain controller (without SYSVOL) |
| Create Sysvol Full NoDefrag %s | Create IFM media with SYSVOL and without defragmenting for a full Active Directory domain controller into folder %s |



EXAM TIP

Remember that IFM does not work across different operating system versions. If you have a domain controller running Windows Server 2008 R2 and you want to promote a server running Windows Server 2012 or Windows Server 2012 R2 by using IFM, you need to upgrade the domain controller first.

MORE INFO To learn more about how to use Offline Domain Join and Djoin.exe, search for “Offline Domain Join (Djoin.exe) Step-by-Step Guide” or visit [http://technet.microsoft.com/en-us/library/offline-domain-join-djoin-step-by-step\(v=ws.10\).aspx](http://technet.microsoft.com/en-us/library/offline-domain-join-djoin-step-by-step(v=ws.10).aspx).



EXAM TIP

Windows Server 2008 R2 introduced a feature named Offline Domain Join that relies on a utility called Djoin.exe. Djoin.exe allows you to join a computer to a domain even if there is no live connection to a domain controller. If you earned your current Windows Server 2008 MCSA certification before the release of Windows Server 2008 R2, you might have missed this feature completely. Even though Djoin.exe isn't new to Windows Server 2012 or Windows Server 2012 R2, it's a good bet that you'll see a simple question about it on the 70-417 exam. Note that for this exam, you don't need to understand the specifics of how to use Djoin.exe; you just should know what it is used for.

Installing a domain controller in Windows Azure

Many companies are starting to use public cloud platforms such as Windows Azure (soon to be renamed Microsoft Azure) to host private servers and applications intended only for use internally within the same company. For example, a multi-branch organization might decide to deploy a line-of-business application in the cloud so that workers in all branch offices can access that application over the Internet. This IT decision then introduces a need for authentication in the cloud. Applications servers need some way to authenticate users, and the authentication provider should normally be local to the application server. A domain controller is often the best answer to provide this authentication.

The procedure of deploying domain controllers in Microsoft's public cloud is now relevant for the 70-417 exam. In early 2014, the Install Domain Controllers objective was updated to include the task "Deploy Active Directory IaaS in Windows Azure." IaaS refers to infrastructure-as-a-service, which, in turn, simply refers to a cloud service feature that lets you interact with hosted VMs at the operating system level but that hides the underlying infrastructure from you. So ultimately, "deploying Active Directory IaaS in Windows Azure" isn't much different from installing a domain controller locally on company premises. However, there are a few Windows Azure-related concepts and preparatory steps that you definitely need to understand for the exam.

NOTE Active Directory IaaS is not the only way to authenticate users in Windows Azure. Windows Azure also provides a cloud-based authentication service branded as "Windows Azure Active Directory." Windows Azure Active Directory is not included in the objectives for 70-417. For now, you just need to know that this alternative authentication service is not AD DS and does not rely on the same set of technologies. Windows Azure Active Directory is most appropriate for applications that are designed to be cloud-based.

Here are the steps required to deploy a domain controller in Windows Azure after you create a Windows Azure account and log on to the Windows Azure Management Portal at <http://manage.windowsazure.com>:

1. Add an affinity group. An *affinity group* is one of the Windows Azure-related concepts that you're most likely to see on the 70-417 exam. An affinity group is a small physical area within one of the Microsoft data centers ("regions") in which you can place VMs. Creating an affinity group will reduce latency among any virtual machines you later add to that affinity group.
 - A. To create an affinity group in the Windows Azure Management Portal, select Settings in the blue menu on the left, select Affinity Groups on the right, and then click Add ("+") on the bottom menu. These options are highlighted in Figure 4-5.

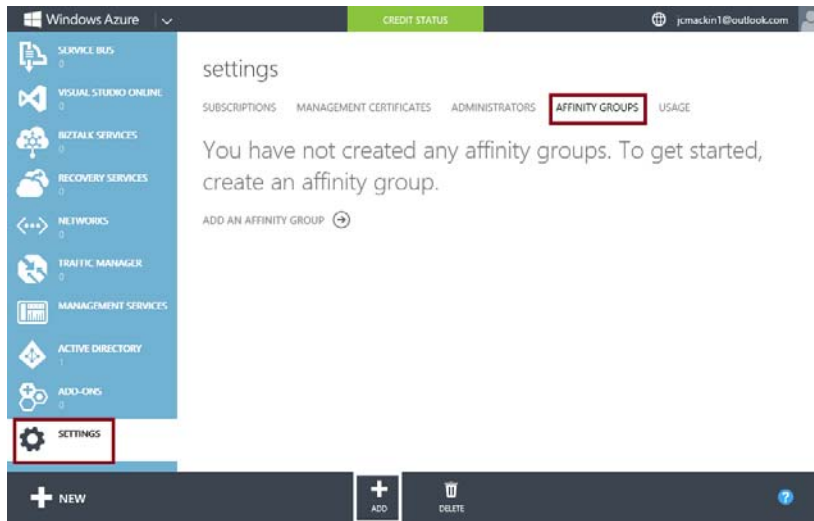


FIGURE 4-5 The first step to adding a new affinity group

- B. When you create a new affinity group, you have to specify a name and the Microsoft regional data center in which you want the affinity group to reside, as shown in Figure 4-6.

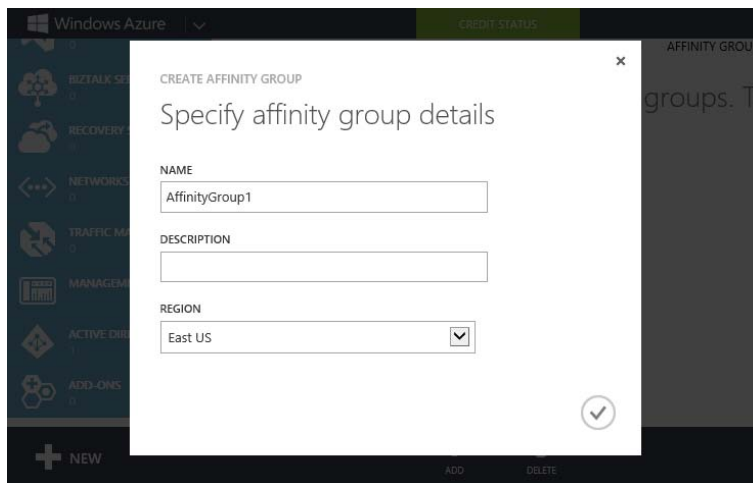


FIGURE 4-6 The second step to adding a new affinity group

2. Create a storage account. Before you create the VM in Windows Azure that will act as a domain controller, you need to create a storage account. Storage accounts in Windows Azure define the data redundancy type for the VMs later associated with the account. There are three redundancy options.
 - *Locally Redundant Replication* ensures that every VM in the storage account is replicated synchronously to two other nodes in the same data center.
 - *Geo-Redundant Replication* adds a second level to this baseline redundancy by performing asynchronous replication to a remote data center hundreds of miles away. At the remote, secondary location, three additional copies of the storage are kept.
 - *Read Access - Geo Redundant Replication* is like the Geo-Redundant Replication option, except that the data that is replicated to the remote location is stored as read-only.
- A. To create a new storage account in the Windows Azure Management Portal, click Storage in the blue menu on the left. On the bottom menu, click +New, as shown in Figure 4-7.

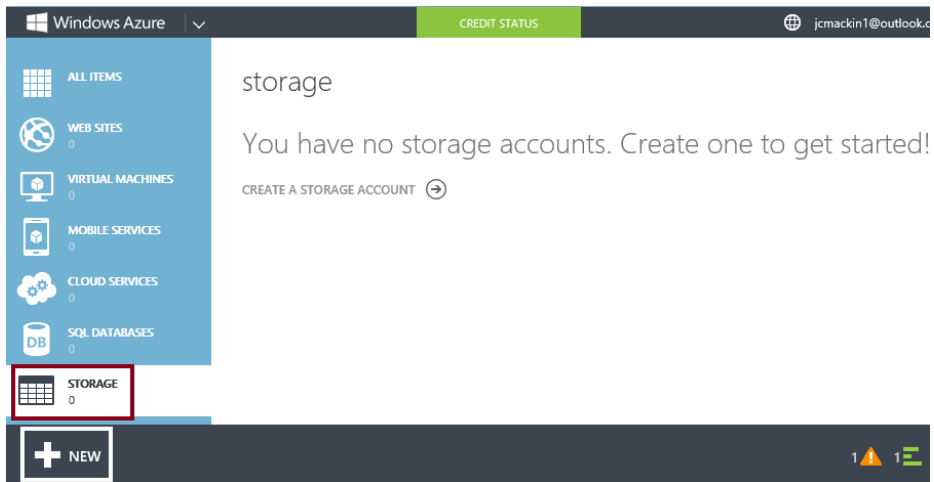


FIGURE 4-7 The first step to creating a storage account

- B. In the menu that expands, select Quick Create. This step opens configuration options on the right, as shown in Figure 4-8. You must then specify a unique URL along with an affinity group and choice of replication.

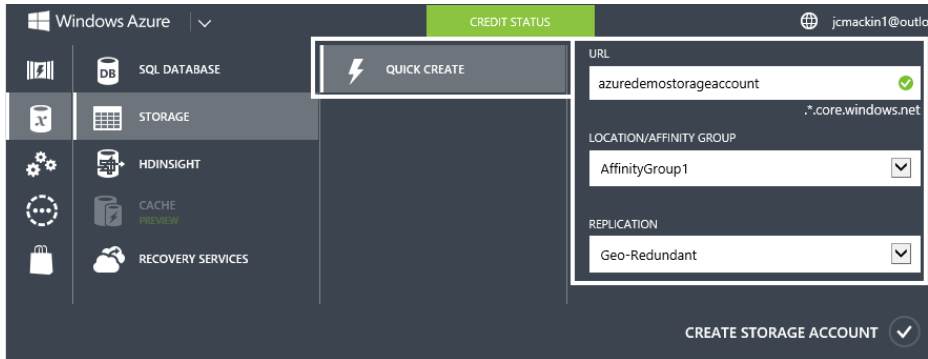


FIGURE 4-8 The second step to creating a storage account

3. Register a DNS server in Windows Azure. To lay the groundwork for the new virtual network that will contain your domain controller in Windows Azure, you first need to register a DNS server address. The new virtual network will use this DNS server. Later, this same internal IP address will be assigned automatically to the first VM in that virtual network.
 - A. To register the DNS server address in the Windows Azure Management Portal, select Networks in the blue menu on the left. Then click +New in the menu on the bottom, as shown in Figure 4-9.

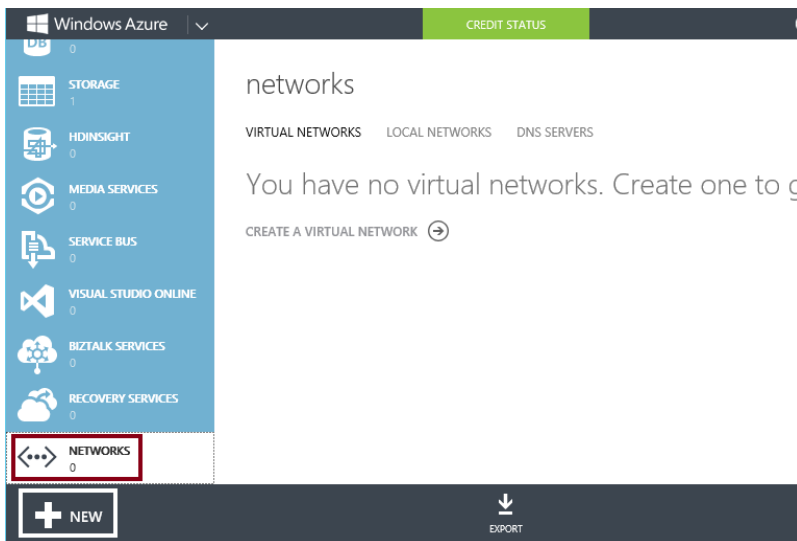


FIGURE 4-9 Registering a DNS server through the Networks settings

- B. Clicking New expands the bottom menu. When you choose Register DNS Server, configuration options that let you choose the name and address of the server appear on the right, as shown in Figure 4-10.

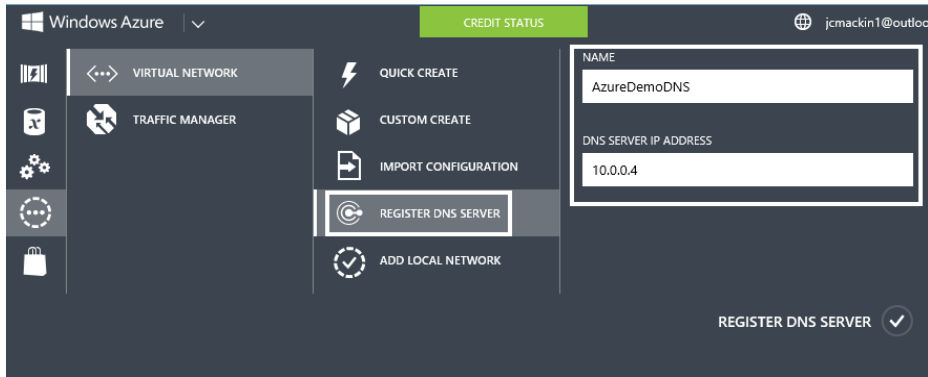


FIGURE 4-10 Registering a DNS server for a virtual network

4. Create a virtual network. A virtual network assigns IPv4 addressing information to an affinity group. When you create the virtual network, you choose a private IPv4 address space (10.---.---.---, 172.16.---.---, or 192.168.---.---), a maximum VM count (defined by a choice of subnet mask), and a DNS server address you have previously registered. Your virtual network will also include automatically configured subnets. Physical connectivity is automatically configured among the members of each subnet.

To create a new virtual network, select Network in the blue menu on the left in the Windows Azure Management Portal and then click New. In the expanded menu that appears, click Quick Create. This step opens the configuration options for the virtual network, as shown in Figure 4-11.

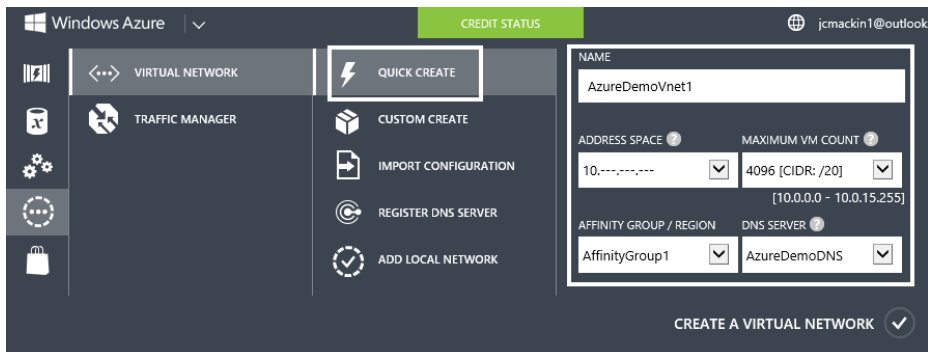


FIGURE 4-11 Creating a new virtual network in Windows Azure

5. Create a new Windows Server 2012 R2 virtual machine. The previous steps create the foundation for the VM that will act as the domain controller, and now you can finally create that VM and assign it to the virtual network and storage accounts you have previously defined.

- A.** To begin creating the VM, select Virtual Machines in the blue menu on the left in the Windows Azure Management Portal and then click +New on the bottom menu. In the expanded bottom menu that opens, click From Gallery, as shown in Figure 4-12. (Note that at the time of this writing, if you choose the Quick Create option, the VM is assigned to a new storage account with a randomly-generated name, not to your previously created storage account.)

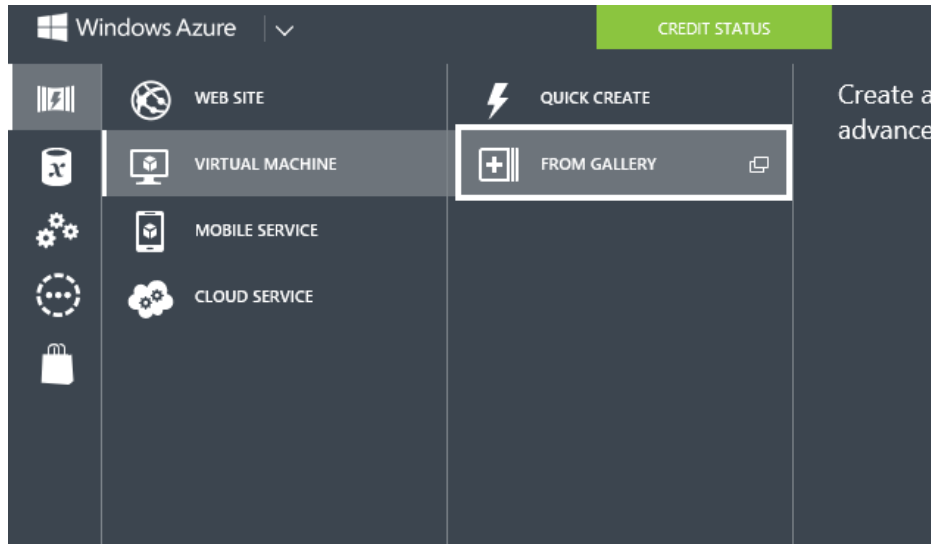


FIGURE 4-12 The first step to creating a new virtual machine in Windows Azure

- B.** At this point, you can select Windows Server 2012 R2 from the gallery of operating system options. (You won't have to install the OS yourself on the new VM.) You will then be given an opportunity to specify a name for the VM, along with an administrator account name and password.
- C.** Next you will see the page in Figure 4-13. Here is where you can assign the new VM to a virtual network, subnet, and storage account. You will also have to assign the VM to a cloud service, which you can create with a name of your choice. You probably don't need to understand cloud services for the 70-417 exam, but for now, just know that each cloud service is assigned one public IP address.

CREATE A VIRTUAL MACHINE

Virtual machine configuration

CLOUD SERVICE ?

Create a new cloud service

CLOUD SERVICE DNS NAME

AzureDemoCloudService .cloudapp.net

REGION/AFFINITY GROUP/VIRTUAL NETWORK ?

AzureDemoVnet1

VIRTUAL NETWORK SUBNETS

Subnet-1(10.0.0.0/23)

STORAGE ACCOUNT

azuredemostorageaccount

AVAILABILITY SET ?

(None)

FIGURE 4-13 The second step to creating a new virtual machine in Windows Azure

6. Connect to the VM and install the domain controller.
 - A. After you create the VM, you can connect to it by selecting it in the Windows Azure Management Portal and then clicking Connect. This step opens up a Remote Desktop connection to the VM, and you can then log on to the server and interact with it through its desktop.
 - B. Once you can connect to the VM, you can proceed with installing the domain controller by following the directions at the beginning of this chapter. However, there is one key difference between installing a domain controller on-premises and in Windows Azure. In Windows Azure, *you leave the VM with an automatically assigned address*. Believe it or not, all servers in Windows Azure should be configured as DHCP clients. The reason this is possible is that DHCP clients in Windows Azure are given extremely long leases, on the order of 99 years or more. In addition, the first server in a virtual network will be assigned the DNS server address that you registered in step 3, so you have in this sense already chosen the address of the domain controller. That address will not change over the lifetime of the VM as long as you perform maintenance shutdowns of the system from within Windows. If you choose the option to shut down the VM from the Windows Azure Management Portal, however, the DHCP lease will be renewed upon restarting, opening the possibility that the address could change.

NOTE If you were to assign a static address to a server, you could eventually lose connectivity to that server because of routing changes internal to Windows Azure.

Connecting a Windows Azure domain controller to AD DS on-premises

The previous section described the process of creating a domain controller in a new forest in Windows Azure. Alternatively, you might also want to connect the domain controller to an existing forest or domain located on the company premises. This configuration requires a VPN connection between the domain controller hosted in Azure and the on-premises site. Both site-to-site and point-to-site VPN types can be configured in Windows Azure, but only a point-to-site VPN can be configured on an existing virtual network. A point-to-site VPN, incidentally, is an option in Windows Azure that lets you to set up a VPN connection between an individual computer and a Windows Azure network without the use of a VPN device. However, a site-to-site VPN is more suitable for domain controller connectivity.

The steps involved in configuring VPNs in Windows Azure is beyond the scope of the 70-417 exam, but you should simply know that you need to configure a VPN connection to connect the domain controller hosted in Windows Azure to an existing AD DS on-prem.



EXAM TIP

Microsoft has announced that Windows Azure is being renamed Microsoft Azure. You can expect the name change to be reflected in the exams when the exams are updated.

Objective summary

- Windows Server 2012 and Windows Server 2012 R2 have new procedures for installing a domain controller. To install a domain controller, first install the Active Directory Domain Services server role. You can accomplish this either by using the Add Roles And Features Wizard or by typing the following at an elevated Windows PowerShell prompt:

```
Install-WindowsFeature -Name AD-Domain-Services -IncludeManagementTools
```
- In the GUI, after you install the Active Directory Domain Services server role, you can choose the option to promote the server to a domain controller that appears both at the end of the Add Roles And Features Wizard and in the notification menu in Server Manager.
- In Windows PowerShell, to install the first domain controller in a new forest, use the Install-ADDSTForest cmdlet. To install the first domain controller in a new domain in an existing forest, use the Install-ADDSDomain cmdlet. To install a new domain controller in an existing domain, use the Install-ADDSDomainController cmdlet.
- Windows Server 2012 and Windows Server 2012 R2 allow you to perform an IFM installation of a domain controller without first performing an offline defrag of the Active Directory database. To achieve this, in the IFM subcommand menu of the Ntdsutil utility, use either the Create Full NoDefrag parameter or the Create Sysvol Full No Defrag parameter.

- To deploy a domain controller in Windows Azure, you first need to create an affinity group, create a storage account, pre-register the DNS server address you are going to assign to the domain controller, create a virtual network where the domain controller will be hosted, and finally create the VM with the Windows Server operating system of your choice. The domain controller in Windows Azure should remain a DHCP client. To connect the domain controller to an existing Active Directory forest, you need to configure the virtual network with VPN connectivity.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the “Answers” section at the end of the chapter.

1. You are a network administrator for Fabrikam.com. You want to set up a new domain that is completely separate from the existing company network. The purpose of the domain is to test functionality and application compatibility in Windows Server 2012 R2.

You have installed Windows Server 2012 R2 on the first server and now want to promote it to a domain controller. Which of the following two commands or cmdlets should you run? (Choose two. Each answer is part of the solution.)

- A. Install-ADDSDomain
 - B. Install-ADDSDomainController
 - C. Install-ADDSEForest
 - D. Install-WindowsFeature
 - E. Dcpromo.exe
2. You are a network administrator for Fabrikam.com. You want to add a new domain controller to the Fabrikam.com domain by using the Installation from Media (IFM) option. You select a storage device with 12 GB of space to store the NTDS database and the SYSVOL folder. The contents of the NTDS database amount to 100 MB, and the SYSVOL folder amounts to 30 MB.

At an elevated command prompt on a Windows Server 2012 R2 domain controller, you enter the Ntdsutil utility and activate the instance of NTDS. You then enter the IFM subcommand menu.

Your goal is to write both the NTDS database and SYSVOL to your storage media, and to achieve this as quickly as possible. Which parameter should you use?

- A. Create Full NoDefrag
- B. Create Sysvol Full NoDefrag
- C. Create Full
- D. Create Sysvol Full

- 3.** You have installed Windows Server 2012 R2 on a new server named RODC1. You now want to use Windows PowerShell to make RODC1 a Read-Only Domain Controller (RODC) in the Fabrikam.com domain.

Which of the following Windows PowerShell cmdlets is *not* necessary to run in order to achieve your goal?

- A.** Install-ADDSDomain
 - B.** Install-WindowsFeature
 - C.** Install-ADDSDomainController
 - D.** Add-ADDSDomainControllerAccount
- 4.** You are a system administrator for Contoso. The Contoso company network includes a single Active Directory domain named Contoso.com. Currently, all servers are located on the company premises.

The company now wants to deploy two new servers in Windows Azure: an application server and a domain controller. The domain controller should be a member of the Contoso.com domain and will provide authentication for the application server.

You create an account for your organization and sign in to the Windows Azure Management Portal.

Which of the following steps should you take? (Choose all that apply.)

- A.** Create an affinity group.
- B.** Assign a static address to the VM that will be the domain controller.
- C.** Create a site-to-site VPN to connect the servers to the company network.
- D.** Manually install Active Directory Domain Services on the VM that will be the domain controller.



Thought experiment

Configuring Hyper-V at Contoso

Your company, Contoso.com, is experiencing rapid growth.

As a way to provide scalability to meet the IT demands of continued growth, the company is building a new virtual infrastructure. The IT department has purchased eight physical servers to host virtualized instances of your servers and of client machines. All eight servers are running Windows Server 2012 R2 with the Hyper-V server role installed. All guest servers are currently running Windows Server 2008 R2.

A company goal is to move all of Contoso's workloads to the virtual infrastructure.

Currently, the network includes a single domain named Contoso.com. Users have complained that login times are slow, especially at the beginning of the workday.

You can find the answers to these questions in the "Answers" section.

- 1.** You create a new virtual machine and install Windows Server 2012 R2. You name the server DC2012A. You now want to promote DC2012A to a new domain controller in the Contoso.com domain. What complete Windows PowerShell command should you run on DC2012A to install the server as an additional domain controller in the Contoso.com domain?
- 2.** You want to promote a computer named DC2012B, which is also running Windows Server 2012 R2. You want DC2012B to be a Read-Only Domain Controller. Which Windows PowerShell cmdlets should you use to create an RODC account for DC2012B and to make it an RODC?
- 3.** The research department has requested the creation of their own domain. You install Windows Server 2012 R2 on a new virtual machine and name the computer ResDC. You now want to promote ResDC to be the first domain controller in a new domain named research.contoso.com. What complete Windows PowerShell command should you run on ResDC to achieve this goal?

Answers

This section contains the answers to the Objective Reviews and the Thought Experiment.

Objective 4.1: Review

1. Correct answers: C, D

- A. Incorrect:** You would use this cmdlet to install the first domain controller in a new domain, but only within an existing forest. You want the new domain to be completely separate from the existing company network.
- B. Incorrect:** You would use this cmdlet to add an additional domain controller to an existing domain.
- C. Correct:** You need to use `Install-ADDSForest` cmdlet because you want to create a new domain that is completely separate from the company network. This cmdlet will automatically promote the server to a domain controller in the new domain.
- D. Correct:** Before using the `Install-ADDSForest` cmdlet, you need to add the Active Directory Domain Services server role. This cmdlet helps you accomplish that goal. The full syntax of the command is:

```
Install-WindowsFeature -Name AD-Domain-Services -IncludeManagementTools
```

- E. Incorrect:** This command can be used only with an answer file in Windows Server 2012 and Windows Server 2012 R2. There is no indication that an answer file is available, and if there were, it would most likely provide a configuration for the Fabrikam.com network, from which the new domain controller must be kept separate.

2. Correct answer: B

- A. Incorrect:** This parameter would copy the NTDS database but not the SYSVOL folder to the media.
- B. Correct:** This parameter would write both the NTDS database and SYSVOL to the media. In addition, it would achieve the result as quickly as possible because it avoids the step of offline defragmentation of this data. You have ample storage space on your media, so offline defragmentation is not necessary.
- C. Incorrect:** This parameter would copy the NTDS database but not the SYSVOL folder to the media.
- D. Incorrect:** This parameter would write both the NTDS database and SYSVOL to the media. However, it would not achieve the result as quickly as possible because offline defragmentation will be performed.

3. Correct answer: A

- A. Correct:** This cmdlet would help you install the first domain controller in a new domain in an existing forest. You cannot use it to install an RODC because the first domain controller in a domain cannot be an RODC.
- B. Incorrect:** You need to use this cmdlet to install Active Directory Domain Services before you install the domain controller.
- C. Incorrect:** You need to use this cmdlet to add an additional domain controller to an existing domain. You would use the `-ReadOnlyReplica` parameter to make the new domain controller an RODC. Perform this step after you have created a computer account for the RODC.
- D. Incorrect:** You would use this cmdlet to create a computer account for the RODC in the domain.

4. Correct answers: A, C, D

- A. Correct:** You want to create an affinity group for the two servers because you want them to be in close physical proximity to each other.
- B. Incorrect:** You shouldn't assign static addresses to VMs in Windows Azure.
- C. Correct:** You need to create a VPN connection to allow communication between the DC in Windows Azure and the domain controllers on the company premises.
- D. Correct:** You should perform the installation of Active Directory Domain Services on the VM as if it were hosted on the local premises.

Thought experiment

1. `Install-ADDSDomainController -DomainName " contoso.com "`
2. To create the RODC account, use the `Add-ADDSDomainControllerAccount` cmdlet. To turn the computer into an RODC, use the `Install-ADDSDomainController` cmdlet with the `-ReadOnlyReplica` parameter.
3. `Install-ADDSDomain -NewDomainName research -ParentDomainName contoso.com`

Deploy, manage, and maintain servers

This chapter contains a single objective: Monitor servers. The vast majority of this topic is probably already familiar to you: Data Collector Sets, event forwarding and collection, alerts, performance counters, and network monitoring: None of these features has changed in any significant way since Windows Server 2008. That's not to say you shouldn't brush up on these topics if you don't feel confident about them. You absolutely should.

But if the question is *what's new* in server monitoring in Windows Server 2012 and Windows Server 2012 R2, there's really just one development you need to know about, and it relates to Resource Metering. We've already covered Resource Metering of virtual machines in Chapter 3, "Configure Hyper-V," and this chapter builds on that knowledge by introducing the related topic of virtual machine (VM) resource pools.

Objectives in this chapter:

- Objective 5.1: Monitor servers

Objective 5.1: Monitor servers

Windows Server 2012 and Windows Server 2012 R2 introduce Windows PowerShell cmdlets to measure virtual machine resource pools, allowing you to monitor usage by all VMs of various resources on a Hyper-V host. The VM resource metering introduced in Chapter 3 is used to measure resource usage by individual VMs. VM resource pools, in contrast, are more useful for measuring usage of a host's resources by all guest VMs.

For the exam, you should learn how to configure and use this new feature.

This section covers the following topics:

- Virtual machine resource pools
- Monitoring servers through Windows PowerShell
- Reviewing older monitoring features

Virtual machine resource pools

Each Hyper-V host includes predefined VM resource pools that correspond to the individual host resource types—such as processor, Ethernet, memory, and disk (VHD)—that are available to the guest VMs on that host. You can use these resource pools to monitor aggregate usage of each resource by all VMs, or you can create new resource pools that measure some subset of a resource, such as a particular set of VHDs. To see the list of all VM resource pools, type **Get-VMResourcePool** at a Windows PowerShell prompt. The following shows a sample output:

| Name | ResourcePoolType | ParentName | ResourceMeteringEnabled |
|------------|------------------------|--------------|-------------------------|
| Primordial | Processor | | False |
| Primordial | Ethernet | | False |
| Primordial | FibreChannelConnection | | False |
| FCSAN | FibreChannelConnection | {Primordial} | False |
| Primordial | FibreChannelPort | | False |
| Primordial | Memory | | False |
| Primordial | VHD | | False |
| Primordial | ISO | | False |
| Primordial | VFD | | False |

Predefined resource pools are named “Primordial” because they represent the total amount of a resource available on the host machine, as opposed to a user-defined subset of that resource. The `ResourceMeteringEnabled` status of all the resource pools in the preceding example is labeled `False` because Resource Metering hasn’t yet been enabled.

Metering virtual machine resource pools

As with VM Resource Metering, covered in Chapter 3, resource pool metering is intended only to provide raw data to be captured by applications developed either in-house or by third-party developers. Still, you can use Windows PowerShell to test the functionality of VM resource pools and view usage data associated with them.

To begin metering usage of a resource pool, first use the `Enable-VMResourceMetering` cmdlet. Specify the resource pool by name with the `-ResourcePoolName` parameter and by type with the `-ResourcePoolType` parameter. For example, to enable metering of the primordial memory resource pool, type the following:

```
Enable-VMResourceMetering -ResourcePoolName Primordial -ResourcePoolType Memory
```

You can also use the wildcard “*” symbol in place of a specific name and ignore the `-ResourcePoolType` parameter if instead you want to enable all resource pools that can possibly be enabled:

```
Enable-VMResourceMetering -ResourcePoolName *
```

A predefined resource pool can still remain disabled after you run this previous command only if no corresponding resource is found on the host computer. After you run the preceding command, the output of Get-VMResourcePool in this example thus changes to the following:

| Name | ResourcePoolType | ParentName | ResourceMeteringEnabled |
|------------|------------------------|--------------|-------------------------|
| Primordial | Processor | | True |
| Primordial | VHD | | True |
| Primordial | Ethernet | | True |
| Primordial | FibreChannelPort | | False |
| Primordial | FibreChannelConnection | | False |
| Primordial | VFD | | False |
| Primordial | Memory | | True |
| Primordial | ISO | | False |
| FCSAN | FibreChannelConnection | {Primordial} | False |

You can measure the resource usage associated with a resource pool after metering is enabled on it. To do so, use the Measure-VMResourcePool cmdlet. For example, the following command will provide usage for network data:

```
Measure-VMResourcePool -Name * -ResourcePoolType Ethernet
```

Sample output is shown here in an abbreviated format:

| Name | ResourcePoolType | NetworkInbound(M) | NetworkOutbound(M) |
|------------|------------------|-------------------|--------------------|
| Primordial | {Ethernet} | 20 | 4 |

Creating virtual machine resource pools

Instead of measuring primordial pool usage, you can use the New-VMResourcePool cmdlet to create resource pools for monitoring and reporting usage of a subset of resources.

For example, the following command creates a new VM resource pool associated with multiple VHD files:

```
PS C:\> New-VMResourcePool "New Resource Pool" VHD -Paths "D:\Hyper-V\Virtual Hard Disks"
```

| Name | ResourcePoolType | ParentName | ResourceMeteringEnabled |
|-------------------|------------------|--------------|-------------------------|
| New Resource Pool | VHD | {Primordial} | False |

As with the primordial resource pools, you can then enable Resource Metering on new VM resource pools and measure their usage.

Monitoring servers through Windows PowerShell

The 70-417 exam is likely to test your knowledge of Windows PowerShell far more than did the exams you took to earn your Windows Server 2008 certification. For this reason, it's a good idea to review the cmdlets related to monitoring.

Two of these cmdlets are new to Windows Server 2012 and Windows Server 2012 R2 and they both relate to *virtual machine eventing*, a minor feature that is still not well documented at the time of this writing. According to Get-Help and Windows PowerShell documentation at <http://technet.microsoft.com/en-us/library/hh848462>, “[v]irtual machine eventing keeps Hyper-V PowerShell objects updated without polling the virtual machine host.” You should also be aware that it is enabled by default, and you can use the cmdlets `Enable-VMEventing` and `Disable-VMEventing` to reenable and disable the feature, respectively.

Use Table 5-1 to review some of the more important Windows PowerShell cmdlets and other commands that relate to server monitoring. (More cmdlets are available to manage events. For a full list, type **get-command *event* | sort noun,verb** at a Windows PowerShell prompt.)

TABLE 5-1 Common command-line tools for server monitoring

| Cmdlet or command-line utility | Description |
|--------------------------------|--|
| Export-Counter | Exports data that is returned by the Get-Counter and Import-Counter cmdlets |
| Get-Counter | Gets performance counter data from local and remote computers |
| Import-Counter | Imports performance counter log files (.blg, .csv, .tsv) and creates the objects that represent each counter sample in the log |
| Get-Event | Gets events in the Windows PowerShell event queue for the current session |
| New-Event | Creates a new custom event |
| Clear-EventLog | Deletes all entries from specified event logs on the local or remote computers |
| Get-EventLog | Manages event logs and displays events contained within those event logs |
| Enable-VMEventing | Enables virtual machine eventing |
| Disable-VMEventing | Disables virtual machine eventing |
| Logman.exe | Manages and schedules performance counter and event trace log collections on local and remote systems |



EXAM TIP

The Deploy, Manage, and Maintain Servers domain is taken from the 70-411 exam. In its original form, it includes two additional objectives covering Windows Deployment Services (WDS) and Windows Server Update Services (WSUS), respectively. Even though these objectives aren't officially indicated for the 70-417 exam, don't be surprised if you see a question about either of these topics on this test. For WDS, you can review its features new to Windows Server 2012 at <http://technet.microsoft.com/en-us/library/hh974416.aspx> and its features new to Windows Server 2012 R2 at <http://technet.microsoft.com/en-us/library/dn281955.aspx>. For WSUS, you can install the server role and browse the options available through the Options node. Pay special attention to the Update Files options described at [http://technet.microsoft.com/en-us/library/cc708431\(v=ws.10\).aspx](http://technet.microsoft.com/en-us/library/cc708431(v=ws.10).aspx).

Reviewing older monitoring features

More than with other objectives, you should expect the Monitor servers objective on the 70-417 exam to be represented by questions about older topics that have not changed since Windows Server 2008. The topic of monitoring virtual machines (as distinct from Resource Metering in Hyper-V) is otherwise too narrow to account for a proportional representation of this objective on the exam. For this reason, make sure you review the older topics also indicated for this objective on the official exam page on the Microsoft website. (You can visit this page at <http://www.microsoft.com/learning/en-us/exam-70-417.aspx>.) These topics include configuring Data Collector Sets (DCS), configuring alerts, monitoring real-time performance, monitoring events, configuring event subscriptions, configuring network monitoring, and scheduling performance monitoring. These are all topics you have already learned well enough to achieve your current certification.

In this section, we'll point out two of these older topics that are particularly important for review.

Creating a Data Collector Set manually

When you create a new Data Collector Set in Performance Monitor in Windows Server 2012 or Windows Server 2012 R2, just as in Windows Server 2008, you are provided with two options: to create the new Data Collector Set from a template, or to create it manually.

If you choose to create the Data Collector Set manually, you are prompted to specify the type of data you want to include. You are provided with the following four options, which have not changed since Windows Server 2008:

- **Performance Counter** This option uses any set of performance counters you specify to gather and save data about the system's performance.
- **Event Trace Data** This option provides information about activities and system events, as opposed to performance counters.
- **System Configuration Information** This option allows you to record the value of specified registry keys as they change.

- **Performance Counter Alert** This option allows you to configure an action to take place (such as running a program) when a selected performance counter crosses a threshold value you specify.

Make sure you remember the purpose of all four of these data collection types for the 70-417 exam.

MORE INFO For additional information about creating a Data Collector Set manually, visit <http://technet.microsoft.com/en-us/library/cc766404.aspx>.



EXAM TIP

Be prepared to answer questions in which you need to diagnose performance problems based on performance counter data. Remember the following guidelines:

- **Memory\Pages/sec** A sustained value of higher than 5 suggests a possible memory bottleneck
- **PhysicalDisk\%Disk Time** A sustained value of higher than 50% suggests a possible disk bottleneck
- **Processor\%Processor Time** A sustained value of higher than 65% suggests a possible CPU bottleneck
- A combination of high memory and processor usage, or a set of seemingly contradictory measurements, suggests a malfunctioning process.

Configuring and managing Data Collector Sets

After you create a new Data Collector Set, you can modify its settings and options. (These settings have not changed since Windows Server 2008.) For example, you can use the Data Manager settings on the shortcut (Actions) menu of a Data Collector Set to specify how (but not where) data is stored for each Data Collector Set. For example, you can configure Data Manager settings to automatically delete the data collected by a Data Collector Set when that data reaches a certain age or folder size.

Another option on the shortcut (Actions) menu of a Data Collector Set is the Save Template option. The Save Template option lets you export a Data Collector Set as an .xml file, which you can then import later on another computer.

For the 70-417 exam, also be sure to review the Data Collector Set properties tabs. You use the Directory tab, for example, to set a directory to store performance log data. You use the Stop Condition tab to determine when the Data Collector Set should stop running.

MORE INFO For more information about Data Manager settings, you can visit <http://technet.microsoft.com/en-us/library/cc765998.aspx>.

Objective summary

- In Windows Server 2012 and Windows Server 2012 R2, you can meter usage of resource pools (such as memory, processor, or Ethernet) on a VM host.
- Use the `Get-VMResourcePool` cmdlet to display a list of all resource pools on a Hyper-V host running Windows Server 2012 or Windows Server 2012 R2.
- Hyper-V hosts running Windows Server 2012 or Windows Server 2012 R2 have pre-configured resource pools, each named `Primordial`, that correspond to the aggregate amount of a resource on the host. You can create new resource pools that correspond to a subset of these primordial resource pools.
- To configure Resource Metering of a resource pool, use the `Enable-VMResourceMetering` cmdlet and specify the name and type of the resource pool with the `-ResourcePoolName` and `-ResourcePoolType` parameters, respectively. Then, to view usage of the resource pool, use the `Measure-VMResourcePool` cmdlet.
- Use the `Enable-VMEventing` and `Disable-VMEventing` cmdlets to enable and disable the feature called virtual machine eventing. Virtual machine eventing keeps Hyper-V PowerShell objects updated without polling the VM host.
- More than with other objectives officially listed for the 70-417 exam, it is important to review the older features about monitoring servers (such as Data Collector Sets and Event Subscriptions) that you learned for your Windows Server 2008 certification.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the “Answers” section at the end of the chapter.

1. The Hyper-V server role has been installed on a server named `HYPV01`, which is hosting 10 VMs and has been configured with 32 GB of RAM. You want to determine average usage of the host server RAM over time by all of your guest VMs together. You have not yet configured Resource Metering of any resource pool or of any VMs. You want to view the usage of the primordial RAM resource pool on `HYPV01`, and you want to do so with the greatest amount of administrative efficiency. Which of the following Windows PowerShell cmdlets should you use to achieve your goal? (Choose all that apply.)
 - A. `Enable-VMResourceMetering`
 - B. `New-VMResourcePool`
 - C. `Measure-VM`
 - D. `Measure-VMResourcePool`

2. The Hyper-V server role has been installed on a server named HYPV02, which is hosting six VMs and has been configured with 16 GB of RAM. One of the guest VMs that is running on HYPV02 is named DBSrv1. DBSrv1 has been allocated 8 GB of RAM. You are concerned that more RAM than necessary has been allocated to DBSrv1. You want to measure average RAM usage by DBSrv1 during a period of high load, and you want to do so from a command-line interface on HYPV02. You have not yet configured Resource Metering of any resource pool or of any VMs. Which of the following Windows PowerShell commands or cmdlets can you use to measure RAM usage by DBSrv1? (Choose all that apply.)
- A. `Enable-VMResourceMetering -ResourcePoolName DBSrv1 -ResourcePoolType Memory`
 - B. `Enable-VMResourceMetering -VMName DBSrv1`
 - C. `Measure-VM`
 - D. `Measure-VMResourcePool`
3. You want to ensure that Hyper-V PowerShell objects are updated without polling the VM host. Which Windows PowerShell cmdlet can you use to achieve this goal?
- A. `Enable-VMEventing`
 - B. `Register-EngineEvent`
 - C. `Register-ObjectEvent`
 - D. `Get-EventSubscriber`



Thought experiment

Monitoring servers at Contoso

You are a network administrator for Contoso.com. One of Contoso's application servers is named App1, which is running a Server Core installation of Windows Server 2012 R2. Company policy dictates that App1 be managed remotely only, and only through Windows PowerShell. None of the firewall exceptions have been enabled on App1 that would be required for remote management through graphical administration consoles.

You can find the answers to these questions in the "Answers" section.

1. Which Windows PowerShell cmdlet should you use to review individual event logs such as the System log on App1?
2. Which command-line utility can you use to create daily counter collection queries on App1?
3. Which Windows PowerShell cmdlet should you use to delete all entries from a specified event log on App1?
4. Which Windows PowerShell cmdlet should you use to display current performance counter data from App1?

Answers

This section contains the answers to the Objective Review and the Thought Experiment.

Objective 5.1: Review

1. Correct answers: A, D

- A. Correct:** If Resource Metering has not yet been configured, you will need to enable it on the primordial RAM resource pool before you can measure its usage.
- B. Incorrect:** You don't need to create a new resource pool because you want to measure the usage of a primordial resource pool, which is built-in.
- C. Incorrect:** You would use the Measure-VM cmdlet if you wanted to measure usage of a particular VM, not of a particular resource pool. You could use the command Measure-VM * to display average RAM usage by each individual VM, but this is not the most efficient solution because it would require you then to add these values together to derive the aggregate RAM usage.
- D. Correct:** To view total usage of a resource pool by all VMs, use the Measure-VM-ResourcePool cmdlet.

2. Correct answers: B, C

- A. Incorrect:** This command would enable Resource Metering of a memory resource pool named DBSrv1. The question does not state that a resource with such a name has been created. In addition, you couldn't use such a command to measure RAM usage by any one VM unless only one VM is running (which is not the case in this question).
- B. Correct:** This command would enable metering of all resources (including RAM) by the VM named DBSrv1.
- C. Correct:** This cmdlet allows you to measure resource usage by a particular VM. In this case, you would enter the command Measure-VM DBSrv1.
- D. Incorrect:** This cmdlet allows you to measure resource usage of a particular resource pool by all guest VMs. You cannot use it to measure resource usage by just one VM unless only one VM is running (which is not the case in this question).

3. Correct answer: A

- A. Correct:** The Enable-VMEventing cmdlet enables virtual machine eventing, which keeps Hyper-V PowerShell objects updated without polling the VM host. Virtual machine eventing is enabled by default, but running this command lets you ensure that it is enabled.
- B. Incorrect:** The Register-EngineEvent cmdlet enables you to subscribe to events that are generated by the Windows PowerShell engine and by the New-Event cmdlet. It doesn't allow you to keep Hyper-V PowerShell objects updated without polling the VM host.
- C. Incorrect:** The Register-ObjectEvent cmdlet enables you to subscribe to the events that are generated by a Microsoft .NET Framework object. It doesn't allow you to keep Hyper-V PowerShell objects updated without polling the VM host.
- D. Incorrect:** The Get-EventSubscriber cmdlet gets the event subscribers in the current session. It doesn't allow you to keep Hyper-V PowerShell objects updated without polling the VM host.

Thought experiment

1. Get-EventLog. (You would normally use parameters to limit the output. For example, the command `Get-EventLog System -Newest 25` retrieves only the most recent 25 events in the System event log.)
2. Logman.exe
3. Clear-EventLog
4. Get-Counter

Configure network services and access

The Configure Network Services and Access domain is another with just one objective tested on the 70-417 exam: Configure DirectAccess. DirectAccess is an improved alternative to a VPN that was first introduced in Windows Server 2008 R2 and Windows 7. If you earned your Windows Server 2008 MCSA before the release of Windows Server 2008 R2, you might have missed this major new technology completely. And if you did learn about DirectAccess in Windows Server 2008 R2, you need to know that this feature has changed significantly in Windows Server 2012 and Windows Server 2012 R2.

Objectives in this chapter:

- Objective 6.1: Configure DirectAccess

Objective 6.1: Configure DirectAccess

DirectAccess in Windows Server 2008 R2 and Windows 7 was a very promising technology that was also difficult to configure. Beginning with Windows Server 2012 and Windows 8, the infrastructure requirements of DirectAccess have been simplified along with the configuration steps. At the same time, its feature set has expanded considerably.

For the 70-417 exam, you first need to understand basic DirectAccess concepts and components. You also need to know how the infrastructure requirements to support DirectAccess clients differ to support various features. Finally, you need to know how to configure DirectAccess.

This section covers the following topics:

- DirectAccess infrastructure options
- Configuring DirectAccess clients
- Configuring DirectAccess servers
- Configuring DirectAccess infrastructure servers

What is DirectAccess?

DirectAccess is an always-on remote access technology that is based on IPv6 communication. Through DirectAccess, a user's computer automatically, transparently, and securely connects to a private corporate network from any location in the world as soon as the computer is connected to the Internet. When a DirectAccess connection is active, remote users connect to resources on the corporate network as if they were on the local premises.

DirectAccess overcomes the limitations of VPNs by providing the following benefits:

- **Always-on connectivity** Unlike with a VPN, a DirectAccess connection is always on, even before the user logs on to his or her computer.
- **Seamless connectivity** To the user, the DirectAccess connection to the corporate network is completely transparent and resembles an always-on VPN connection.
- **Bidirectional access** With DirectAccess, the user's remote computer not only has access to the corporate intranet, but the intranet can also see the user's computer. This means that the remote computer can be managed by using Group Policy and other management tools (such as System Center 2012 R2 Configuration Manager ["Configuration Manager" for short]) in exactly the same way that computers located on the internal network are managed.

In addition, DirectAccess includes the following security features:

- DirectAccess uses IPsec to authenticate both the computer and user. If you want, you can require a smart card for user authentication.
- DirectAccess also uses IPsec to provide encryption for communications across the Internet.

Understanding IPv6 and DirectAccess

A DirectAccess connection from a remote client to an internal resource includes two legs. In the first half of the connection, the DirectAccess client always uses IPv6 to initiate contact with the DirectAccess server, typically found at the edge of the private network. IPv6 transition technologies are used to assist this connection when necessary. The second half of the connection occurs between the DirectAccess server and the internal network resource. This part of the connection can proceed either over IPv4 (only if the DirectAccess server is running Windows Server 2012 or Windows Server 2012 R2 and acting as a NAT64/DNS64 device) or over IPv6.

Figure 6-1 shows the two legs of a DirectAccess connection between a remote client and an internal network resource.




FIGURE 6-1 A DirectAccess connection to an internal resource

First leg: External client to private network edge

If the DirectAccess client can obtain a global IPv6 address from its environment, then the connection to the DirectAccess server proceeds over the IPv6 Internet in a straightforward manner. However, IPv6 is not widely implemented yet on public networks, so three IPv6 transition technologies are used to assist in establishing the IPv6 connection to the DirectAccess server. If all three of the following transition technologies are enabled on the client through Group Policy, they are attempted in the following order of preference:

1. **6to4** For DirectAccess clients that have a *public* IPv4 address, 6to4 can be used to connect to the DirectAccess server via IPv6 across the public IPv4 Internet. 6to4 achieves this by tunneling or encapsulating IPv6 data within an IPv4 header, in a technique known as IPv6-over-IPv4. 6to4 requires any intervening router or firewall to be configured so that outbound traffic for Protocol 41 is allowed. Also note that 6to4 does not work if the client is behind a network address translation (NAT) device.
2. **Teredo** For DirectAccess clients behind a NAT device and configured with a *private* IPv4 address, Teredo can be used to connect to the DirectAccess server via IPv6 across the public IPv4 Internet. Like 6to4, Teredo tunnels IPv6 traffic in IPv4. The intervening routers and firewalls must be configured to allow outbound traffic through User Datagram Protocol (UDP) port 3544.
3. **IP-HTTPS** For DirectAccess clients that cannot effectively establish IPv6 connectivity to the DirectAccess server through 6to4 or Teredo, IP-HTTPS is used. With IP-HTTPS, DirectAccess clients encapsulate IPv6 traffic within HTTPS traffic. Virtually all routers allow outbound HTTPS traffic, so this option is almost always possible.

NOTE Since the release Windows Server 2012 and Windows 8, the performance of IP-HTTPS is close to that of Teredo because a “null encryption” option is used for HTTPS communication. However, in Windows Server 2008 R2 and Windows 7, IP-HTTPS used Secure Sockets Layer (SSL) encryption on top of the IPsec encryption that was used to secure the connection between the DirectAccess client and server. This “double encryption” significantly degraded network performance.

Second leg: Private network edge to internal resource

Between the network edge and the internal network resource, the connection can proceed over either IPv6 or IPv4. You don’t have to deploy global IPv6 on your internal network because Windows Server 2012 and Windows Server 2012 R2 can act as a NAT64/DNS64 device when deployed as a DirectAccess server at the network edge. (A NAT64/DNS64 device translates between IPv6 and IPv4.) However, an all-IPv6 connection still provides the best performance and is the preferred scenario.

NOTE Windows Server 2008 R2 didn’t provide NAT64/DNS64 functionality, but you could use Microsoft Forefront Unified Access Gateway 2010 or a third-party device to provide NAT64/DNS64 translation. Otherwise, to implement DirectAccess, you had to deploy global IPv6 on your internal network or use the IPv6 transition technology ISATAP. You can still use ISATAP in Windows Server 2012 and Windows Server 2012 R2, but it is not recommended.

Understanding the DirectAccess connection process

A DirectAccess connection to a target intranet resource is initiated when the DirectAccess client connects to the DirectAccess server through IPv6. IPsec is then negotiated between the client and server. Finally, the connection is established between the DirectAccess client and the target resource.

This general process can be summarized in the following steps:

1. The DirectAccess client computer attempts to connect to an internal computer configured as the *network location server*. If the network location server is available, the DirectAccess client determines that it is already connected to the intranet, and the DirectAccess connection process stops. If the network location server is not available, the DirectAccess client determines that it is connected to the Internet and the DirectAccess connection process continues.

NETWORK LOCATION SERVER A network location server is an intranet web server that a DirectAccess client attempts to access as a way to determine whether the client is located on the intranet or Internet. An internal address of the DirectAccess server can be configured as the network location server, but using a separate, high-availability internal web server for the network location server is preferred. If you configure a separate web server as a network location server, the web server does not have to be dedicated to this one service.

2. The DirectAccess client computer connects to the DirectAccess server by using IPv6 and IPsec. If a native IPv6 network isn't available, the client establishes an IPv6-over-IPv4 tunnel by using 6to4, Teredo, or IP-HTTPS. The user does not have to be logged in for this step to complete.
3. As part of establishing the IPsec session, the DirectAccess client and server authenticate each other by using Kerberos or computer certificates.
4. By validating Active Directory Domain Services group memberships, the DirectAccess server verifies that the computer and user are authorized to connect using DirectAccess.
5. If Network Access Protection (NAP) is enabled and configured for health validation, the DirectAccess client obtains a health certificate from a Health Registration Authority (HRA) located on the Internet prior to connecting to the DirectAccess server. The HRA forwards the DirectAccess client's health status information to a NAP health policy server. The NAP health policy server processes the policies defined within the Network Policy Server (NPS) and determines whether the client is compliant with system health requirements. If so, the HRA obtains a health certificate for the DirectAccess client. When the DirectAccess client connects to the DirectAccess server, it submits its health certificate for authentication.
6. The DirectAccess server begins forwarding traffic from the DirectAccess client to the intranet resources to which the user has been granted access.

Understanding DirectAccess infrastructure options

You can deploy DirectAccess in a number of network scenarios, ranging from the very simple to the very complex. A few of these options are illustrated in the examples that follow.

Simple DirectAccess infrastructure

A simple DirectAccess infrastructure includes a DirectAccess server that is running Windows Server 2012 or Windows Server 2012 R2 and is deployed at the network edge. This DirectAccess server is configured as a Kerberos proxy and NAT64/DNS64 translation device. The external interface is configured with a public IP address. (Two would be necessary to support Teredo.) The internal address is associated with the network location server.

Within the internal network is a domain controller/DNS server and at least one internal network resource such as a file server or application server. Note that this simple infrastructure supports only Windows 8 clients and later because Windows 7 clients do not support Kerberos authentication for DirectAccess.

Figure 6-2 shows a simple DirectAccess infrastructure.

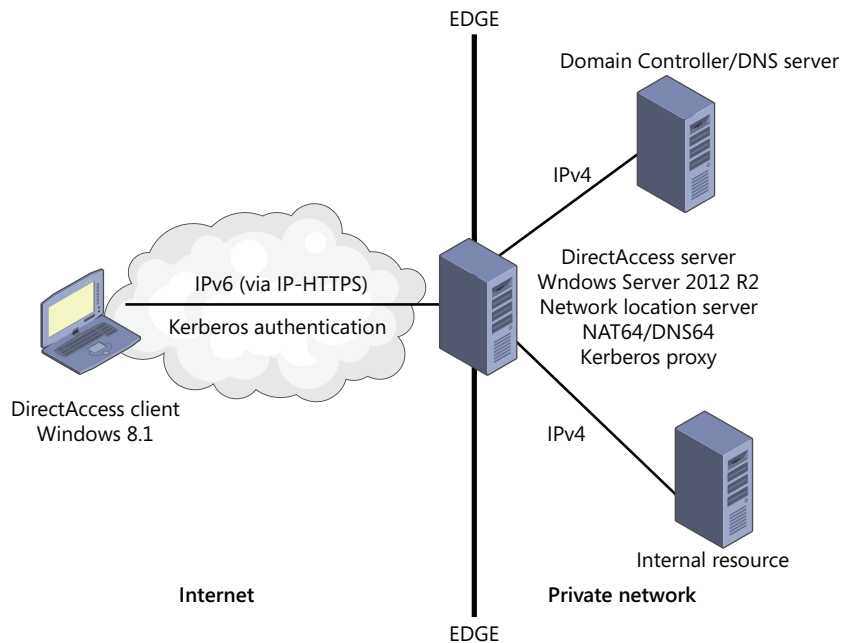


FIGURE 6-2 A simple DirectAccess infrastructure



EXAM TIP

Remember that only Windows Server 2012 and later, and Windows 8 and later, support Kerberos proxy, a feature that greatly simplifies authentication for DirectAccess clients. In addition, only Windows Server 2012 and Windows Server 2012 R2 include built-in support for NAT64/DNS64 translation, which lets you use DirectAccess with your existing internal IPv4 infrastructure.

DirectAccess server behind NAT

Beginning with Windows Server 2008 R2, all versions of DirectAccess allow you to deploy a DirectAccess server behind the network edge in a perimeter network. However, only in Windows Server 2012 and Windows Server 2012 R2 can you deploy the DirectAccess server behind a NAT device. In such a scenario, the DirectAccess server needs only a single network adapter and a single address. Connections from the DirectAccess clients through the NAT device to the DirectAccess server are established through IP-HTTPS.

Figure 6-3 illustrates a DirectAccess network topology in which a DirectAccess server is deployed behind a NAT device.

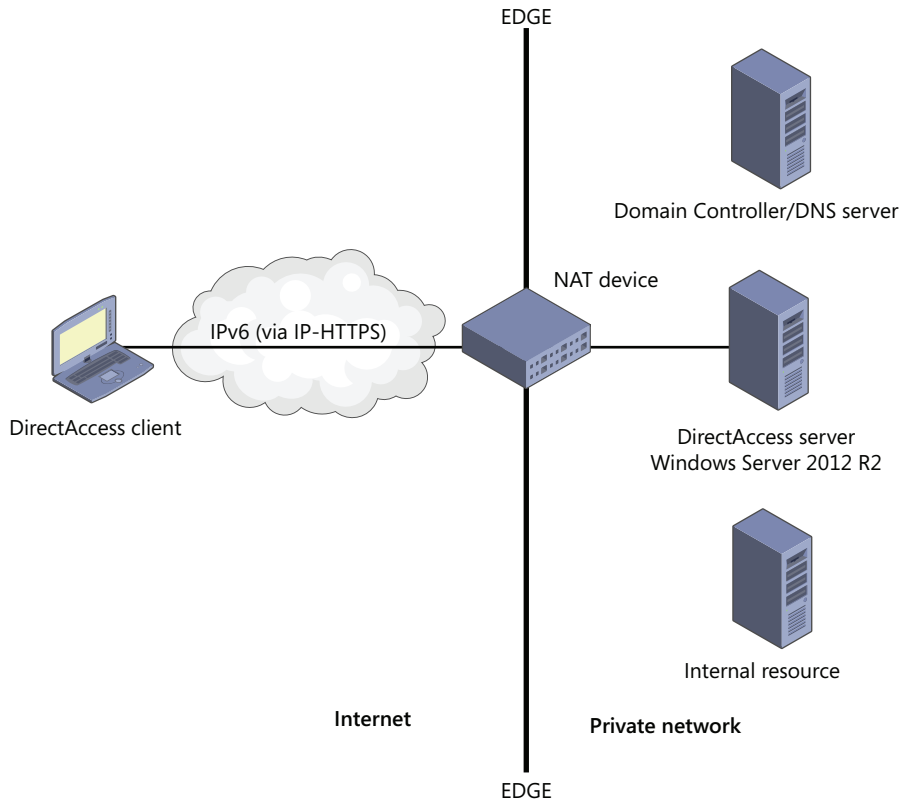


FIGURE 6-3 DirectAccess server deployed behind a NAT device

Multisite/Multidomain DirectAccess infrastructure

Another infrastructure option that is new to Windows Server 2012 and Windows Server 2012 R2 is the ability for DirectAccess to be deployed across multiple sites. A multisite deployment of DirectAccess requires a public key infrastructure (PKI) and computer authentication through certificates. In addition, *multidomain* support is now a built-in feature of DirectAccess that requires no extra configuration.

When you configure a multisite deployment, the DirectAccess clients are provided with a list of the DirectAccess servers that act as entry points to the private network at each site. Before connecting, DirectAccess clients running Windows 8 or later then ping each of these DirectAccess servers. The Windows 8 or later client then initiates contact with the server whose latency is determined to be the shortest. (Windows 7 clients in a multisite deployment simply use a single, preconfigured DirectAccess server address.)

Figure 6-4 shows a multisite DirectAccess infrastructure.

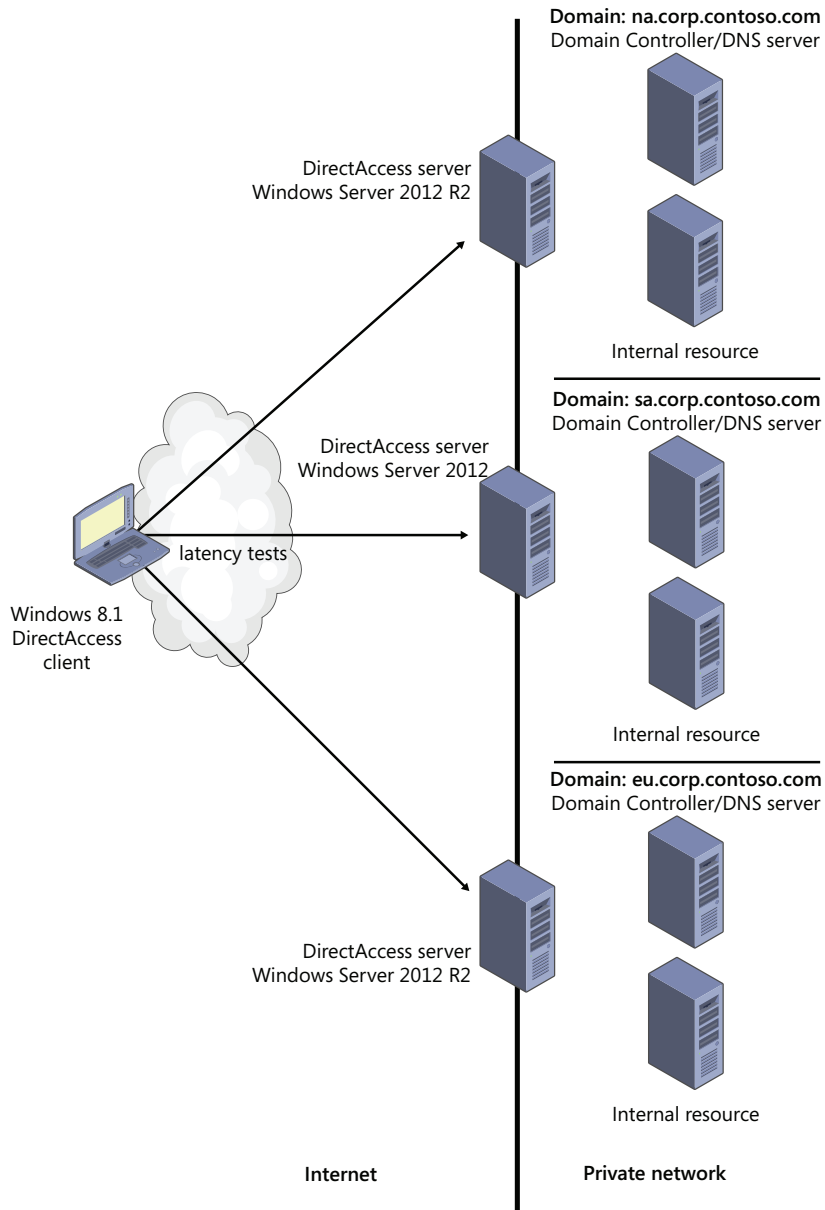


FIGURE 6-4 A multisite DirectAccess infrastructure

NOTE To enable multisite capability in DirectAccess, in the Remote Access Management Console click Enable Multisite (as shown later in this chapter in Figure 6-20), or use the Windows PowerShell cmdlet Enable-DAMultiSite.

Complex DirectAccess infrastructure

Although Windows Server 2012 and Windows Server 2012 R2 greatly simplify the basic infrastructure requirements for DirectAccess, the infrastructure can become complex if you need functionality that is not included by default in a basic setup. For example, one new feature introduced in Windows Server 2012 is the ability to deploy DirectAccess servers in a Network Load Balancing (NLB) cluster. This functionality naturally adds complexity to your infrastructure but is often necessary to support many remote clients. Another requirement that adds more elements to your infrastructure is the need to support Windows 7 clients. Windows 7 clients can authenticate DirectAccess connections only with computer certificates, so your DirectAccess infrastructure would require a PKI in such a scenario. Next, DirectAccess can also be deployed with NAP, which is another factor that adds complexity but might be required by your IT policies. Additional features such as two-factor authentication with one-time passwords (OTPs) would raise the infrastructure requirements even more. Figure 6-5 illustrates a more complex DirectAccess infrastructure that supports all three IPv6 transition technologies, improves load capacity with an NLB cluster, supports Windows 7 clients with a PKI/certification authority, includes a NAP infrastructure, and has a network location server deployed apart from the DirectAccess server.

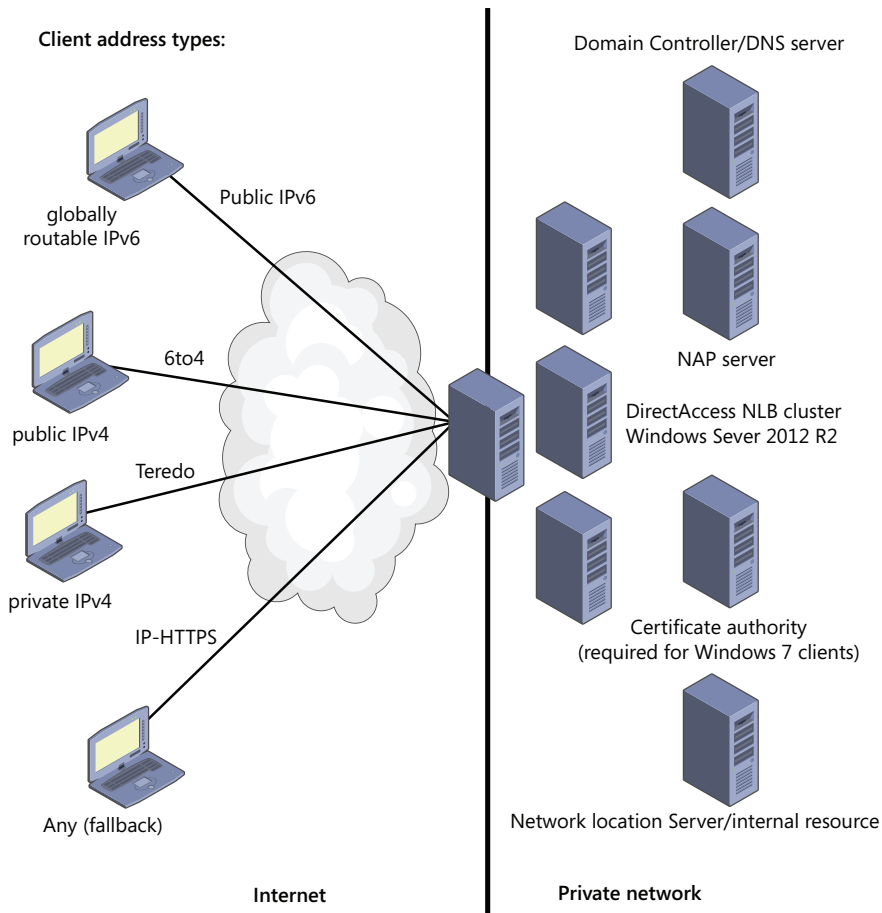


FIGURE 6-5 A complex DirectAccess infrastructure

NOTE To enable load balancing in DirectAccess, in the Remote Access Management Console click Enable Load Balancing (as shown later in this chapter in Figure 6-20) or use the Windows PowerShell cmdlets `Set-RemoteAccessLoadBalancer` and `Add-RemoteAccessLoadBalancerNode`.

Installing and configuring DirectAccess

Windows Server 2012 and Windows Server 2012 R2 have greatly simplified the process of installing and configuring DirectAccess. DirectAccess is now unified with traditional VPNs in a new Remote Access server role and managed with the same tool, the Remote Access Management Console. In fact, you can now configure a Windows Server to act as both a DirectAccess server and a traditional VPN server at the same time, an option that was not possible in Windows Server 2008 R2. Even more significant than unified management are the new configuration wizards first introduced in Windows Server 2012 that make the process of deploying and configuring DirectAccess and VPNs relatively easy.

Installing DirectAccess

DirectAccess now belongs to the Remote Access server role. You can install the DirectAccess component of the Remote Access role through the Add Roles and Features Wizard or by typing the following at an elevated Windows PowerShell prompt:

```
Install-WindowsFeature DirectAccess-VPN -IncludeManagementTools
```

You can then configure DirectAccess using the Remote Access Management Console, shown in Figure 6-6, or by using Windows PowerShell commands.

MORE INFO To review the cmdlets used to configure DirectAccess, visit <http://technet.microsoft.com/en-us/library/hh918399> or type the following at a Windows PowerShell prompt when the “DirectAccess and VPN (RAS)” component of the Remote Access role is installed:

```
Get-Command -Module RemoteAccess *da*
```

Note also that installing the Remote Access role and its role management tools installs the Windows PowerShell module named `DirectAccessClientComponents`, which provides the additional client cmdlets listed at <http://technet.microsoft.com/en-us/library/hh848426>.

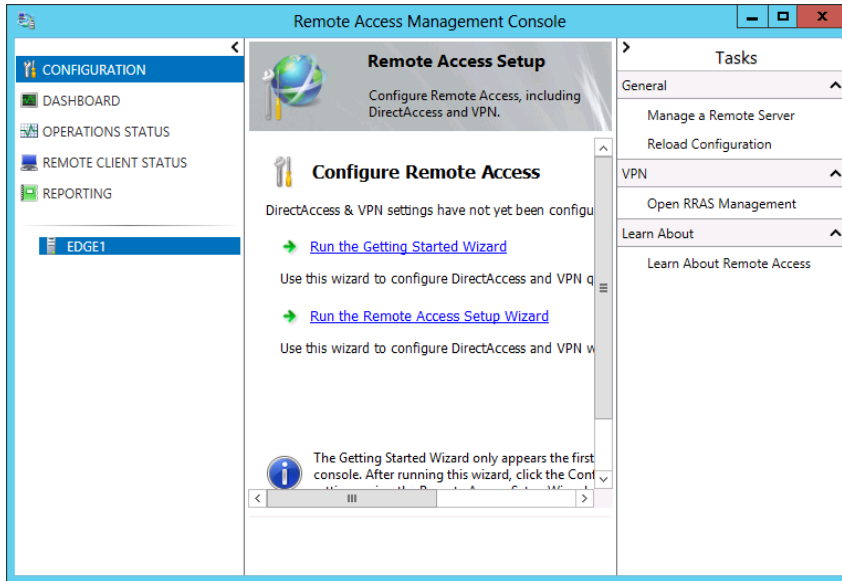


FIGURE 6-6 The Remote Access Management Console provides a unified configuration and management tool for all remote access technologies

Configuring DirectAccess

Figure 6-6 shows the Remote Access Management Console before you take any configuration steps. The central pane shows two options for wizards: the Getting Started Wizard and the Remote Access Setup Wizard. Whichever wizard you choose to begin configuration, you are next presented with an option to configure just DirectAccess, just a VPN, or both, as shown in Figure 6-7.

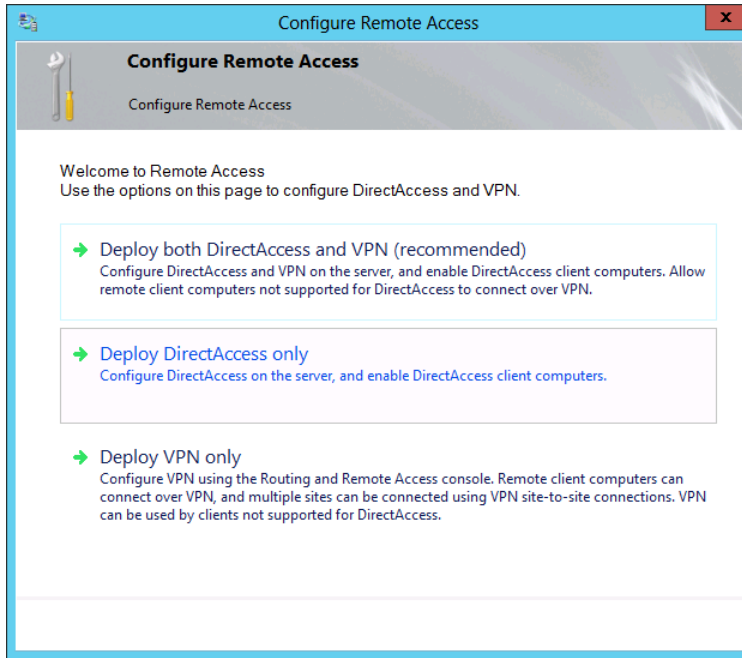


FIGURE 6-7 The Remote Access configuration wizards allow you to configure just DirectAccess, just a VPN, or both.

The Getting Started Wizard, which was first introduced in Windows Server 2012, is an excellent tool that helps you deploy a remote access solution quickly. However, it is not especially useful for exam preparation precisely because it hides the very configuration options you need to know and understand for the test. In addition, VPN configuration has not changed since you earned your Windows Server 2008 MCSA in any way that is significant for the 70-417 exam. For these reasons, to prepare for the Configure DirectAccess objective for the 70-417 exam, you should focus on configuration options that appear after you click Run The Remote Access Setup Wizard shown in Figure 6-6 and then click Deploy DirectAccess Only shown in Figure 6-7.

After you click Deploy DirectAccess Only, the Remote Access Management Console reappears with the center pane replaced by an image similar to the one shown in Figure 6-8. The four steps in the map are associated with four configuration wizards that you must complete in order: The first is for configuring DirectAccess clients, the second is for configuring the DirectAccess server, the third is for configuring infrastructure servers, and the fourth is for configuring the application servers (if desired). These wizards create and configure Group Policy Objects (GPOs) for DirectAccess servers and clients.

It's unlikely you'll be asked to identify any of these wizards by name on the 70-417 exam. However, you might be asked about *any configuration option* that appears in any of these four wizards. These four wizards, then, provide a useful way to organize the new configuration options that you need to learn and understand for the 70-417 exam, so we will look at them in order.

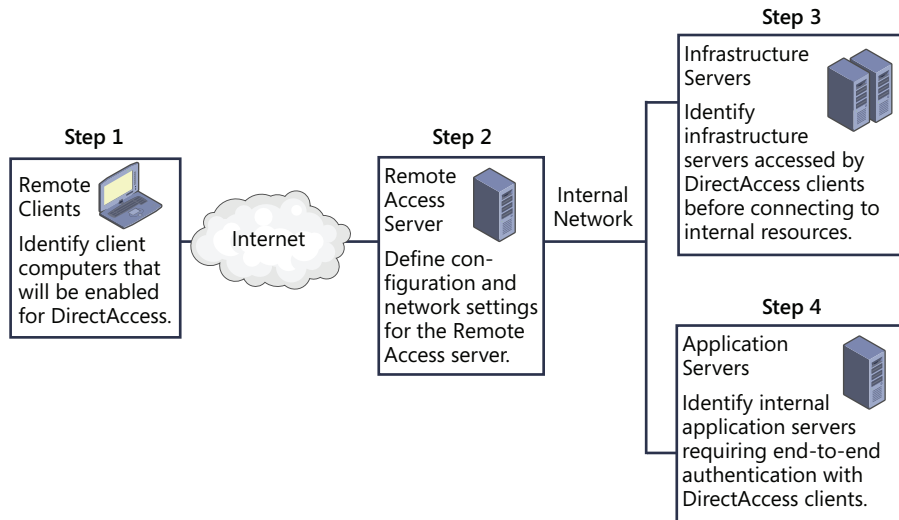


FIGURE 6-8 The four DirectAccess configuration wizards

STEP 1: DIRECTACCESS CLIENT SETUP

The first page of the DirectAccess Client Setup Wizard is shown in Figure 6-9. This Deployment Scenario page allows you the option to configure DirectAccess clients either for both remote access and remote management, or for remote management only. The first, default option configures bidirectional communication for DirectAccess servers and clients. Choosing the second option, however, would allow administrators to manage remote DirectAccess clients through tools such as Configuration Manager, but it would prevent those clients from accessing the internal corporate network. Note that the option to deploy DirectAccess clients for remote management only is new to Windows Server 2012 and Windows Server 2012 R2.




FIGURE 6-9 The Deployment Scenario page of the DirectAccess Client Setup Wizard

NOTE To choose the deployment scenario in Windows PowerShell, use the `Set-DAServer` cmdlet with the `-DAInstallType` parameter and either the `FullInstall` or `ManageOut` value. For example, to configure the DirectAccess deployment for remote management only, type the following at an elevated Windows PowerShell prompt on the DirectAccess server:

```
Set-DAServer -DAInstallType ManageOut
```

The second page of the DirectAccess Client Setup Wizard is the Select Groups page, shown in Figure 6-10. The first function of this page is to let you specify the security groups containing the computer accounts that you want to enable for DirectAccess. This is an important step to remember: No DirectAccess client is allowed access to the internal network if you don't assign that client the right to do so. To perform this task in Windows PowerShell, use the `Add-DAClient` cmdlet with the `-SecurityGroupNameList` parameter.

A second option on this page is to enable DirectAccess for mobile computers only. Interestingly, this option is selected by default if you run the Getting Started Wizard. Computers connecting remotely through DirectAccess are most likely to be mobile computers, but there are exceptions, and these exceptions could easily form the premise of an exam question. (Scenario: Some users working on domain-joined desktop computers from remote sites can't connect through DirectAccess. Why not? The option to enable DirectAccess for mobile computers only is selected.)



EXAM TIP

If only laptops are able to connect through DirectAccess, you can change this setting by modifying the the DirectAccess Client Settings GPO. Specifically, you need to remove the DirectAccess - Laptop Only WMI filter that is linked to this GPO in the Security Filtering settings. Note also that both creating and removing these WMI filters requires Domain Admin permissions.

The third option on this page is Use Force Tunneling. This option forces the DirectAccess client to tunnel *all* network traffic through the private network, regardless of where that traffic is ultimately destined. This behavior, for example, could be used to ensure that all web traffic from DirectAccess clients passes through an internal web proxy server. In Windows PowerShell, this option is configured by using the Set-DAClient cmdlet with the -ForceTunnel parameter.



EXAM TIP

Expect to see the Use Force Tunneling option appear in an exam question about DirectAccess.

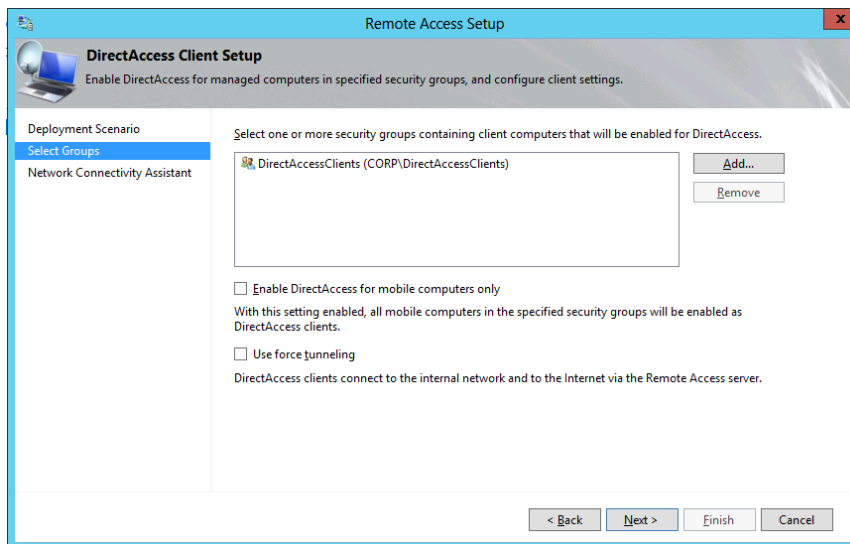


FIGURE 6-10 The Select Groups page of the DirectAccess Client Setup Wizard

The final page in the DirectAccess Client Setup Wizard is the Network Connectivity Assistant page, shown in Figure 6-11. The Network Connectivity Assistant is client software embedded in Windows 8 and later that determines whether DirectAccess is functioning. (This feature is not the same as the network location server, which helps a client determine whether it is on the Internet or intranet.)

The first setting on this page is the host address for the Network Connectivity Assistant. DirectAccess client computers use this address to verify that the client can successfully connect to the internal network. This setting is unlikely to appear on the 70-417 exam (except maybe as an incorrect answer choice), but if you need to enter this resource manually, you should specify the address of a corporate URL or FQDN that is always available to DirectAccess clients. Your internal DNS should resolve this address to the internal address of the Remote Access server.

The most testable setting on this page is the option to allow DirectAccess clients to use local name resolution. Local name resolution in this case refers to the broadcast-based protocols of NetBIOS over TCP/IP and Link-Local Multicast Name Resolution (LLMNR). When this option is enabled, DirectAccess clients are allowed to resolve single label names such as App1 using local name resolution if they can't be resolved through DNS. Local name resolution must also be configured in the Infrastructure Server Setup Wizard.

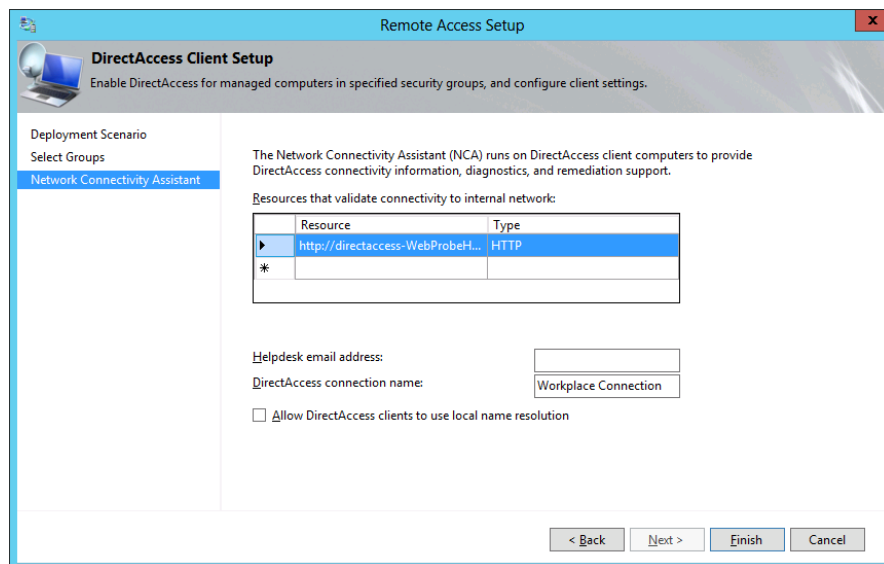


FIGURE 6-11 The Network Connectivity Assistant page of the DirectAccess Client Setup Wizard

STEP 2: REMOTE ACCESS SERVER SETUP

The first page of the Remote Access Server Setup Wizard is the Network Topology page, shown in Figure 6-12. This page lets you specify where in your network you are going to deploy your DirectAccess server.

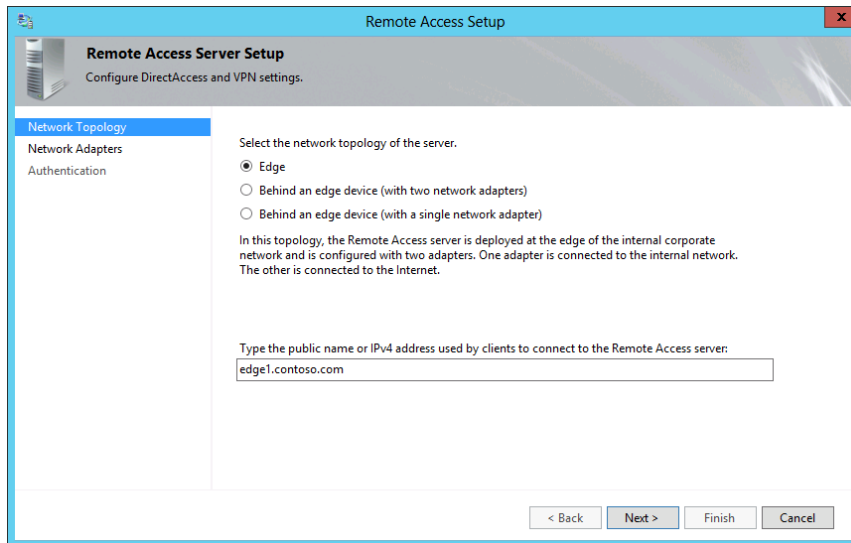


FIGURE 6-12 The Network Topology page of the Remote Access Server Setup Wizard

The first option is Edge. Choosing this option requires the DirectAccess server to be configured with two network adapters, one connected directly to the Internet and one connected to the internal network. The external interface needs to be assigned two consecutive public IPv4 addresses if you need to support Teredo.

The second option is Behind An Edge Device (With Two Network Adapters). Select this option if you want to deploy the DirectAccess server in a perimeter network behind a firewall or router. In this topology, the network adapter attached to the perimeter network is assigned one or two consecutive public IPv4 addresses, and the second adapter attached to the internal network can be assigned a private address.

The third option is Behind An Edge Device (With A Single Network Adapter). Choose this option if you want to deploy your DirectAccess server behind a NAT device. In this case, your DirectAccess server is assigned a single private IP address.

Finally, the Network Topology page requires you to specify the name or IPv4 address the DirectAccess clients will use to connect to the DirectAccess server. Be sure to specify a name that can be resolved through public DNS or an IPv4 address that is reachable from the public network.

The second page of the Remote Access Server Setup Wizard is the Network Adapters page, shown in Figure 6-13. This page requires you to choose the network adapter or adapters that will be assigned to internal network and external network, as required by your specified topology.

This page also requires you to specify a certificate that the DirectAccess server will use to authenticate IP-HTTPS connections. If your organization has deployed a PKI, you can browse to a copy of the computer certificate for the local server. If you don't have a PKI, you need

to choose the option to use a self-signed certificate instead. Note that the availability of this latter option was first introduced in Windows Server 2012 and could easily serve as the basis for a test question.

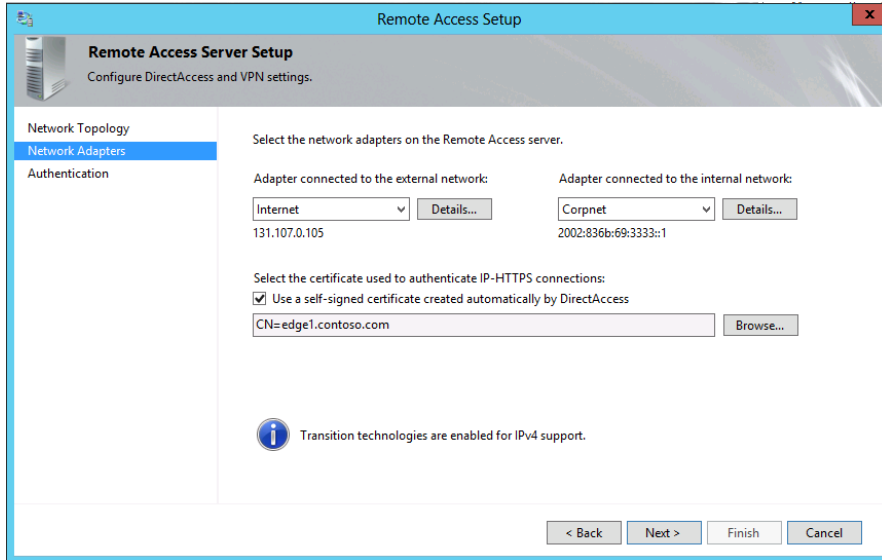


FIGURE 6-13 The Network Adapters page of the Remote Access Server Setup Wizard

The final page of the Remote Access Server Setup Wizard is the Authentication page, shown in Figure 6-14. This page lets you configure the following settings related to DirectAccess client authentication:

- **User Authentication** By default, users authenticate only with Active Directory credentials. However, you can choose the option here to require two-factor authentication. Typically, two-factor authentication requires a user to insert a smart card in addition to typing his or her Active Directory credentials. Note that in Windows Server 2012 and Windows Server 2012 R2, however, the Trusted Platform Module (TPM) of client computers can act as a virtual smart card for two-factor authentication. Alternatively, you can also configure two-factor authentication so that users must enter an OTP such as one provided through RSA SecurID in addition to their Active Directory credentials. OTP requires a PKI and RADIUS server, along with a number of configuration steps that you don't need to understand for the 70-417 exam. For the 70-417 exam, you merely need to know that OTP is an alternative to smart cards for two-factor authentication in DirectAccess.
- **Use Computer Certificates** If you configure DirectAccess in the GUI, client computers are authenticated through Kerberos by default. However, you can select an option to require computer authentication through the use of certificates. Computer certificate authentication is required to support two-factor authentication, a multisite deployment of DirectAccess, or Windows 7 DirectAccess clients.

- **Enable Windows 7 Client Computers To Connect Via DirectAccess** By default, Windows 7 client computers cannot connect to a Windows Server 2012 or Windows Server 2012 R2 Remote Access deployment. You need to enable that functionality here.
- **Enable Corporate Compliance For DirectAccess Clients With NAP** This page allows you to require a health check of client computers through NAP. To configure this setting in Windows PowerShell, use the Set-DAServer cmdlet with the -HealthCheck parameter.

MORE INFO The authentication settings displayed on this page can all be configured through the Set-DAServer cmdlet. For more information, use Get-Help or visit <http://technet.microsoft.com/en-us/library/hh918371>.

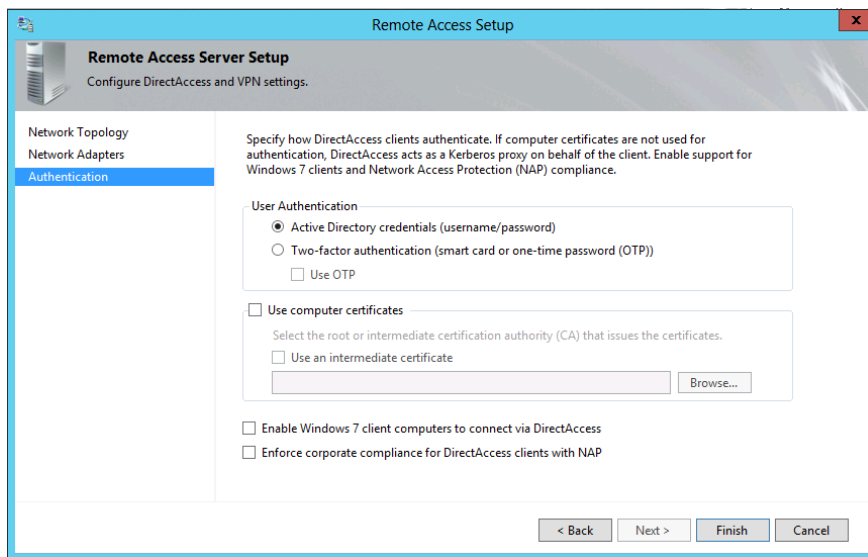


FIGURE 6-14 The Authentication page of the Remote Access Server Setup Wizard

STEP 3: INFRASTRUCTURE SERVER SETUP

The Infrastructure Server Setup Wizard allows you to configure settings related to the network location server, the DNS server, and management servers such as update or antivirus servers.

The first page of this wizard is the Network Location Server page, shown in Figure 6-15. As explained earlier in this chapter, DirectAccess clients use this server to determine whether they are on the company network. It's recommended that you use an internal web server other than the DirectAccess (Remote Access) server for this purpose. (The DNS address and associated IP address in this case is naturally associated with the interface attached to the internal network.) If you do specify the DirectAccess server as the network location server, it must be authenticated by a computer certificate—a self-signed one, if necessary. To configure the network location server using Windows PowerShell, use the Set-DANetworkLocationServer cmdlet.

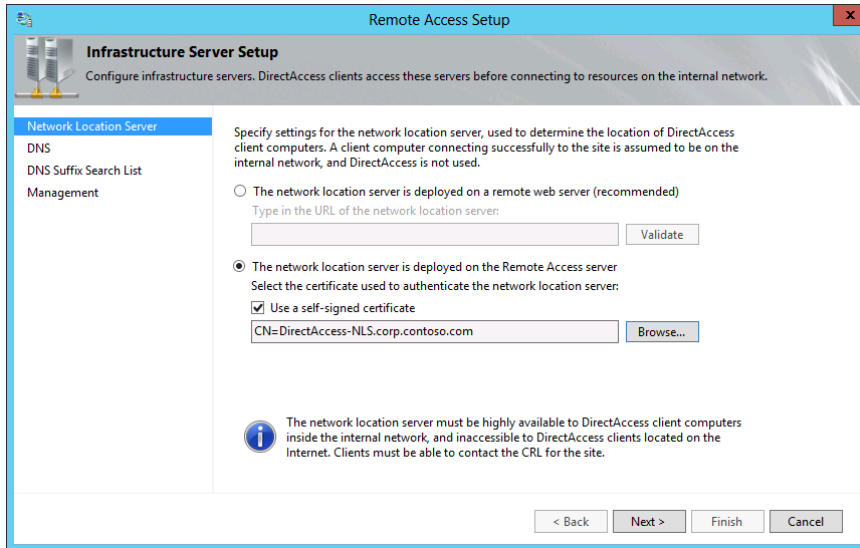


FIGURE 6-15 The Network Location Server page of the Infrastructure Server Setup Wizard

The second page of the Infrastructure Server Setup Wizard is the DNS page shown in Figure 6-16. The main function of this page is to allow you to configure the Name Resolution Policy Table (NRPT). The entries you create here are written to the GPO used to configure DirectAccess clients.

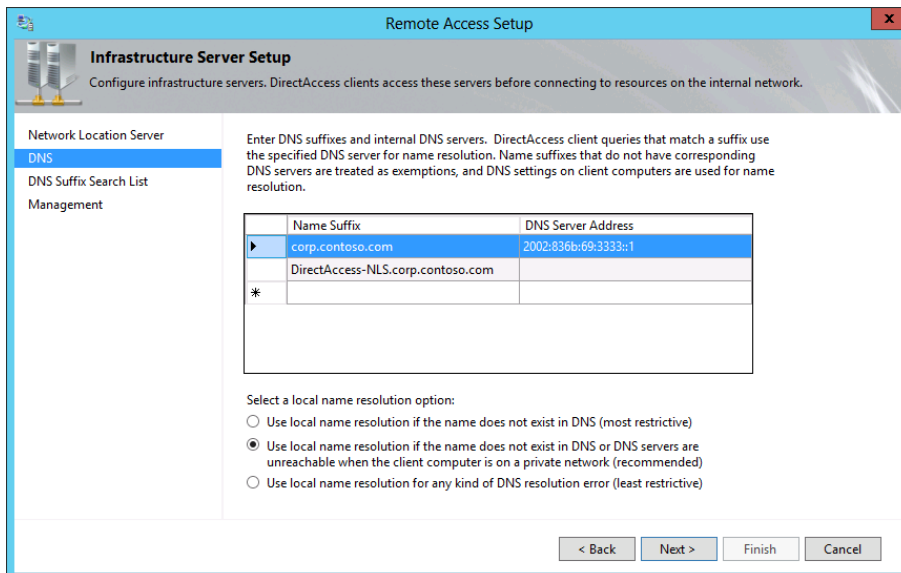


FIGURE 6-16 The DNS page of the Infrastructure Server Setup Wizard

The NRPT is a feature that allows a DNS client to assign a DNS server address to particular namespaces rather than to particular interfaces. The NRPT essentially stores a list of name resolution rules that are applied to clients through Group Policy. Each rule defines a DNS

namespace (a domain name or FQDN) and DNS client behavior for that namespace. Together, these name resolution rules are called a *Name Resolution Policy*. When a DirectAccess client is on the Internet, each name query request is compared against the namespace rules stored in the Name Resolution Policy. If a match is found, the request is processed according to the settings in the Name Resolution Policy rule. The settings determine the DNS servers to which each request will be sent. If a name query request does not match a namespace listed in the NRPT, it is sent to the DNS servers configured in the TCP/IP settings for the specified network interface.

You might need to configure Name Resolution Policy entries if, for example, you need to enable DNS clients to resolve DNS suffixes found only within your intranet namespace. Another reason might be if you have a split public/private DNS environment based on the same domain name, and you need to ensure that DirectAccess clients don't contact your company's public servers (such as a web server) through the DirectAccess connection.



EXAM TIP

You need to understand the function of a Name Resolution Policy and the NRPT for the 70-417 exam. Also know that you can view the NRPT by using the `Get-DnsClientNrptPolicy` cmdlet in Windows PowerShell.

The second configuration decision you need to make on the DNS page relates to the DirectAccess clients' use of local name resolution methods such as NetBIOS and LLMNR. Unlike the setting in the DirectAccess Client Setup Wizard, which merely allows (does not block) the use of local name resolution, the setting here determines how local name resolution will be used if allowed. You have three options. The most restrictive is to use local name resolution only if the name does not exist in DNS. This option is considered the most secure because if the intranet DNS servers cannot be reached, or if there are other types of DNS errors, the intranet server names are not leaked to the subnet through local name resolution. The second and recommended option is to use local name resolution if the name doesn't exist in DNS or DNS servers are unreachable when the client computer is on a private network. The final and least restrictive option is to use local name resolution for any kind of DNS resolution error. This option is considered the least secure because the names of intranet network servers can be leaked to the local subnet through local name resolution.

To configure local name resolution for clients in Windows PowerShell, use the `Set-DAClientDNSConfiguration` cmdlet with the `-Local` parameter. The three choices available in the GUI are designated by the `FallbackSecure`, `FallbackPrivate`, or `FallbackUnsecure` values, respectively.

MORE INFO For more information about the `Set-DAClientDNSConfiguration` cmdlet, use `Get-Help` or visit <http://technet.microsoft.com/en-us/library/hh918389>.

The third page of the Infrastructure Server Setup Wizard is the DNS Suffix Search List Page, shown in Figure 6-17.

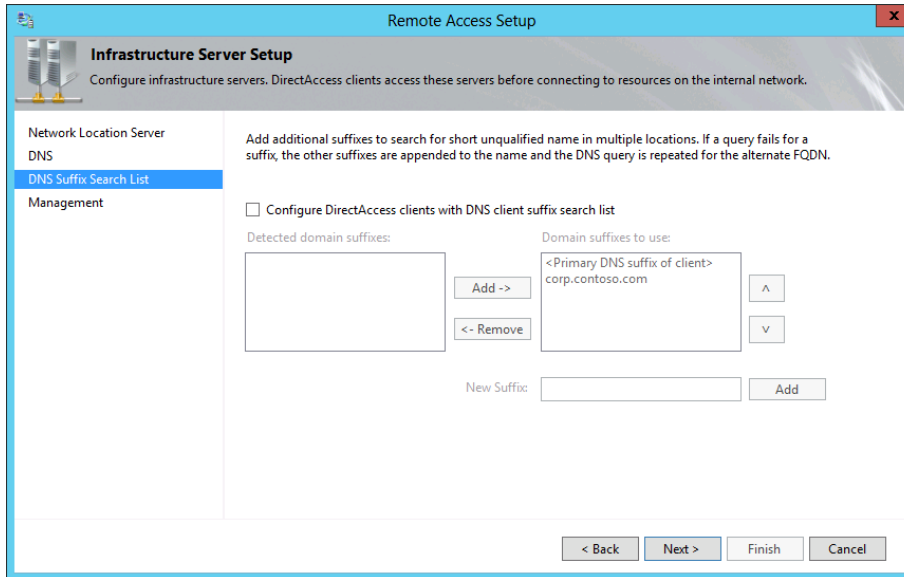


FIGURE 6-17 The DNS Suffix Search List page of the Infrastructure Server Setup Wizard

DirectAccess clients use the list you configure here to resolve single label names, such as `http://finance`. DNS cannot resolve single label names unless the DNS client first appends a suffix or if a GlobalNames zone is in use on the DNS servers. By default, the suffix is the domain to which the computer is joined.



EXAM TIP

Your DNS suffix search list should normally match the namespace rules in your NRPT. This is especially important in split-brained DNS scenarios, in which both an organization's internal private network and its publicly accessible resources use the same DNS domain name (such as `contoso.com`). To help DirectAccess clients resolve internal names correctly from the Internet, you can enter the full name of internal resources in the Name Suffix list (shown in Figure 6-16) and then specify for these resources a DNS server address corresponding to the IPv6 address of the internal DNS server. Likewise, you can enter the full name of *external* resources in the Name Suffix list and then leave the DNS server address blank. A blank entry in the DNS server address directs the client to use the DNS server currently assigned to its network connection for the suffix or FQDN specified.

The fourth and final page of the Infrastructure Server Setup Wizard is the Management page, shown in Figure 6-18.

You don't need to enter any domain controllers or Configuration Manager servers here because they are automatically detected the first time that DirectAccess is configured. Instead, use this page to configure DNS clients with the names of management servers that cannot be detected automatically, such as Windows Server Update Services (WSUS)

update servers and antivirus servers. Note that if the list of available domain controllers or Configuration Manager servers is modified after you configure DirectAccess, you can simply click Update Management Servers in the Remote Access Management Console to refresh the management server list.

There is one other point to be aware of: Management servers that initiate connections to DirectAccess clients must support IPv6 either natively or through ISATAP.

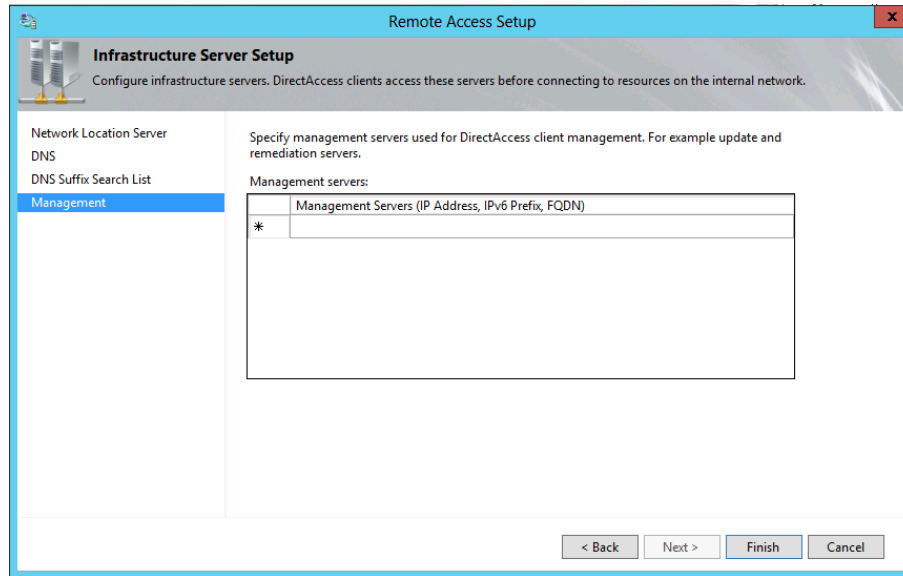


FIGURE 6-18 The Management page of the Infrastructure Server Setup Wizard

STEP 4: DIRECTACCESS APPLICATION SERVER SETUP

DirectAccess Application Server Setup is a single configuration page, as shown in Figure 6-19. You can use this page to configure encryption between the application servers you specify here and the DirectAccess server. (By default, of course, traffic is already encrypted between the DirectAccess client and server.)

To configure the list of application servers using Windows PowerShell, use the Add-DAAppServer cmdlet.

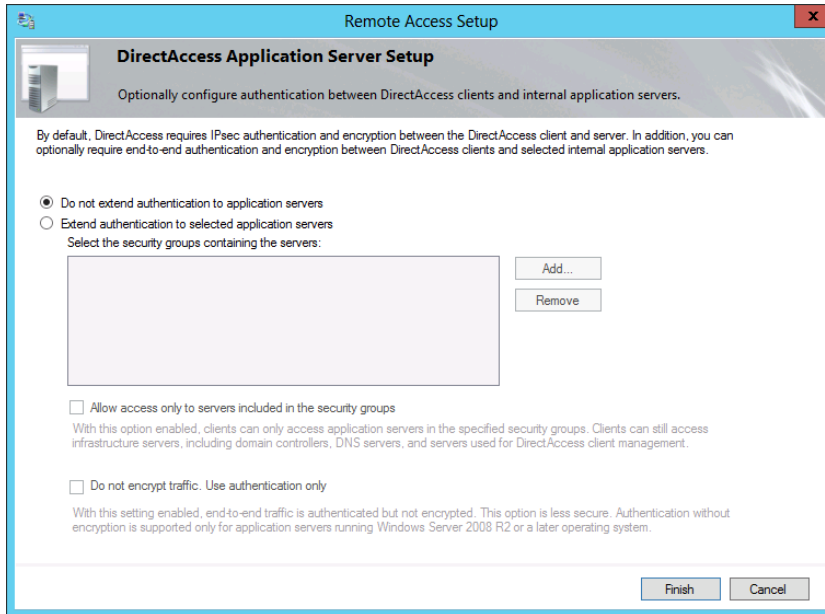


FIGURE 6-19 The DirectAccess Application Server Setup page

STEP 5: ADVANCED CONFIGURATION OPTIONS

After you complete DirectAccess Application Server Setup, the Remote Access Management Console appears, as shown in Figure 6-20. At this point, you can start new wizards to configure advanced options such as a multisite deployment or load balancing by clicking the related options in the Tasks pane.

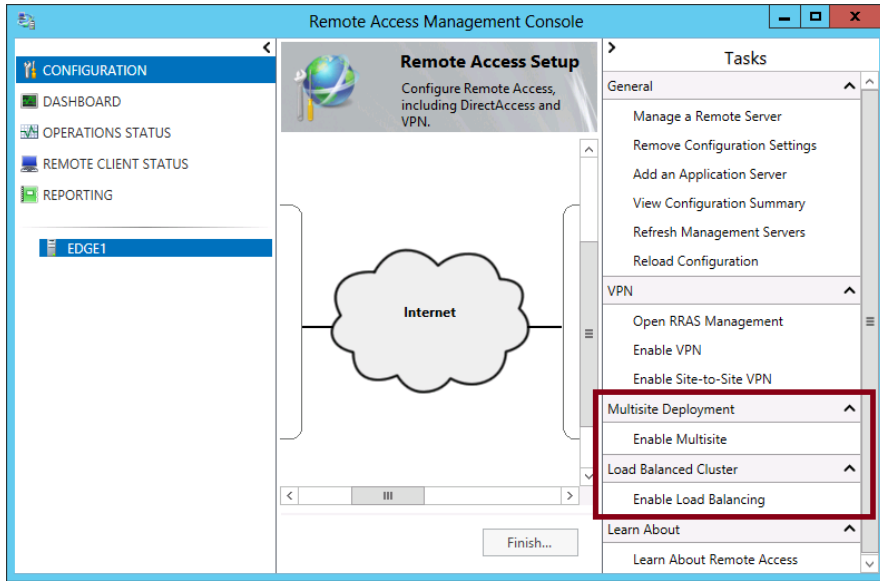


FIGURE 6-20 Configuring advanced DirectAccess options

Verifying the configuration

After you have completed your configuration, you can use the Operations Status item in the left pane of the Remote Access Management Console to verify that DirectAccess is ready to use. Remote clients can begin to connect to the network through DirectAccess after the operations status of all components is shown to be working, as shown in Figure 6-21. This process can take several minutes or more after you complete the final configuration wizard.

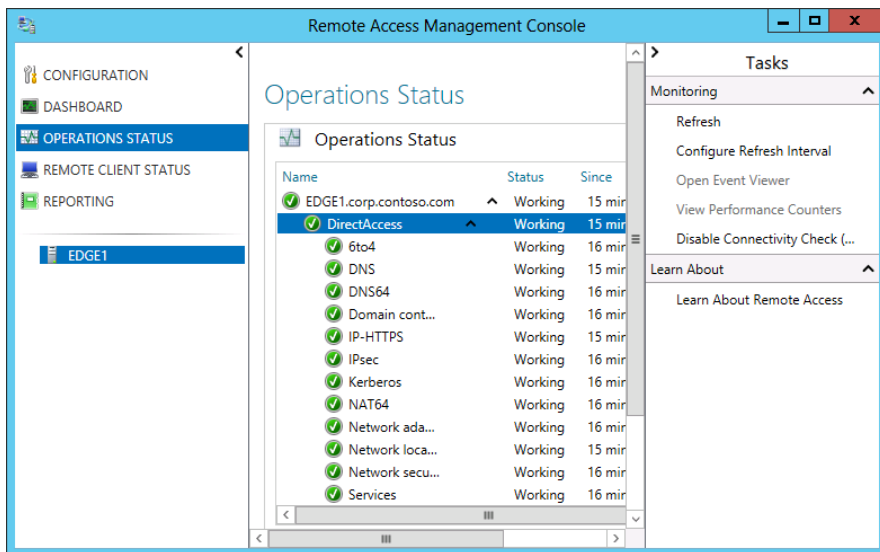
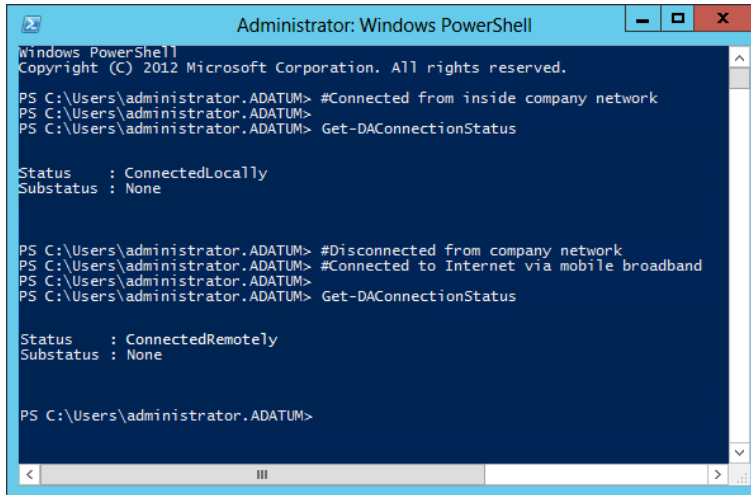


FIGURE 6-21 Using the Operations Status item

After the server components are working, you can verify DirectAccess functionality from the client end. First, you can use the `Get-DAConnectionStatus` cmdlet to determine whether DirectAccess can properly determine the location of the client. Figure 6-22 shows a Windows PowerShell console session from a portable client that is initially plugged in to a corporate network. The first time the cmdlet is run, the client is shown to be connected locally. When the laptop is disconnected and an Internet broadband connection is enabled, the cmdlet is run again. This time, the client is determined to be connected remotely, and connectivity to the intranet is established through DirectAccess. Another way to verify DirectAccess functionality on the client end is to look at the connection status in the Networks bar. Figure 6-23 shows how a DirectAccess connection appears when it is available.



```
Administrator: Windows PowerShell
Windows PowerShell
Copyright (C) 2012 Microsoft Corporation. All rights reserved.

PS C:\Users\administrator.ADATUM> #Connected from inside company network
PS C:\Users\administrator.ADATUM>
PS C:\Users\administrator.ADATUM> Get-DAConnectionStatus

Status      : ConnectedLocally
Substatus   : None

PS C:\Users\administrator.ADATUM> #Disconnected from company network
PS C:\Users\administrator.ADATUM> #Connected to Internet via mobile broadband
PS C:\Users\administrator.ADATUM>
PS C:\Users\administrator.ADATUM> Get-DAConnectionStatus

Status      : ConnectedRemotely
Substatus   : None

PS C:\Users\administrator.ADATUM>
```

FIGURE 6-22 DirectAccess automatically determines when a client is connected locally or remotely




FIGURE 6-23 A DirectAccess connection as it appears in the Networks bar in Windows 8

Notice that the icon representing a DirectAccess connection in Figure 6-23 resembles a server. Contrast this DirectAccess icon with the VPN icon shown in Figure 6-24. These icons were first introduced in Windows Server 2012 and Windows 8. You need to be able to recognize both the DirectAccess and VPN icons for the 70-417 exam.

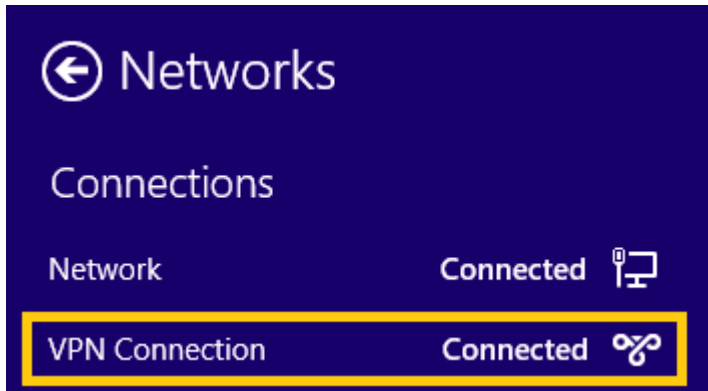


FIGURE 6-24 A VPN connection as it appears in the Networks bar in Windows 8



EXAM TIP

You should read the descriptions of the eight Group Policy settings available in the DirectAccess Client Experience Settings container within a GPO. You can find these settings in Computer Configuration\Policies\Administrative Templates\Network\DirectAccess Client Experience Settings. These settings allow you to tweak DirectAccess client behavior. For example, the Friendly Name setting allows you to specify a name other than Work-place Connection for the DirectAccess network connection. Another setting—Prefer Local Names Allowed—allows users to disconnect from NRPT rules and use normal, local DNS to resolve names and connect to local resources.

Objective summary

- DirectAccess is a bidirectional, always-on alternative to a VPN that clients can use to connect to corporate resources while they are connected to the Internet. DirectAccess first appeared as a feature in Windows Server 2008 R2 and Windows 7, but since the release of Windows Server 2012 and Windows 8, DirectAccess deployment has been greatly simplified.
- Windows Server 2012 and Windows 8 removed the requirement that DirectAccess clients authenticate themselves through computer certificates. Instead, Kerberos now is used as the default option.

- Windows Server 2012 introduced several new infrastructure and topology options for DirectAccess, including support for multiple domains, support for multiple sites, deploying the DirectAccess server behind a NAT device, and load balancing through an NLB cluster.
- In Windows Server 2012, DirectAccess and VPNs were unified in a new server role named Remote Access. To add the DirectAccess component of the Remote Access role, type the following at an elevated Windows PowerShell prompt:


```
Install-WindowsFeature DirectAccess-VPN -IncludeManagementTools
```
- You can configure DirectAccess by completing four wizards corresponding to DirectAccess clients, the DirectAccess server, infrastructure servers, and application servers. These wizards include a number of features and options that can plausibly appear in test questions on the 70-417 exam. For this reason, it is recommended that you learn about all of the options in these wizards to prepare for the exam.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the “Answers” section at the end of the chapter.

1. Which of the following is required to establish a DirectAccess connection between a Windows 8.1 client and a DirectAccess server running Windows Server 2012 R2?
 - A. A computer certificate on the client.
 - B. A user certificate on the client.
 - C. An IPv6 address on the client.
 - D. An IPv4 address on the client.

2. You are an administrator for a company with a network that includes 300 computers running Windows 8.1 and 20 servers running Windows Server 2012 R2. The network consists of a single domain named Contoso.com.

Your manager has asked you to begin testing DirectAccess with a group of 20 trial users in your organization. You deploy a single DirectAccess server on the company network edge and choose to implement computer authentication through Kerberos. You later ask the trial users to attempt to connect to the corporate network from outside the company premises. All users attempt to connect on domain-joined computers running Windows 8.1. Although most users are able to connect remotely to the corporate network, certain users working on desktop computers or virtual machines report that they cannot establish a connection. You would like to enable these users to connect to the corporate network through DirectAccess.

Which of the following Windows PowerShell commands is most likely to help you meet your goal?

- A.** Set-DAClient -OnlyRemoteComputers "Enabled"
 - B.** Set-DAClient -OnlyRemoteComputers "Disabled"
 - C.** Set-DAClient -ForceTunnel "Enabled"
 - D.** Set-DAClient -ForceTunnel "Disabled"
- 3.** You are an administrator for a company named Contoso.com with a network that includes 500 computers running Windows 8.1 and 30 servers running Windows Server 2012 R2. The network consists of a single domain named Contoso.com.

Many Contoso employees work on the road and only rarely visit the company premises. They currently connect to the company network by means of a VPN. You want to deploy DirectAccess so that you can apply software patches through System Center Configuration Manager. You don't want to enable computers to access resources on the company network through the DirectAccess connection.

Which of the following Windows PowerShell commands will help you meet your goal?

- A.** Set-DAServer -DAInstallType ManageOut
- B.** Set-DAServer -DAInstallType FullInstall
- C.** Set-DAServer -HealthCheck "Enabled"
- D.** Set-DAServer -HealthCheck "Disabled"



Thought experiment

Configuring DirectAccess at Fabrikam

You work as a network administrator for Fabrikam.com, which is based in New York and has a branch office in London. The Fabrikam.com network includes three Active Directory domains. The Fabrikam.com domain includes resources in both the New York office and the London office. The Na.fabrikam.com domain includes resources that are mostly situated in the New York office, and Eu.fabrikam.com includes resources that are mostly situated in the London office. The domain name Fabrikam.com is also used for the company's public website.

The servers on the network are running a combination of Windows Server 2008 R2 and Windows Server 2012 R2. The clients are running a combination of Windows 7 and Windows 8.1.

You are working with the rest of the IT department in planning for a DirectAccess deployment. Currently, users connect to the network remotely through a VPN. The VPN servers in both offices are running Windows Server 2008 R2.

You can find the answers to these questions in the "Answers" section.

- 1.** The New York and London offices each include two resources within the Fabrikam.com domain, resources that some remote users might need to access through a DirectAccess connection. You want to ensure that DirectAccess clients connecting to resources within the Fabrikam.com domain perform DNS lookup of these resources by contacting internal DNS servers. You also want to make sure that DirectAccess clients connect to public DNS servers when attempting to connect to public website at www.fabrikam.com. What can you do to ensure that DirectAccess clients always contact the proper DNS servers when attempting to access resources with a domain suffix of fabrikam.com?
- 2.** You want remote users to be able to automatically connect through a DirectAccess connection to the nearest entry point to the company network, whether it is in London or New York. How can you achieve this, and what requirements must first be met?
- 3.** You want to ensure that when remote users from the New York office are connected through a DirectAccess connection and enter an address such as <http://app1>, that the address for app1.na.contoso.com is first queried for in DNS, followed by app1.eu.contoso.com. How can you achieve this?
- 4.** Certain users connect to confidential resources when working remotely. For these users you want to configure two-factor authentication. However, you want to avoid the expense and administrative complexity of traditional smart cards. Which two alternative features can you consider in your environment to provide two-factor authentication?

Answers

This section contains the answers to the Objective Review and the Thought Experiment.

Objective 6.1: Review

1. Correct answer: C

- A. Incorrect:** In Windows Server 2012 R2 and Windows 8.1, Kerberos can be used in place of a computer certificate.
- B. Incorrect:** A user certificate is not required to establish a DirectAccess connection.
- C. Correct:** DirectAccess connections are based on IPv6 communication. If the DirectAccess client cannot obtain a global IPv6 address from its environment, the client must obtain one with the aid of an IPv6 transition technology.
- D. Incorrect:** IPv4 communication is not required for DirectAccess.

2. Correct answer: B

- A. Incorrect:** If only desktop computers and virtual machines are having trouble connecting through DirectAccess, this setting is most likely already enabled.
- B. Correct:** This command would disable the setting that limits DirectAccess connectivity to mobile computers only.
- C. Incorrect:** This setting would force all traffic from the client to pass through the DirectAccess connection. It would not help desktop and virtual computers establish a DirectAccess connection.
- D. Incorrect:** This setting would remove the requirement that clients force all traffic to pass through the DirectAccess connection. It would not help desktop and virtual computers establish a DirectAccess connection.

3. Correct answer: A

- A. Correct:** This command would deploy DirectAccess for remote management only.
- B. Incorrect:** This command would deploy full DirectAccess for client access and remote management.
- C. Incorrect:** This command would require NAP health checks on DirectAccess clients. It would not configure DirectAccess clients for management only.
- D. Incorrect:** This command would disable NAP health checks on DirectAccess clients. It would not configure DirectAccess clients for management only.

Thought experiment

1. You can configure the NPRT so that the four internal Fabrikam.com resources are associated with internal DNS servers.
2. You can enable a multisite deployment. You first need to make sure that the DirectAccess servers are running Windows Server 2012 or later, that the clients are running Windows 8 or later, and that your company has deployed a PKI.
3. Configure a DNS suffix search list in the Infrastructure Server Setup for the DirectAccess deployment for the Na.fabrikam.com domain.
4. You can consider virtual smart cards or OTPs.

This page intentionally left blank

Configure a network policy server infrastructure

Network Access Protection (NAP), as you know, is a Windows Server technology that enforces health requirements on client computers as they attempt to connect to a company network. These health requirements can relate to the status of software updates, of anti-virus protection, of host firewall status, or of spyware protection. NAP was first introduced in Windows Server 2008.

In a move that surprised many, Microsoft announced with the release of Windows Server 2012 R2 that NAP has been officially deprecated (set on a path to obsolescence). Some improved alternative to NAP might very well appear in a future version of Windows Server, but for now, you still have to deal with NAP on the 70-417 exam. Questions about NAP are *not* being phased out.

Although NAP doesn't include any significant new features in Windows Server 2012 or Windows Server 2012 R2, one important new feature, System Health Validator (SHV) Multi-configuration, did appear in Windows Server 2008 R2. This new feature falls within "Configure Network Access Protection," the one NAP objective listed for the 70-417 exam.

Objectives in this chapter:

- Objective 7.1: Configure Network Access Protection (NAP)

Objective 7.1: Configure Network Access Protection

NAP can be deployed in many different configurations, depending on whether it is enforced through DHCP, virtual private networks (VPNs), IPSec, Remote Desktop Services Gateway, or 802.1x. It's important to review how NAP enforcement is configured.

Most of NAP has remained the same since Windows Server 2008, but there is one new feature in NAP that falls within the Configure Network Access Protection objective: SHV Multi-configuration.

This section covers the following topics:

- How NAP works
- Configuring NAP
- SHV Multi-configuration

How NAP works

First, let's review some basic NAP concepts. When a client computer first attempts to connect to a network, its first point of contact could be a DHCP server, a VPN server, or another type of device. In a NAP infrastructure, this first point of contact is configured as a NAP enforcement point, and the NAP client is configured to report its system status (called a *statement of health* or *SoH*) to this NAP enforcement point.

The NAP enforcement point uses the RADIUS protocol to forward the SoH and connection request to a Network Policy Server (NPS). The NPS server uses connection request policies to determine whether the client connection request will be processed by NAP. If evaluated by NAP, the client request is next processed by network policies, which provide potential instructions about whether to allow the connection, block the connection, or allow restricted access only to a remediation server or set of servers. Of all the instructions defined in various network policies, only one set is applied to a connection: that of the first network policy whose conditions match the connection request.

Figure 7-1 shows an example of a simple NAP infrastructure.

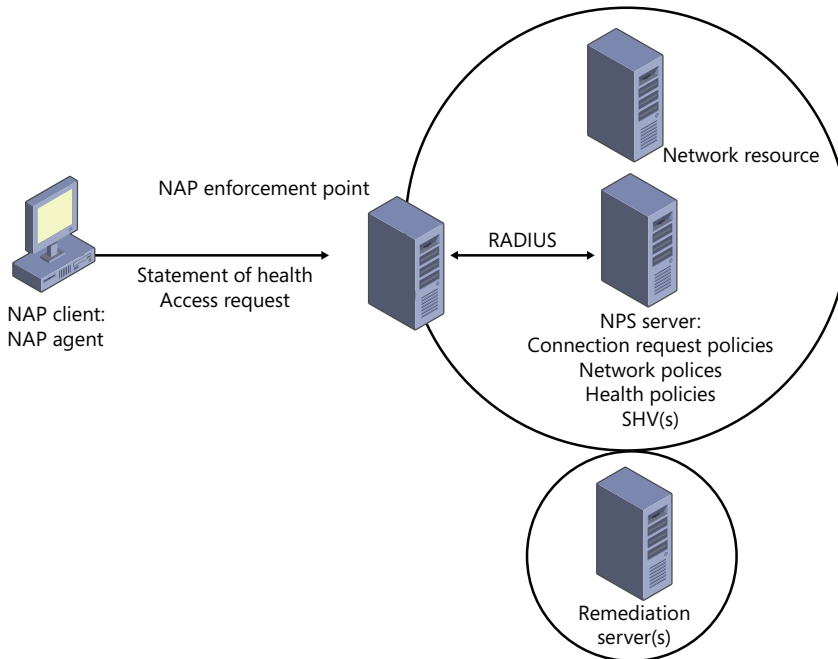


FIGURE 7-1 A NAP infrastructure



EXAM TIP

NAP enforcement types can have particular requirements that affect a NAP infrastructure. For example, both VPN and 802.1x enforcement require the NAP server to have installed a computer certificate to support server authentication through Protected Extensible Authentication Protocol (PEAP). This certificate typically is provided by a local enterprise certificate authority (CA). Another example is that IPsec enforcement requires you to install at least one Health Registration Authority (HRA) server in addition to the other NPS components. These HRA servers also require a computer certificate and comprise what is called a trusted server group for NAP IPsec clients.

Network policies usually include health policies as matching conditions. Health policies, for their part, determine whether a NAP client matches an indicated state of failing or passing a health check according to an SHV. Windows Server includes one built-in SHV, Windows Security Health Validator.

Besides the network policies that assess the health compliance of NAP clients, an additional network policy is normally also included to match clients that are not NAP-capable. These network policies include a condition named "NAP-Capable" (meaning "NAP capability status") whose value is configured as "Computer is non NAP-capable." (NAP-capable computers are ones that send an SoH.) Network policies created to match non-NAP-capable clients may be configured either to allow or block the connection request.



EXAM TIP

A network policy can include both a Health Policy condition and a NAP-Capable condition. If you configure the NAP-Capable condition with a value of "Computer is NAP-capable", then the network policy checks for health requirements only on clients that send an SoH.

The following list further describes these components involved in NAP processing:

- **Connection request policies** Rules that determine whether a connection request will be processed by network policies.
- **Network policies** Rules that compare the health of connection requests to health policy statements and accordingly allow access, block access, or allow remediated access to those requests. Network policies include conditions and condition values configured to match different types of clients. The Health Policy condition uses a health policy check to match a client. The NAP-Capable condition matches clients based on whether they have sent an SoH. The MS-Service class condition is used to match particular DHCP scopes.
- **Health policies** A statement of health compliance or noncompliance according to a particular SHV.
- **SHVs** A software component that performs a particular set of tests about the safety of a client connection.
- **Windows SHVs** The default SHV and only SHV built into Windows Server.



EXAM TIP

Make sure you understand network policies and conditions well for the 70-417 exam.

Figure 7-2 illustrates how these components could work together in a particular example of NAP processing.

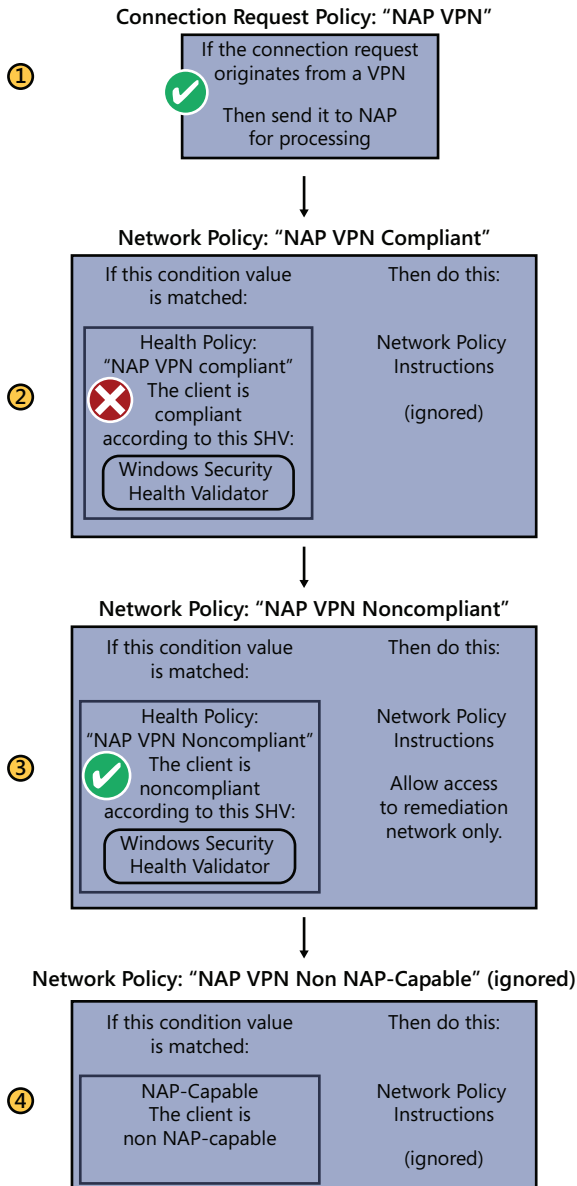


FIGURE 7-2 The first network policy that accurately describes a stated health policy condition about a NAP client provides the instructions about how to handle the NAP client request

Configuring NAP

The procedures for configuring the various NAP enforcement types all differ from each other, but they do share common steps. In general, you first configure the NAP server by using the Configure NAP Wizard. You start this wizard by clicking Configure NAP in the details pane when the NPS (Local) node is selected in the console, as shown in Figure 7-3.

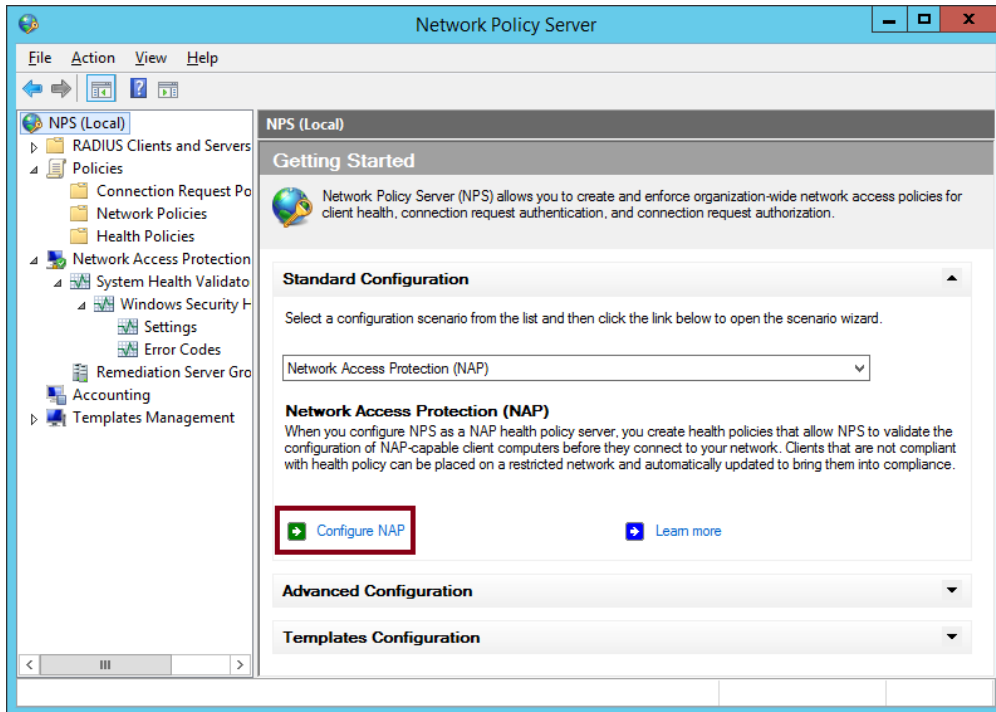


FIGURE 7-3 Configuring the NAP server

You use the Configure NAP Wizard to specify the NAP enforcement type you want to implement and to create the required connection request policies, network policies, and health policies. After running the wizard, you create security groups for NAP and configure Group Policy. You can also modify the policies created by the wizard, for example, by adding an MS-Service class condition to match the profile name you have assigned a DHCP scope on your DHCP server. This condition would accompany the Health Policy condition automatically added by the Configure NAP Wizard, as shown in Figure 7-4.



EXAM TIP

Remember that the network policy conditions shown in Figure 7-4 are used to match a state of health compliance for a particular DHCP scope. To use the MS-Service class condition shown you have to assign the scope a profile name (here, "scope1") in the scope properties on the DHCP server. The MS-Service class condition lets you apply different network policies (and therefore different levels of access protection) to different scopes.

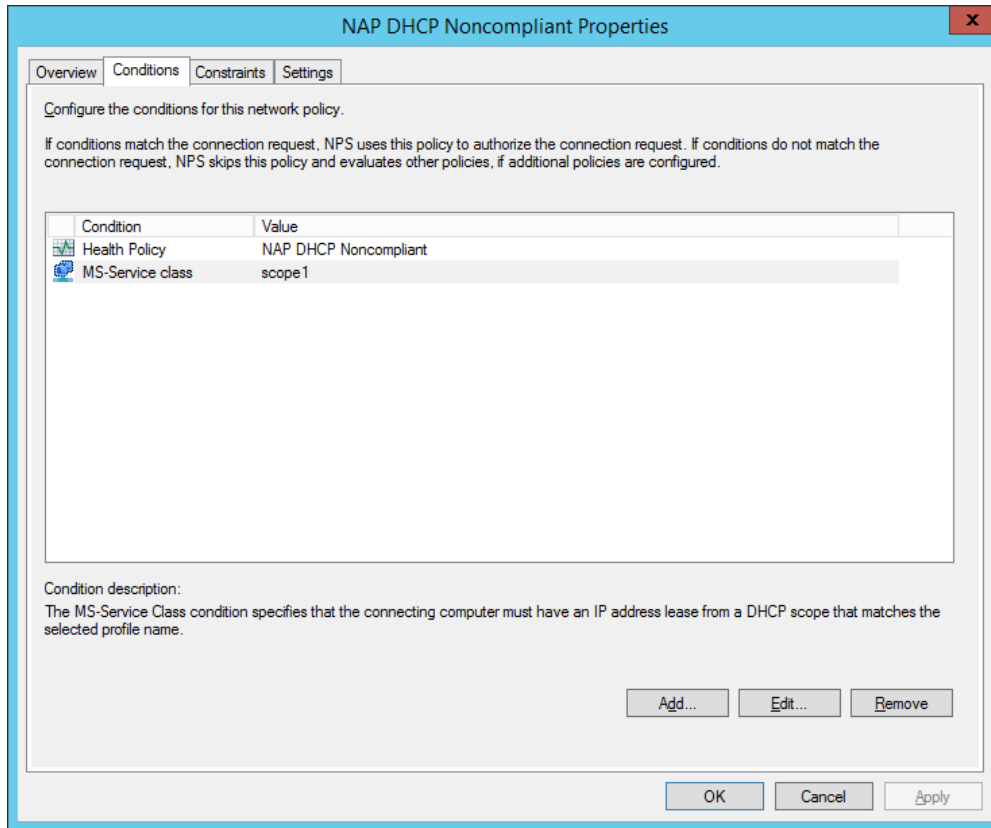


FIGURE 7-4 The matching conditions for a particular network policy named NAP DHCP Noncompliant

MORE INFO For more information about configuring the various NAP enforcement types, visit [http://technet.microsoft.com/en-us/library/dd314175\(v=ws.10\).aspx](http://technet.microsoft.com/en-us/library/dd314175(v=ws.10).aspx).



EXAM TIP

In particular, configuring NAP with IPsec enforcement requires extra steps beyond those stated above. Although the configuration steps haven't changed since Windows Server 2008, it's recommended that you review IPsec-specific NAP topics such as HRAs and HRA automatic discovery as they are described at [http://technet.microsoft.com/en-us/library/dd125312\(v=ws.10\).aspx](http://technet.microsoft.com/en-us/library/dd125312(v=ws.10).aspx).

SHV multi-configuration

Windows Server 2008 allowed you to configure just one set of health tests for each SHV. As a result, an NPS server couldn't normally adjust its health checks to suit different NAP client types.

This limitation could sometimes present a problem. In some scenarios, you might prefer to apply different health checks to different enforcement methods, computers, or users. For example, you might want to require all VPN-connected computers to have their antivirus software both enabled and up-to-date but require local DHCP-based connections to have their antivirus software only enabled. To meet such a requirement in Windows Server 2008, you normally needed to use two NPS servers.

In Windows Server 2008 R2 and later, however, you can now create multiple configurations for each SHV. After you create additional configurations beyond the default configuration, you can specify which SHV configuration you want to use for a particular health policy. Figure 7-5 shows an example of multiple configurations created for the built-in SHV, Windows Security Health Validator.

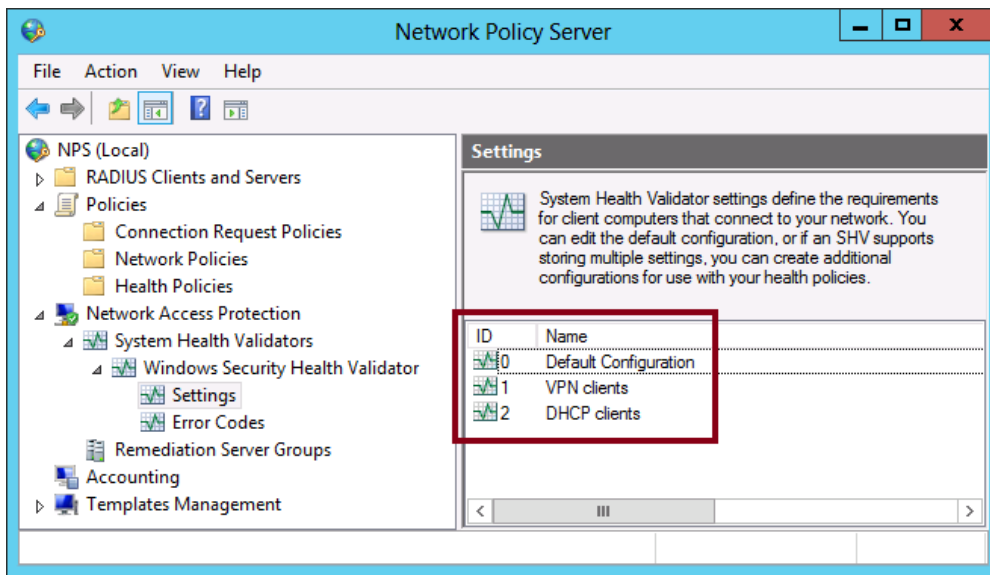


FIGURE 7-5 An SHV with three configured sets of health requirements

Default configuration

Since Windows Server 2008 R2, a Settings node now appears in the Network Policy Server console beneath the default Windows Security Health Validator (and beneath any additional SHVs you have installed that are also compatible with multiple configurations). When you select the Settings node, only the Default Configuration appears by default. This configuration can't be deleted or renamed.

Creating additional SHV configurations

To create an additional configuration for an SHV, perform the following steps. (These steps demonstrate the procedure using the built-in Windows Security Health Validator as the SHV.)

1. In the Network Policy Server console tree, navigate to Network Access Protection\System Health Validators\Windows Security Health Validator\Settings.
2. Right-click Settings and then click New, as shown in Figure 7-6.

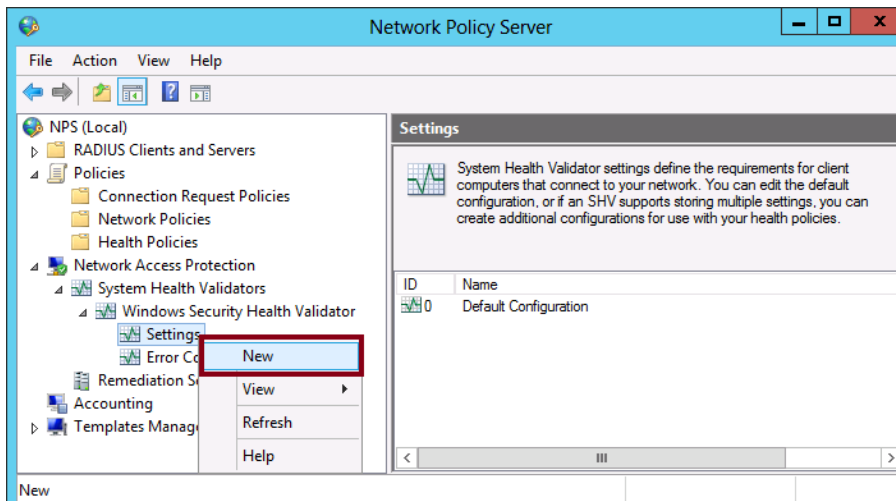


FIGURE 7-6 Creating an additional SHV configuration

3. In the Configuration Friendly Name dialog box, type a name for the new configuration and then click OK.
4. In the Windows Security Health Validator window, shown in Figure 7-7, specify the desired system health requirements for the configuration.



FIGURE 7-7 Specifying settings for a new SHV configuration

You can enable any of the following health checks:

- **A Firewall Is Enabled For All Network Connections** If this check box is selected, the client computer must have a firewall that is registered with Windows Security Center and that is enabled for all network connections.
- **An Antivirus Application Is On** If this check box is selected, the client computer must have an antivirus application installed, registered with Windows Security Center, and turned on.
- **Antivirus Is Up To Date** If this check box is selected, the client computer can also be checked to ensure that the antivirus signature file is up-to-date.
- **An Antispyware Application Is On** If this check box is selected, the client computer must have an antispyware application installed, registered with Windows Security Center, and turned on. (Not available for Windows XP)
- **Antispyware Is Up To Date** If this check box is selected, the client computer can also be checked to ensure that the antispyware signature file is up-to-date. (Not available for Windows XP.)
- **Automatic Updating Is Enabled** If this check box is selected, the client computer must be configured to check for updates from Windows Update. You can choose whether to download and install them.

- **Security Update Settings** Use this section to define health checks related to security updates. If you select the option to restrict access for clients that do not have all available security updates installed, clients will be designated as noncompliant if they do not meet this requirement according to the criteria you specify. You can specify the minimum severity level required for the updates and the minimum number of hours allowed since the client has checked for security updates. You can also choose require clients to use Windows Server Update Services (WSUS), Windows Update, or both sources.



EXAM TIP

Remember that the two anti-spyware checks are not available for Windows XP.

Assigning an SHV configuration to a health policy

To assign different health checks to different NAP client types, you can assign different SHV configurations to the health policies created for these different client types. For example, you might want to assign one SHV configuration to your VPN client health policies and another to your DHCP client health policies.

It's best to use the Configure NAP Wizard to generate your health policies automatically. The health policies created by the Configure NAP Wizard will be assigned appropriate names and be set as conditions in new, correctly configured network policies. Normally there will be two health policies for each client type, one compliant and one noncompliant. For example, if you run the Configure NAP Wizard twice and specify first VPN and then DHCP as the network connection methods, the wizard will generate the four health policies shown in Figure 7-8. For each client type, the noncompliant health policy serves as a matching condition for clients that do not pass one of the health checks.

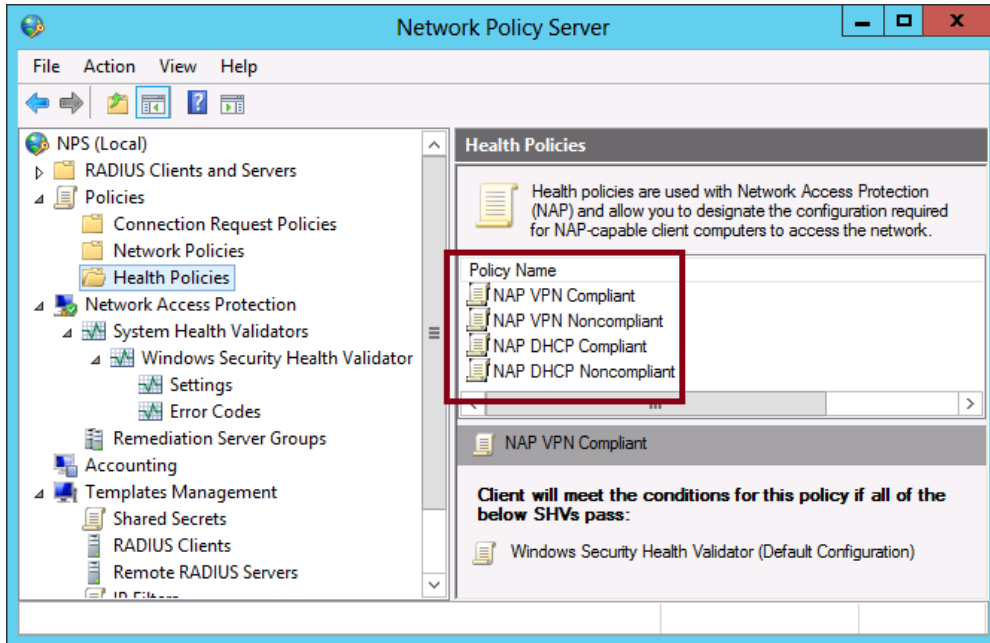


FIGURE 7-8 The Configure NAP Wizard creates a compliant and noncompliant health policy for each network connection method

If you want to assign a custom SHV configuration to a certain type of client, the only thing you have to do after running the Configure NAP Wizard is to modify the properties of the newly created health policies. You want to specify the same SHV configuration for both the compliant and noncompliant versions of the same NAP client type (for example, VPN or DHCP).

By default, when a new health policy is created, the Default Configuration of the SHV is used to define the health checks for that health policy. To assign a nondefault SHV configuration instead, perform the following steps:

1. In the Network Policy Server console, navigate to Policies\Health Policies and then double-click the name of the health policy that you want to modify.
2. On the Settings tab, in the SHVs Used In This Health Policy list, click the drop-down arrow in the Setting column for the Windows Security Health Validator SHV to see a list of available configurations. (Figure 7-9 shows an example.)
3. Select the desired configuration in the Setting drop-down list and then click OK.

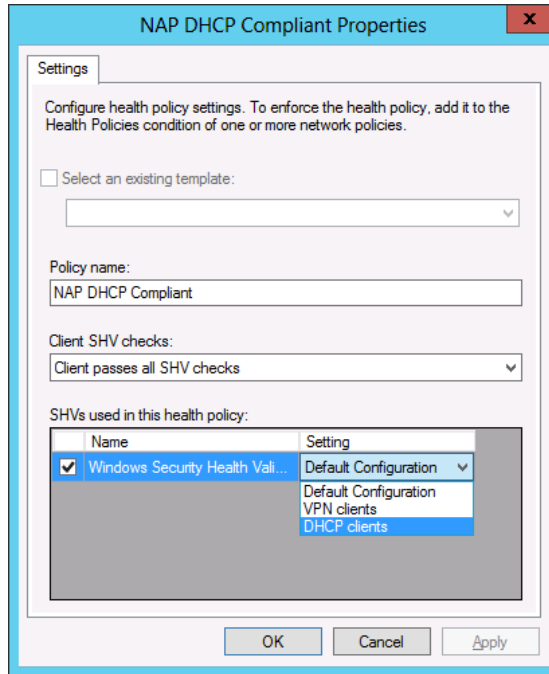


FIGURE 7-9 Assigning an SHV configuration to a health policy

Objective summary

- NAP is a technology that enforces health requirements on client computers as they attempt to connect to a network.
- The NAP feature most likely to be tested is SHV Multi-configuration. This feature first appeared in Windows Server 2008 R2. With SHV Multi-configuration, you can define different sets of health checks for a single SHV. You might use this feature to assign a higher health standard for certain types of NAP clients, such as VPN clients.
- After you create a new configuration for an SHV, you can assign that configuration to health policies. The configuration is applied to a particular NAP client type if you modify the health policies created for that client type.
- NAP has not changed much since Windows Server 2008, so you should be prepared to answer some of the same types of questions about this feature that you saw when you last earned your certification.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the “Answers” section at the end of the chapter.

1. You have deployed NAP in your network with VPN enforcement. You have deployed a single NPS on a computer running Windows Server 2012 R2 and are using Windows Security Health Validator as the only SHV.

Through your NAP policies, VPN clients that are evaluated as noncompliant are allowed access only to a set of remediation servers.

You now want to implement NAP with DHCP enforcement. However, you only want to log noncompliant DHCP clients. You don't want to block noncompliant DHCP clients from accessing any part of the network.

What should you do?

- A. Create a new configuration for the SHV for the DHCP clients.
 - B. Install an additional SHV and configure it for the DHCP clients.
 - C. Modify the default NAP DHCP connection request policy.
 - D. Modify the default NAP DHCP Noncompliant network policy.
2. You have deployed NAP in your network with VPN enforcement. You have deployed a single NPS on a computer running Windows Server 2012 R2 and are using Windows Security Health Validator as the only SHV. Your VPN clients are allowed only restricted access to the network if either security updates or virus definitions are not up-to-date. You now want to implement NAP with DHCP enforcement. However, you want to use NAP to ensure only that automatic updates and antivirus software are enabled on the DHCP client.

What should you do?

- A. Create a new configuration for the SHV for the DHCP clients.
 - B. Install an additional SHV and configure it for the DHCP clients.
 - C. Modify the default NAP DHCP connection request policy.
 - D. Modify the default NAP DHCP Noncompliant network policy.
3. You have been testing a new deployment of NAP in your network. NAP is currently configured so that VPN clients with antivirus software that is not up-to-date log their status with the NPS. These clients are currently not blocked from network access. You now want to change your NAP configuration so that the access of the same VPN clients is now restricted to a set of remediation servers on the company network.
- How can you achieve this goal?
- A. Modify the NAP VPN Compliant network policy.
 - B. Modify the NAP VPN Noncompliant network policy.
 - C. Modify the NAP VPN Compliant health policy.
 - D. Modify the NAP VPN Noncompliant health policy.



Thought experiment

Configuring NAP at Fabrikam

You work as a network administrator for Fabrikam.com. A month ago, you began testing NAP on VPN clients. You configured the Windows Security Health Validator to determine whether VPN clients had antivirus software enabled. In your current configuration, clients that are determined to be noncompliant simply report their status. They are not denied access to the network.

In the month since you implemented NAP, you have successfully remediated the client computers that have reported their antivirus application as disabled. You now are ready to move beyond the testing and want to modify your configuration to enforce a stricter NAP policy. You can find the answers to these questions in the “Answers” section.

- 1.** Currently, your NAP policies determine only whether an antivirus application is enabled on the client. You now want to add a second health check, to ensure that a firewall is enabled on the client. You also want to ensure that the client firewall is automatically enabled if it is determined to be in a disabled state. How can you achieve this?
- 2.** You want to completely block access to VPN clients that are determined to be infected. How can you achieve this without blocking other clients?
- 3.** You want to assign additional security update health checks to users who are members of the Finance group who connect through a VPN. How can you achieve this?

Answers

This section contains the answers to the Objective Review and the Thought Experiment.

Objective 7.1: Review

1. Correct answer: D

- A. Incorrect:** An SHV configuration does not affect how a client request is handled. It is used only to perform health checks on a client.
- B. Incorrect:** A new SHV would not determine whether a client connection is allowed or blocked. That behavior is determined by network policies.
- C. Incorrect:** Connection request policies do not determine how noncompliant client requests are handled. They determine whether connection requests are evaluated by NAP.
- D. Correct:** The Configure NAP Wizard creates a NAP DHCP Noncompliant policy that determines how noncompliant DHCP client requests are handled. To allow noncompliant DHCP clients to access the network and simply log the noncompliance, you need to modify the properties of this policy.

2. Correct answer: A

- A. Correct:** You can create an additional configuration for the SHV that performs only the desired checks on DHCP clients. You then need to assign this configuration to the health policies created for DHCP clients.
- B. Incorrect:** You don't need to install an additional SHV. The built-in SHV includes the health checks you need. You need only to create a second configuration of the built-in SHV.
- C. Incorrect:** A connection request policy doesn't determine which particular health checks are performed. It determines whether a connection request is evaluated by NAP.
- D. Incorrect:** A network policy doesn't allow you to specify particular health checks to be performed. It specifies a health policy that in turn specifies a SHV configuration. To modify which health checks are performed, you need to change the SHV configuration.

3. Correct answer: B

- A. Incorrect:** You don't want to change how compliant VPN clients are handled. You want to change how noncompliant VPN clients are handled.
- B. Correct:** A network policy determines how compliant or noncompliant connection requests are handled. In this case, you want to change how noncompliant VPN clients are handled. To achieve your goal, modify the NAP Enforcement setting on the Settings tab of the NAP VPN Noncompliant network policy. Change the setting from Allow Full Network Access to Allow Limited Access.
- C. Incorrect:** A health policy doesn't change how the connection requests from compliant or noncompliant clients are handled. It changes only how connection requests are evaluated.
- D. Incorrect:** A health policy doesn't change how the connection requests from compliant or noncompliant clients are handled. It changes only how connection requests are evaluated.

Thought experiment

- 1.** Modify the Windows Security Health Validator policy so that it verifies that a firewall is enabled for all network connections. Next, in the network policy that matches the VPN clients that are noncompliant, select the option on the Settings tab to enable auto-remediation of client computers.
- 2.** Create a new health policy that specifies the client SHV check as Client Reported As Infected By One Or More SHVs. Create a new network policy that specifies the new health policy as a condition, and configure the new network policy to deny access. Move the new network policy to the top of the list of network policies.
- 3.** First run the Configure NAP Wizard and specify VPN as the connection method and the Finance group as the user group to which the policy should apply. Next, create a second configuration for the Windows Security Health Validator that performs a check of security updates in the manner you wish. Finally, attach the new configuration of the Windows Security Health Validator to the new health policies just created by the Configure NAP Wizard.

Configure and manage Active Directory

This domain is inherited from Exam 70-411: Administering Windows Server 2012. Two of the four original objectives from that domain are officially listed for the 70-417 exam: Configure Domain Controllers and Maintain Active Directory. More important for our upgrade exam purposes, these objectives together include just two new features that have appeared since Windows Server 2008: domain controller cloning and the Active Directory Recycle Bin.

Objectives in this chapter:

- Objective 8.1: Configure domain controllers
- Objective 8.2: Maintain Active Directory

Objective 8.1: Configure domain controllers

Our interest in “configuring domain controllers” is quite narrow here. Although you might see a few questions on the 70-417 exam that tests only what you needed to know in this area for your Windows Server 2008 certification (how to configure read-only domain controllers, for example), the most important topic within this objective will be domain controller cloning. Domain controller cloning is a genuinely new feature in Windows Server 2012 and Windows Server 2012 R2, and you will almost certainly see questions about this topic on the 70-417 exam. The first thing you need to know? The cloning capability applies only to *virtualized* domain controllers.

This section covers the following topic:

- Configure domain controller cloning

Cloning domain controllers

In Windows Server 2012 and Windows Server 2012 R2, you can now deploy additional domain controllers in a domain by safely copying an existing virtual domain controller. In earlier versions of Windows Server, the fastest way to deploy a new domain controller, physical or virtual, was to start a fresh sysprepped image, promote the new server based on

that image to a domain controller, and then complete additional configuration requirements as necessary. Cloning is much faster.

Cloning a domain controller has certain environmental requirements. In addition, cloning a domain controller is a special, specific procedure that includes adding the source virtual machine (VM) to a built-in group account and running certain Windows PowerShell cmdlets prior to exporting and importing.

It's also worth noting that cloning seems perfect for questions on the 70-417 exam. Not only is this feature new in Windows Server 2012 and Windows Server 2012 R2, but the cloning prerequisites and the cloning procedure also are made up of details that are just specific enough for the exam. For this reason, be sure to learn these details well, such as the specific names of cmdlets and xml files.

Prerequisites for cloning

To clone a domain controller, three different servers need to be running Windows Server 2012 or Windows Server 2012 R2:

- The host server with the Hyper-V server role installed, on which the source VM is running as a guest
- A second server, physical or virtual, that is itself a domain controller in the same domain as the VM to be cloned, and that is hosting the PDC Emulator operations master role
- The source VM to be cloned that is a domain controller, and that cannot be hosting the PDC Emulator operations master role

Note also that the domain controller clone that results from the cloning procedure will be located in the same site as the source domain controller.

Add the source domain controller to the Cloneable Domain Controllers group

If your environment meets the prerequisites, you are ready to begin the cloning procedure. The first step in this procedure is to add the source VM that is a domain controller to the Cloneable Domain Controllers global security group. This built-in group account is new to Windows Server 2012 and Windows Server 2012 R2 and is found in the Users container within a domain, as shown in Figure 8-1.

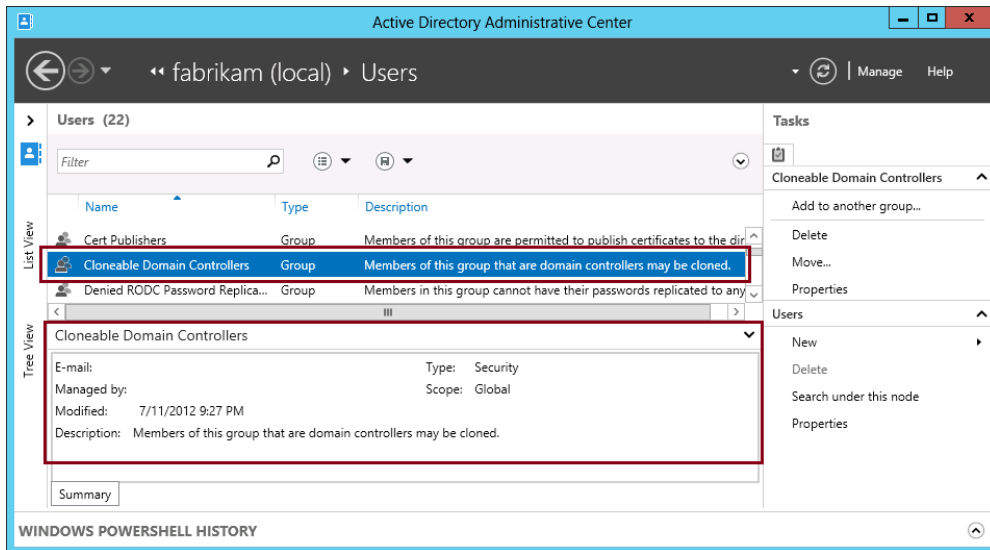


FIGURE 8-1 Windows Server 2012 and Windows Server 2012 R2 include a new Cloneable Domain Controllers global security group

POWERSHELL To add a user or computer account to a new security group, you can use the `Add-ADGroupMember` cmdlet. For more information about this cmdlet, visit <http://technet.microsoft.com/en-us/library/ee617210.aspx>.

Review applications with the `Get-ADDCCloningExcludedApplicationList` cmdlet

The next step in cloning a virtual domain controller is to run the `Get-ADDCCloningExcludedApplicationList` cmdlet on the source VM. The purpose of this cmdlet is to present a list of applications or services that are not evaluated for cloning and that are installed on the source VM. If no such applications are found, the cmdlet provides the output shown in Figure 8-2. If an unevaluated application is found, an output similar to the one in Figure 8-3 is displayed.

A screenshot of a Windows PowerShell window. The title bar reads "Windows PowerShell". The text inside the window shows the command prompt: "PS C:\Users\jcmackin.FABRIKAM> Get-ADDCCloningExcludedApplicationList". The output of the command is "No excluded applications were detected.", which is highlighted in yellow. The prompt "PS C:\Users\jcmackin.FABRIKAM> _" is visible below the output.

FIGURE 8-2 Output of Get-ADDCCloningExcludedApplicationList revealing no unsafe applications

A screenshot of a Windows PowerShell window. The title bar reads "Windows PowerShell". The text inside the window shows the command prompt: "PS C:\Users\jcmackin.FABRIKAM> Get-ADDCCloningExcludedApplicationList". The output is a table with two columns: "Name" and "Type". The first row of data is "ISO Recorder" under "Name" and "Program" under "Type". The prompt "PS C:\Users\jcmackin.FABRIKAM> _" is visible below the output.

| Name | Type |
|--------------|---------|
| ISO Recorder | Program |

FIGURE 8-3 Output of Get-ADDCCloningExcludedApplicationList revealing a potentially unsafe application

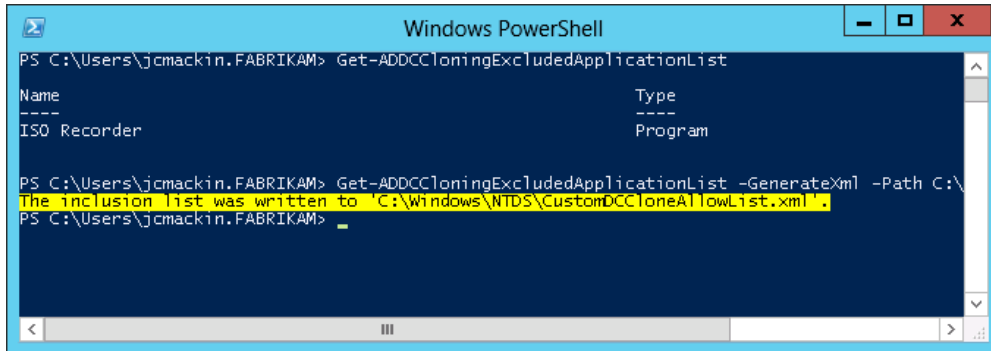
If the cmdlet returns a list of services and installed programs, review the list. Consult the software vendor of each application to determine whether it can be safely cloned. If any applications or services in the list cannot be safely cloned, you must uninstall them from the source domain controller at this point.

Next, if you determine that the services or applications returned by the Get-ADDCCloningExcludedApplicationList cmdlet are safe for cloning, you can add them to an inclusion list stored in a file named CustomDCCloneAllowList.xml. To do so, use the

-GenerateXml parameter with the same cmdlet. For example, the following command generates the excluded application list as a file named CustomDCCloneAllowList.xml at the specified folder path (C:\Windows\NTDS) and forces an overwrite if a file by that name is found to already exist at that path location:

```
Get-ADDCCloningExcludedApplicationList -GenerateXml -Path C:\Windows\NTDS -Force
```

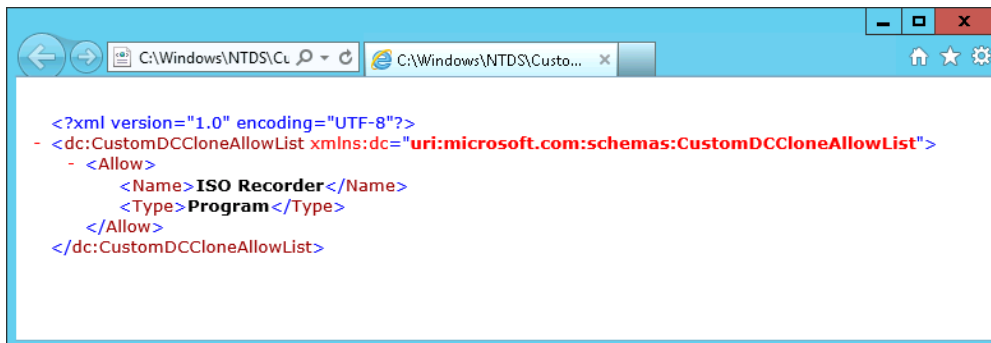
The output of this command is shown in Figure 8-4, and the contents of the CustomDCCloneAllowList.xml file are shown in Figure 8-5.



```
Windows PowerShell
PS C:\Users\jcmackin.FABRIKAM> Get-ADDCCloningExcludedApplicationList
Name                               Type
----                               -
ISO Recorder                        Program

PS C:\Users\jcmackin.FABRIKAM> Get-ADDCCloningExcludedApplicationList -GenerateXml -Path C:\Windows\NTDS -Force
The inclusion list was written to 'C:\Windows\NTDS\CustomDCCloneAllowList.xml'.
PS C:\Users\jcmackin.FABRIKAM>
```

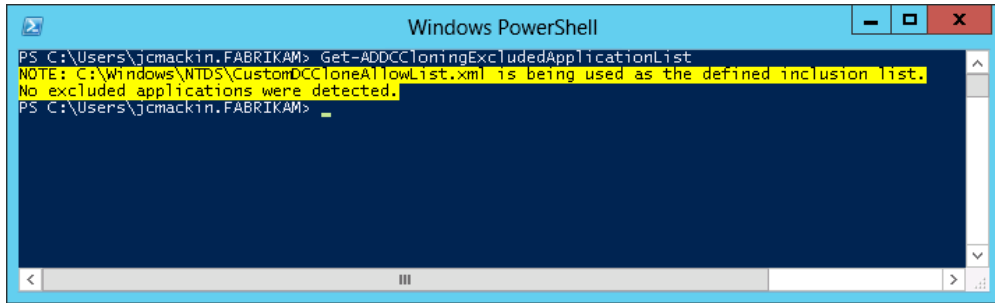
FIGURE 8-4 Adding detected applications to the inclusion list



```
<?xml version="1.0" encoding="UTF-8"?>
- <dc:CustomDCCloneAllowList xmlns:dc="uri:microsoft.com:schemas:CustomDCCloneAllowList">
  - <Allow>
    <Name>ISO Recorder</Name>
    <Type>Program</Type>
  </Allow>
</dc:CustomDCCloneAllowList>
```

FIGURE 8-5 The CustomDCCloneAllowList.xml file

After you perform this step, the Get-ADDCCloningExcludedApplicationList cmdlet will provide the output shown in Figure 8-6.



```
Windows PowerShell
PS C:\Users\jcmack\in.FABRIKAM> Get-ADDCCloningExcludedApplicationList
NOTE: C:\Windows\NTDS\CustomDCCloneAllowList.xml is being used as the defined inclusion list.
No excluded applications were detected.
PS C:\Users\jcmack\in.FABRIKAM> _
```

FIGURE 8-6 The output of `Get-ADDCCloningExcludedApplicationList` after adding a detected application to the inclusion list

Note that if any programs originally returned by the `Get-ADDCCloningExcludedApplicationList` cmdlet are not added to the inclusion list (`CustomDCCloneAllowList.xml`), the next step will fail.

MORE INFO For more information about the `Get-ADDCCloningExcludedApplicationList` cmdlet, visit <http://technet.microsoft.com/en-us/library/hh852291>.

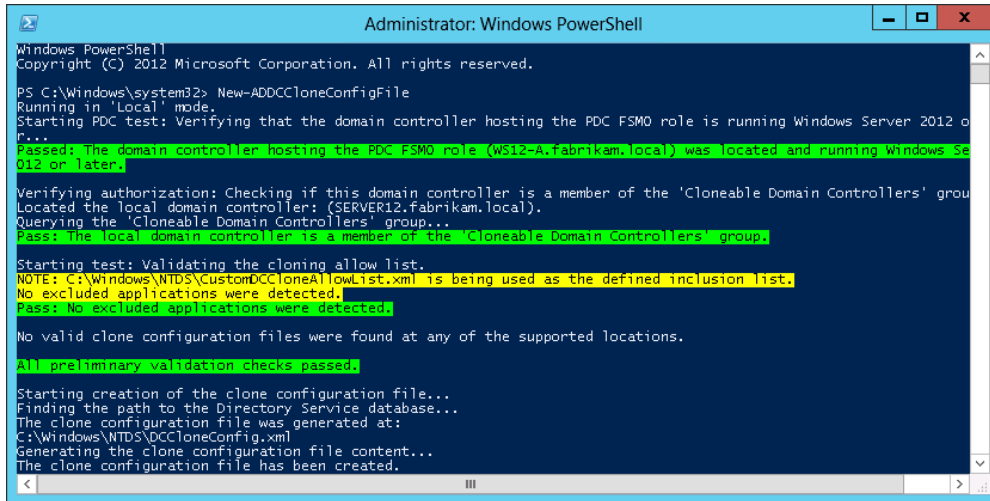
Run the `New-ADDCCloneConfigFile` cmdlet on the source VM

The `New-ADDCCloneConfigFile` cmdlet runs a list of prerequisite checks on the source VM and generates a clone configuration file, `DCCloneConfig.xml`, if all the checks succeed. The clone configuration file includes settings that you have specified for the new clone VM, such as an IP address configuration and computer name, by using parameters with the cmdlet. If the command runs successfully, the `DCCloneConfig.xml` file is saved in a location (`C:\Windows\NTDS`) that will automatically configure the clone with these settings when you later start the clone for the first time, so you don't need to look for the file or move it from its default location. (If you don't specify a name for the clone, one will be chosen automatically. If you don't specify a static IP configuration, it will be set dynamically.)

The checks succeed if the answers to the following three questions are all "yes":

- Is the PDC Emulator operations master role hosted on a domain controller running Windows Server 2012 or later?
- Is the source domain controller a member of the Cloneable Domain Controllers group?
- Are all programs and services originally listed in the output of the `Get-ADDCCloningExcludedApplicationList` cmdlet now either removed or captured in `CustomDCCloneAllowList.xml`?

A successful check of a source domain controller is shown in Figure 8-7.



```
Administrator: Windows PowerShell
Windows PowerShell
Copyright (C) 2012 Microsoft Corporation. All rights reserved.

PS C:\Windows\system32> New-ADDCCloneConfigFile
Running in 'Local' mode.
Starting PDC test: Verifying that the domain controller hosting the PDC FSMO role is running Windows Server 2012 or later...
Passed: The domain controller hosting the PDC FSMO role (WS12-A.fabrikam.local) was located and running Windows Server 2012 or later.
Verifying authorization: Checking if this domain controller is a member of the 'Cloneable Domain Controllers' group...
Located the local domain controller: (SERVER12.fabrikam.local).
Querying the 'Cloneable Domain Controllers' group...
Pass: The local domain controller is a member of the 'Cloneable Domain Controllers' group.
Starting test: Validating the cloning allow list.
NOTE: C:\Windows\NTDS\CustomDCCloneAllowList.xml is being used as the defined inclusion list.
No excluded applications were detected.
Pass: No excluded applications were detected.
No valid clone configuration files were found at any of the supported locations.
All preliminary validation checks passed.
Starting creation of the clone configuration file...
Finding the path to the Directory Service database...
The clone configuration file was generated at:
C:\Windows\NTDS\DCCloneConfig.xml
Generating the clone configuration file content...
The clone configuration file has been created.
```

FIGURE 8-7 The output of the New-ADDCCloneConfigFile cmdlet



EXAM TIP

You need to remember the function of both the CustomDCCloneAllowList.xml file and the DCCloneConfig.xml file, and that these files are saved in %Systemroot%\NTDS.

MORE INFO For more information about the New-ADDCCloneConfigFile cmdlet, visit <http://technet.microsoft.com/en-us/library/jj158947>.

Export and then import the VM of the source domain controller

To clone the VM, first shut it down. Then, you can use the Export command in Hyper-V Manager to copy the VM files to a location you choose. To export the VM using Windows PowerShell instead, use the Export-VM cmdlet as in the following example:

```
Export-VM -Name Test -Path D:\
```

At this point, *you must delete all the snapshots in the Snapshots subdirectory of the exported VM*. If desired, you can then copy the exported VM and its associated files to another computer running Windows Server 2012 or later that has the Hyper-V role installed.

Next, you can use the Import command in Hyper-V Manager to import the exported VM. Use the Copy The Virtual Machine (Create A New Unique ID) option when importing the VM, as shown in Figure 8-8. To perform this step in Windows PowerShell, use the Import-VM cmdlet as in the following example:

```
Import-VM -Path 'D:\Test2\Virtual Machines\8F148B6D-C674-413E-9FCC-4FBED185C52D.XML' - Copy -GenerateNewId
```

Finally, after importing the copy of the source VM, you can restart the source VM and then start the new clone VM.

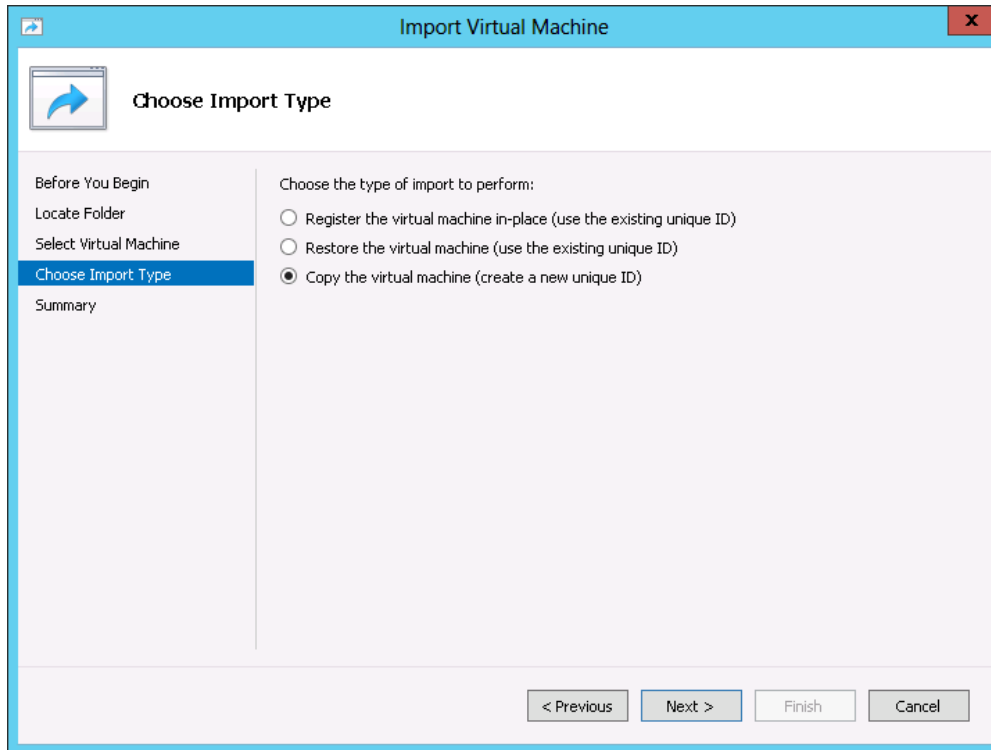


FIGURE 8-8 Creating a unique ID when importing an image allows you to use the source image again



EXAM TIP

Remember that you can use Dcpromo with the `/CreateDCAccount` option to create an RODC account in Active Directory. For more information about the options still used with Dcpromo in Windows Server 2012 and Windows Server 2012 R2, type **Dcpromo /?** at a command prompt.

Objective summary

- Windows Server 2012 and Windows Server 2012 R2 allow you to clone a virtualized domain controller for rapid deployment of a new domain controller.
- The source VM must be a member of the Cloneable Domain Controllers global security group.

- Three computers must be running Windows Server 2012 or later: the host server running Hyper-V, the guest VM that is the domain controller to be cloned, and a third domain controller that owns the PDC Emulator operations master role for the domain.
- You need to use the `Get-ADDCCloningExcludedApplicationList` cmdlet to determine whether any applications or services on the source domain controller have not yet been determined to be safe for cloning. You must either uninstall such applications or add them to the inclusion list by running the same cmdlet again with the `-GenerateXml` switch.
- Next, run the `New-ADDCCloneConfigFile` cmdlet to run prerequisite checks to determine whether the domain controller is ready to be cloned.
- When the domain controller passes the prerequisite checks, use Hyper-V Manager or the `Export-VM` and `Import-VM` cmdlets to copy the VM. Be sure to delete the snapshots of the exported VM before you import. When importing, choose the option to copy the VM and create a new unique ID.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the “Answers” section at the end of the chapter.

1. You are a network administrator for Fabrikam.com. The Fabrikam.com network includes a private cloud built on six physical servers with the Hyper-V role installed, three of which are running Windows Server 2012 R2, and three of which are running Windows Server 2008 R2. These physical servers host a total of 24 virtualized guest servers, including domain controllers, infrastructure servers, database servers, and application servers. All VMs are members of the Fabrikam.com domain.
One of the virtualized domain controllers is named DC1. DC1 is hosted on a physical server named HYPV1. You want to clone DC1 to add another domain controller to the Fabrikam.com domain.
Which of the following is *not* a requirement for cloning DC1?
 - A. DC1 must own the PDC emulator role.
 - B. HYPV1 must be running Windows Server 2012 or later.
 - C. DC1 must be running Windows Server 2012 or later.
 - D. All programs and services originally listed in the output of the `Get-ADDCCloningExcludedApplicationList` cmdlet when it is run on DC1 must be added to the Allow list in `CustomDCCloneAllowList.xml`.

2. You want to clone a domain controller named VDC1 that is running Windows Server 2012 R2 within a VM. VDC1 is hosted on a server named Host01, which is also running Windows Server 2012 R2.

VDC1 belongs to the Contoso.local domain and owns both the PDC Emulator and Infrastructure Master operations master roles for that domain. Contoso.local currently includes just one other domain controller, named VDC2. VDC2 is running Windows Server 2008 R2 and is the only global catalog server in Contoso.local.

You want to create a new domain controller named DC3 that is based on a clone of VDC1. Which steps do you need to take before you can achieve this? (Choose all that apply.)

- A. Move the Infrastructure Master operations master role to VDC2.
 - B. Make VDC1 a global catalog server.
 - C. Upgrade VDC2 to Windows Server 2012 or Windows Server 2012 R2.
 - D. Make Host01 a member of the Cloneable Domain Controllers global security group.
 - E. Make VDC1 a member of the Cloneable Domain Controllers global security group.
3. You want to clone a domain controller named DCA.fabrikam.local that is running Windows Server 2012 in a VM. DCA is hosted on a server named HV01, which is also running Windows Server 2012.

When you run the cmdlet `Get-ADDCCloningExcludedApplicationList` on DCA, the output displays the name of a single application, App1. You want to ensure that App1 is made available on all future domain controllers that result from cloning DCA. You have verified that App1 is safe for cloning.

What should you do next?

- A. Export the DCA VM.
- B. Add App1 to the `CustomDCCloneAllowList.xml` file.
- C. Run the `New-ADDCCloneConfigFile` cmdlet.
- D. Run the `New-VirtualDiskClone` cmdlet.

Objective 8.2: Maintain Active Directory

Windows Server 2008 R2 introduced Active Directory Recycle Bin, a Windows PowerShell–based feature that allowed you to restore objects deleted from the Active Directory Domain Services database. Windows Server 2012 brings this functionality of Active Directory Recycle Bin to Active Directory Administrative Center, the graphical tool for managing Active Directory Domain Services that also first appeared in Windows Server 2008 R2.

This section covers the following topics:

- Restoring deleted objects in Active directory
- Using Active Directory Recycle Bin to restore deleted objects
- Setting the deleted object lifetime in a domain

Restoring deleted objects in Active Directory

Before Windows Server 2008 R2, there were just two methods you could use to restore an object that had accidentally been deleted from Active Directory Domain Services: You could perform an authoritative restore with the Ntdsutil command-line utility, or you could use the LDP utility to perform a procedure called tombstone reanimation. Both of these methods, however, had significant drawbacks. With Ntdsutil, the drawbacks were that you first had to boot the domain controller into Directory Services Restore Mode (making the domain controller temporarily unavailable to clients on the network) and that you could only restore deleted objects that you had previously backed up. With tombstone reanimation, the drawbacks were that it was a complicated procedure and that it couldn't be relied on to restore an object's group memberships.



EXAM TIP

Remember that you use Ntdsutil to perform an authoritative restore and LDP to perform tombstone reanimation.

MORE INFO For more information about performing an authoritative restore with Ntdsutil, visit [http://technet.microsoft.com/en-us/library/cc755296\(v=WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc755296(v=WS.10).aspx). For more information about reanimating tombstoned objects, visit <http://technet.microsoft.com/en-us/magazine/2007.09.tombstones.aspx>.

Active Directory Recycle Bin

Windows Server 2008 R2 first removed these drawbacks with Active Directory Recycle Bin. With Active Directory Recycle Bin, you don't have to take the domain controller offline to restore a deleted object, and the original group memberships of the deleted objects are preserved when you restore them.

Windows Server 2008 R2 introduced Active Directory Recycle Bin in a Windows PowerShell-only mode. Windows Server 2012 and Windows Server 2012 R2 make this new feature more accessible by bringing its functionality to the graphical Active Directory Administrative Center tool. For the exam, you need to know how to enable and use Active Directory Recycle Bin in both Windows PowerShell and Active Directory Administrative Center.

ENABLING ACTIVE DIRECTORY RECYCLE BIN

For the exam and the real world, remember that the Active Directory Recycle Bin is not enabled by default. You can use Active Directory Recycle Bin to restore only those objects that have been deleted after the feature is enabled. Objects you deleted before then can be restored only through authoritative restore or tombstone reanimation.

To enable Active Directory Recycle Bin in Windows PowerShell, first make sure that all domain controllers in the domain are running Windows Server 2008 R2 or later. In addition, the functional level of your forest must be set to Windows Server 2008 R2 or higher. You can use the `Get-ADForest` cmdlet to check the functional level of your forest:

```
Get-ADForest ForestName
```

If you need to raise the functional level of the forest, you can use the `Set-ADForestMode` cmdlet with the following syntax:

```
Set-ADForestMode -Identity ForestName -ForestMode Windows2008R2Forest
```

Once your environment meets the prerequisites of Active Directory Recycle Bin, you can enable the feature by using the following Windows PowerShell command:

```
Enable-ADOptionalFeature 'Recycle Bin Feature' -scope ForestOrConfigurationSet -target DomainName -server DomainControllerName
```

To enable Active Directory Recycle Bin in the graphical user interface (GUI) in Windows Server 2012 and Windows Server 2012 R2, open Active Directory Administrative Center from the Tools menu in Server Manager. Then, in Active Directory Administrative Center, right-click the domain icon in the console tree and select Enable Recycle Bin from the shortcut menu, as shown in Figure 8-9.



EXAM TIP

Remember that Active Directory Recycle Bin requires all domain controllers to be running Windows Server 2008 R2 or later, and similarly, that the forest functional level must be set to Windows Server 2008 R2 or higher.

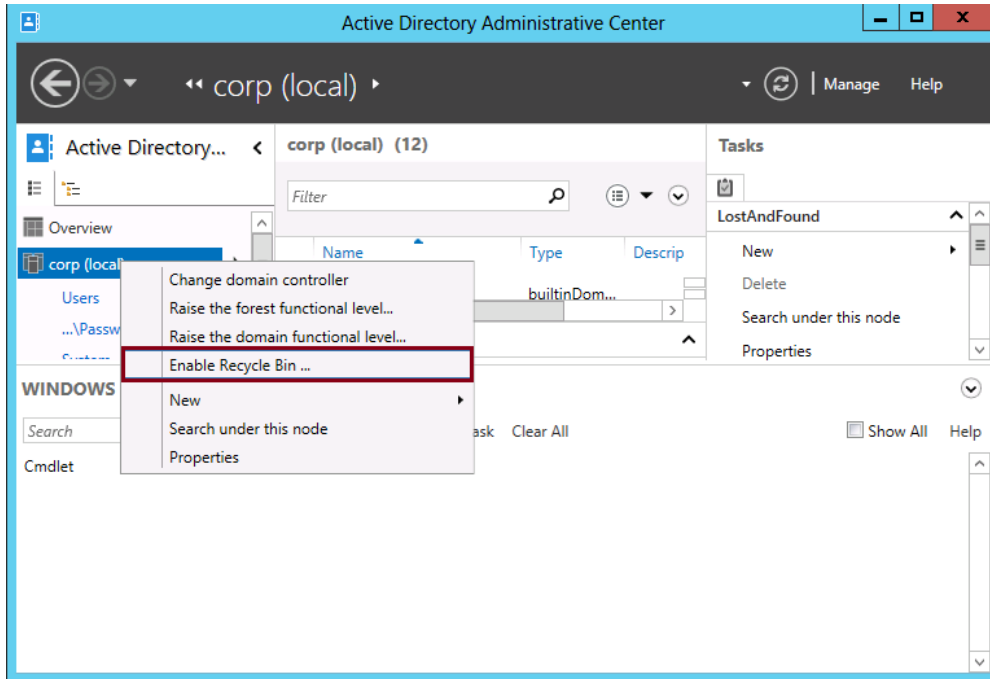


FIGURE 8-9 Enabling Active Directory Recycle Bin in Windows Server 2012 and Windows Server 2012 R2

Enabling Active Directory Recycle Bin is irreversible. In some environments, allowing administrators to see previously deleted objects might be undesirable. Consequently, you should make sure that Active Directory Recycle Bin is compatible with your organization's security policy before enabling the feature.

RESTORING DELETED OBJECTS IN ACTIVE DIRECTORY ADMINISTRATIVE CENTER

A new Deleted Objects container appears in Active Directory Administrative Center at the root of the domain container after you enable Active Directory Recycle Bin, as shown in Figure 8-10. Objects that you delete appear in this container for a period of time called the deleted object lifetime, which is 180 days by default.

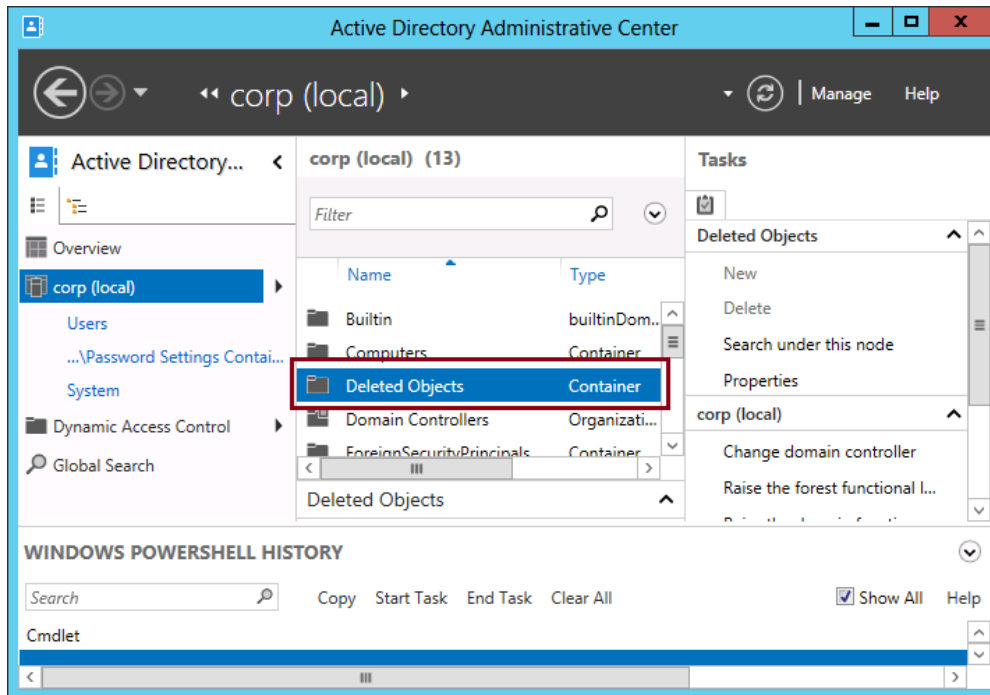


FIGURE 8-10 A Deleted Objects container appears after you enable Active Directory Recycle Bin

Restoring an object in the GUI is simple—so simple, in fact, that it might be challenging for the exam writers to devise difficult questions about restoring objects from the GUI. To restore the object to its last known parent container, just right-click the object and select Restore from the shortcut menu, as shown in Figure 8-11.



EXAM TIP

In preparing for the Maintain Active Directory objective, make sure you review the functions and features of the Ntdsutil utility. Even though this utility has not changed since Windows Server 2008, it's an essential Active Directory maintenance tool. A good way to review the capabilities of Ntdsutil is to type `ntdsutil help` at a command prompt on a domain controller. Alternatively, you can search for the utility online at the TechNet web site. Be sure to review all of the Ntdsutil contexts, such as Snapshot, which allows you to create and manage snapshots of the Active Directory database, and Files, which allows you to move the Active Directory database (Ntds.dit) to a new location.

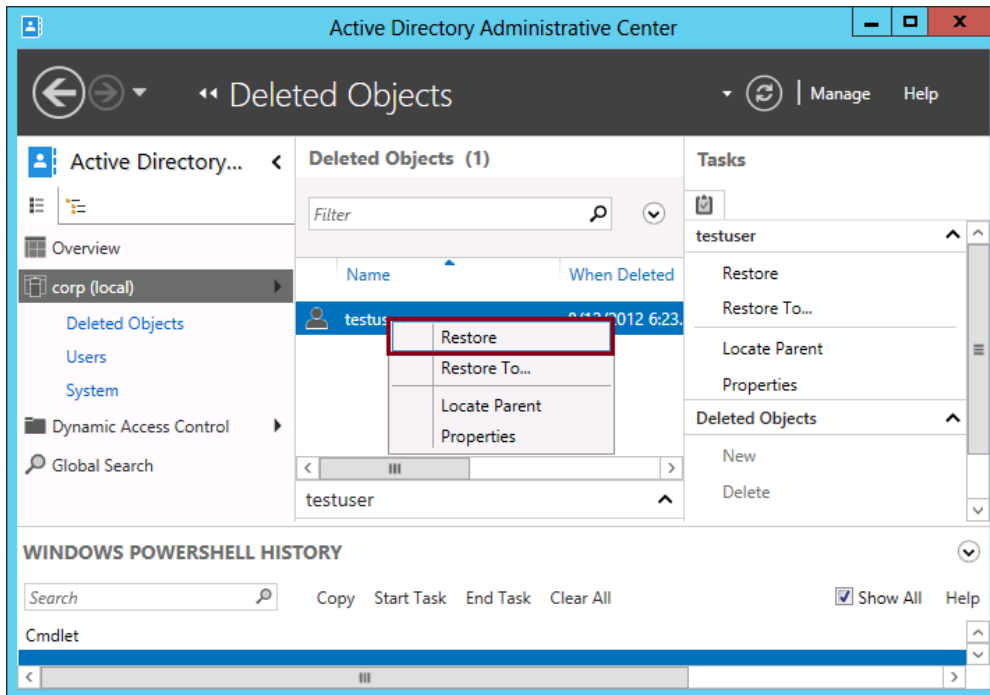


FIGURE 8-11 Restoring a deleted object in Active Directory

To restore an object to a different container, select **Restore To** and select the new container in which you want the object to appear. The **Locate Parent** option opens the former parent container in the console.

One potential complication in restoring an object might occur if you have deleted both the container and the object. In this case, you need to restore the parent before the child object, or choose to restore the object to another container.

RESTORING DELETED OBJECTS IN WINDOWS POWERSHELL

To restore a deleted object in Windows PowerShell, first use the `Get-ADObject` cmdlet with the `-Filter` and `-IncludeDeletedObjects` parameters, and then pipe the result to the `Restore-ADObject` cmdlet. For example, to restore a deleted user with the display name "Mary," type the following command at an elevated Windows PowerShell prompt:

```
Get-ADObject -Filter {DisplayName -eq "Mary"} -IncludeDeletedObjects | Restore-ADObject
```

Here's another example: To restore a user whose canonical name (CN) is like "Jorge," type the following:

```
Get-ADObject -Filter {CN -like "Jorge"} -IncludeDeletedObjects | Restore-ADObject
```

In the real world, it doesn't make much sense to restore a deleted object by using Windows PowerShell if you don't have to. However, don't be surprised if the Windows PowerShell method of Active Directory Recycle Bin still appears on the exam.



EXAM TIP

Don't forget about the Dsomain tool, which was introduced in Windows Server 2008. Dsomain allows you to mount a shadow copy-enabled backup or Ntdsutil snapshot of the Active Directory database. After using Dsomain to mount an Active Directory backup or Ntdsutil snapshot, you can use Ldp.exe to view the historical contents of the database, including object attributes. When you use Dsomain, be sure to mount the backup or snapshot in a port other than port 389, which is the reserved LDAP port used by the currently running instance of the Active Directory database. For more information about Dsomain, search for "Active Directory Domain Services Database Mounting Tool (Snapshot Viewer or Snapshot Browser) Step-by-Step Guide" or visit [http://technet.microsoft.com/en-us/library/cc753609\(v=ws.10\).aspx](http://technet.microsoft.com/en-us/library/cc753609(v=ws.10).aspx).

MORE INFO For more information about how to use Get-ADObject, visit <http://technet.microsoft.com/en-us/library/ee617198.aspx>.

DELETED OBJECT LIFETIME

By default, you have only 180 days to restore an object after it is deleted. This period is known as the deleted object lifetime and is governed by the msDS-DeletedObjectLifetime attribute assigned to the domain. To change the value of this attribute, use the Set-ADObject cmdlet in the following manner:

```
Set-ADObject -Identity "CN=Directory Service,CN=Windows NT,CN=Services,CN=Configuration,DC=<mydomain>,DC=<com>" -Partition "CN=Configuration,DC=<mydomain>,DC=<com>" -Replace:@{ "msDS-DeletedObjectLifetime" = <value> }
```



EXAM TIP

The Configure and Manage Active Directory domain is taken from the 70-411 exam. In its original form, the domain includes two additional objectives: the Configure Service Authentication objective and the Configure Account Policies objective. Thus, do not be surprised if you see a question related to one of these objectives on the 70-417 exam. You should review older topics that fall within the scope of these objectives, such as Managed Service Accounts and Password Settings Objects, especially if you haven't dealt with these features since you earned your Windows Server 2008 certification.

Replace `DC=<mydomain>,DC=<com>` with the appropriate forest root domain name of your Active Directory environment and replace `<value>` with the new value of the deleted object lifetime.

For example, to set the deleted object lifetime to 365 days, run the following command:

```
Set-ADObject -Identity "CN=Directory Service,CN=Windows NT,CN=Services,CN=Configuration,DC=contoso,DC=com" -Partition "CN=Configuration,DC=contoso,DC=com" -Replace:@{ "msDS-DeletedObjectLifetime" = 365 }
```



EXAM TIP

Remember that domains with a functional level of Windows Server 2008 or higher perform replication of the SYSVOL folder by using the Distributed File System Replication (DFSR) engine. Lower functional levels use the File Replication Service (FRS) instead.

Objective summary

- Windows Server 2008 R2 introduced Active Directory Recycle Bin, a Windows PowerShell–based tool that allowed you to restore a deleted object such as a user without taking the domain controller offline and without risk of losing the object’s group memberships.
- Windows Server 2012 and Windows Server 2012 R2 bring the functionality of Active Directory Recycle Bin to the GUI, in the Active Directory Administrative Center.
- Active Directory Recycle Bin must first be enabled. Once enabled, it can’t be disabled.
- Once Active Directory Recycle Bin is enabled, a Deleted Object container appears at the root of the domain in Active Directory Administrative Center. You can easily find and restore objects from this location.
- Deleted objects have a default lifetime of 180 days. This period is a configurable attribute of a domain, called the deleted object lifetime. To change the deleted object lifetime, use the Set-ADObject cmdlet.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the “Answers” section at the end of the chapter.

1. You are a network administrator for Contoso.com. You have learned that a user account was accidentally deleted from the Contoso.com Active Directory domain. The domain controllers in your network are all running Windows Server 2012 R2. Active Directory Recycle Bin is not yet enabled.
You want to restore the deleted user account without taking any domain controller offline. What should you do?
 - A. Perform an authoritative restore of the deleted user account with the Ntdsutil utility.
 - B. Reanimate the tombstone of the deleted object.
 - C. Enable Active Directory Recycle Bin, and use Active Directory Administrative Center to restore the object.
 - D. Enable Active Directory Recycle Bin, and use Windows PowerShell to restore the object.

2. You are a network administrator for Contoso.com. You have learned that a user account for a user named Dan Park was accidentally deleted from the Contoso.com Active Directory domain. The domain controllers in your network are all running Windows Server 2012 R2. Active Directory Recycle Bin was enabled at the time the object was deleted.

You want to restore the deleted user account without taking any domain controller offline. What should you do?

- A. Restore the object from the Deleted Objects container in Active Directory Administrative Center.
- B. Perform an authoritative restore using the Ntdsutil utility.
- C. Reanimate the tombstone of the deleted object.
- D. Run the following command:

```
Get-ADObject -Filter {displayName -eq "Dan Park"} | Restore-ADObject
```

3. You are a network administrator for Adatum.com. You have learned that all 10 user accounts in the Finance department were accidentally deleted from the Adatum.com Active Directory domain. The domain controllers in your network are all running Windows Server 2012 R2. Active Directory Recycle Bin was enabled at the time the user accounts were deleted.

You attempt to restore the deleted user accounts in Active Directory Administrative Center, but you receive errors indicating that object's parent is deleted.

You want to restore the deleted user accounts. What should you do?

- A. Use the Set-ADObject cmdlet to extend the deleted object lifetime in the domain.
- B. Re-create the parent organizational unit of the user accounts and then restore the user accounts.
- C. Restore the parent organizational unit of the user accounts and then restore the user accounts.
- D. Restart a domain controller in Directory Services Restore Mode and perform an authoritative restore of the deleted user accounts.



Thought experiment

Configuring and managing Active Directory at Proseware

You are a network administrator for Proseware.com. The Proseware.com network consists of a single Active Directory domain, including two domain controllers. One domain controller named DC1 is running Windows Server 2012 R2, and the other named DC2 is running Windows Server 2008 R2. Both domain controllers are running in GUI mode in VMs, on host servers running Windows Server 2012 R2.

One of your goals as a network administrator is to help improve both the resiliency and scalability of the network. Proseware.com has been experiencing rapid growth, and you want the network to handle increasingly heavy workloads in the coming months and years. You also want to be able to improve fault tolerance so that configuration errors such as accidental deletions can be quickly reversed.

With this information in mind, answer the following questions. You can find the answers to these questions in the “Answers” section.

1. Is DC1 eligible for cloning? Why or why not?
2. You install a third domain controller named DC3. DC3 is running a fresh installation of Windows Server 2012, with no additional software or applications. Which step(s) must you take before DC3 can be cloned?
3. You export the DC3 VM. Which step(s) must you take before importing the exported VM?
4. How would you find out if Active Directory Recycle Bin is enabled on the Proseware.com domain?
5. If Active Directory Recycle Bin is not enabled, how would you determine whether this feature can be enabled in the domain?

Answers

This section contains the answers to the Objective Reviews and the Thought Experiment.

Objective 8.1: Review

1. Correct answer: A

- A. Correct:** DC1 must not own the PDC Emulator operations master role.
- B. Incorrect:** The host server must be running Windows Server 2012 or Windows Server 2012 R2.
- C. Incorrect:** The source virtualized domain controller must be running Windows Server 2012 or later.
- D. Incorrect:** Cloning cannot occur unless every program listed in the output of the `Get-ADDCCloningExcludedApplicationList` cmdlet also appears on the inclusion list, `CustomDCCloneAllowList.xml`.

2. Correct answers: C, E

- A. Incorrect:** The cloned domain controller can be the Infrastructure Master. It cannot be the PDC Emulator. Also, the Infrastructure Master role should not be placed on a global catalog server.
- B. Incorrect:** It is not necessary to make VDC1 a global catalog server to clone it.
- C. Correct:** The PDC Emulator in the domain needs to be running Windows Server 2012 or Windows Server 2012 R2.
- D. Incorrect:** The host computer doesn't need to be in any global security group.
- E. Correct:** The domain controller you want to clone must be a member of the Cloneable Domain Controllers global security group.

3. Correct answer: B

- A. Incorrect:** You shouldn't export the VM until you add App1 to the Allow list and run the `New-ADDCCloneConfigFile` cmdlet.
- B. Correct:** The next step is to add App1 to the Allow list, called `CustomDCCloneAllowList.xml`, by running the `Get-ADDCCloningExcludedApplicationList` with the `-GenerateXml` switch.
- C. Incorrect:** You should run the `New-ADDCCloneConfigFile` cmdlet only after you add App1 to the inclusion list. This cmdlet runs a list of prerequisite checks on the source VM.
- D. Incorrect:** You don't use this cmdlet to clone a domain controller. You would use the cmdlet to clone a virtual disk.

Objective 8.2: Review

1. Correct answer: B

- A. Incorrect:** The solution requires that no domain controller should be taken offline. To perform an authoritative restore with the Ntdsutil utility, you first need to take a domain controller offline by booting it in Directory Services Restore Mode.
- B. Correct:** You cannot restore the object by using Active Directory Recycle Bin because this feature was not enabled when the object was deleted. To restore the object without taking a domain controller offline, you will have to reanimate the tombstone of the object.
- C. Incorrect:** You cannot restore the object by using Active Directory Recycle Bin because this feature was not enabled when the object was deleted.
- D. Incorrect:** You cannot restore the object by using Active Directory Recycle Bin because this feature was not enabled when the object was deleted.

2. Correct answer: A

- A. Correct:** Restoring the deleted object in the GUI is by far the simplest option. You can restore the object in Active Directory Administrative Center because Active Directory Recycle Bin was enabled when the object was deleted and at least one of your domain controllers is running Windows Server 2012 or Windows Server 2012 R2.
- B. Incorrect:** You should not perform an authoritative restore because this procedure requires you to take a domain controller offline.
- C. Incorrect:** Although you can perform this procedure, it is unnecessarily complicated. In addition, through this procedure the object you restore might be stripped of its group memberships, so it is not the best option.
- D. Incorrect:** This command will not work without the `-IncludeDeletedObjects` switch.

3. Correct answer: C

- A. Incorrect:** The errors indicate that the objects' parent is deleted. Extending the deleted object lifetime will have no effect on the state of the parent container.
- B. Incorrect:** You need to restore the original parent container, not re-create one with the same name.
- C. Correct:** If the parent container is deleted, you are able to restore it from the Deleted Objects container. After it is restored, you are able to restore its child objects.
- D. Incorrect:** This step would not help. Restoring the deleted objects in Directory Services Restore Mode would not affect the underlying problem that the parent container is missing.

Thought experiment

1. No, it is not. Although DC1 and the host computer both meet the requirements of running Windows Server 2012 or later, the network environment does not meet the requirement of a second domain controller running Windows Server 2012 or later with the PDC Emulator role.
2. You need to ensure that DC1 holds the PDC Emulator role and transfer the role to that computer if necessary. In addition, you need to add DC3 to the Cloneable Domain Controllers global security group.
3. You need to delete the snapshots in the Snapshots subdirectory.
4. You can look in Active Directory Administrative Center and determine whether a Deleted Objects container exists at the root of the domain.
5. Use the Get-Forest cmdlet to determine the functional level of the forest. The forest needs to be running at the Windows Server 2008 R2 functional level or higher.

Configure and manage Group Policy

Windows Server 2012 and Windows Server 2012 R2 include a varied assortment of enhancements to Group Policy, but only the narrow topic of configuring Group Policy processing has been singled out for the 70-417 exam. As it turns out, the most important new Group Policy feature is the only one to fall within this “processing” objective: remote Group Policy updating. This chapter introduces you to this useful new functionality introduced in Windows Server 2012.

In addition, Windows Server 2008 R2 introduced a new GroupPolicy module for Windows PowerShell. The module includes 26 cmdlets, and some of them are likely to appear on the exam. The chapter introduces these cmdlets as a reference for your exam preparation.

Finally, Windows 8.1 introduces a new feature you need to know called Group Policy Caching. Even though Group Policy Caching relates specifically to Windows 8.1, you still might see an exam question on this topic.

Objectives in this chapter:

- Objective 9.1: Configure Group Policy processing

Objective 9.1: Configure Group Policy processing

Remote Group Policy updating is the most important new Group Policy feature you need to learn about for the 70-417 exam. It appears in two guises: GUI and Windows PowerShell. The GUI feature might be a bit too straightforward by itself to serve as the basis for exam questions, so be sure to learn the extra, “complicating” details explained next, such as how the feature actually works, and which service and open ports are required for the feature to function. In Windows PowerShell, the feature hits a “sweet spot” of exam-level difficulty, so it’s more likely you’ll see exam questions based on this version of the feature. As a result, make sure you can understand the syntax and the various options that can accompany the `Invoke-GPUdate` cmdlet.

You don’t need to understand the syntax of other Windows PowerShell cmdlets in the Group Policy module. You just need to be able to recognize their purpose.

This section covers the following topics:

- Remote Group Policy update
- Group Policy cmdlets in Windows PowerShell
- Group Policy caching

Remote Group Policy update

Windows Server 2012 and Windows Server 2012 R2 introduce a handy feature that is sure to please network administrators: the ability to perform a Group Policy update on many remote computers at once. You can accomplish this task by using either the Group Policy Management Console or Windows PowerShell. Before now, of course, you had to use the `GPUpdate` command on a local computer to refresh policy for just that computer and the locally logged-on user. If you wanted to update many computers at once, you had to use a script or a third-party tool.

Updating Group Policy in an organizational unit with Group Policy Management Console

To remotely refresh Group Policy in the Group Policy Management Console, simply right-click an organizational unit (OU) container in the console tree and select Group Policy Update from the shortcut menu, as shown in Figure 9-1. This step schedules `GPUpdate.exe` to be run within 10 minutes on all clients running Windows Vista or later and on all servers running Windows Server 2008 or later in that OU.

Note the limitations: You can force a Group Policy refresh on all computers *within a single OU and all subcontainers only*. You cannot single out specific computers or update Group Policy on computers that are not located in an OU. (This restriction applies only to the Group Policy Management Console, not to Windows PowerShell.) Also, you cannot use this feature to update computers running operating systems earlier than Windows Vista and Windows Server 2008, whether through Group Policy Management or through Windows PowerShell.

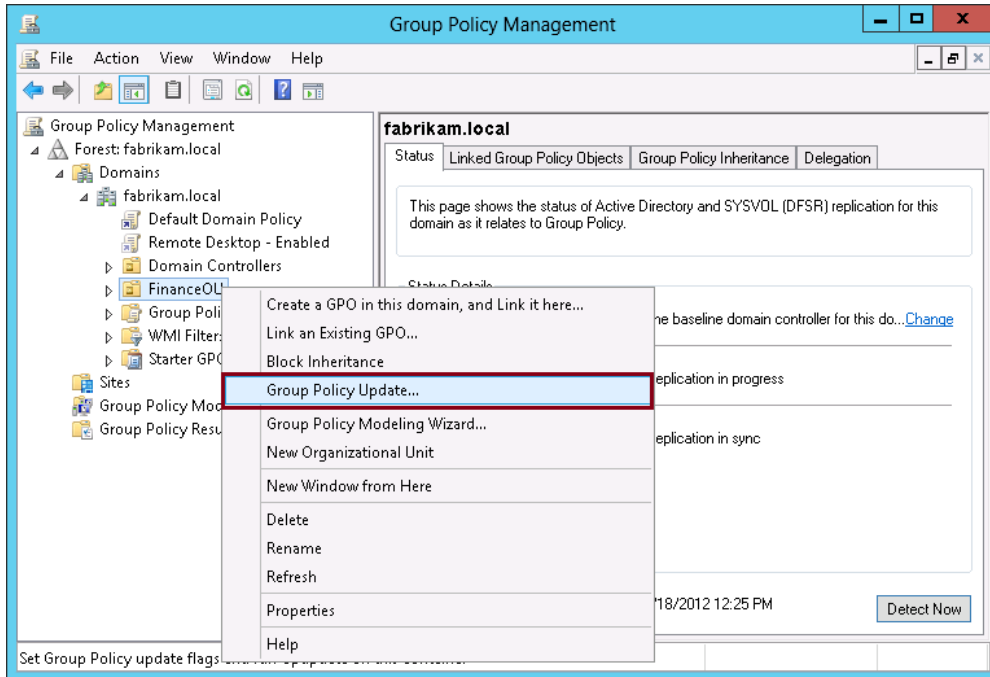


FIGURE 9-1 Updating Group Policy on all computers in an OU

After you select the Group Policy Update option, a message box appears indicating the number of computers that will be affected and asking you to confirm the update, as shown in Figure 9-2.

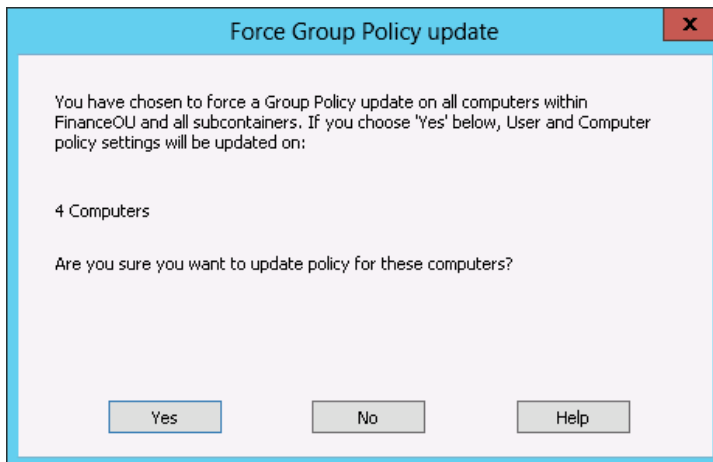


FIGURE 9-2 A remote update in Group Policy Management forces the update for all computers in an OU

When you give your consent, a window (shown in Figure 9-3) appears, indicating the success or failure of the *scheduling* of the update. The update itself is not immediate. As shown in Figure 9-3, the message indicates that a Group Policy update will be forced on all computers in the OU and all subcontainers within 10 minutes. This slight delay is a good thing when there are many computers in the OU: The computers will not all update at the same time and strain the resources of domain controllers.

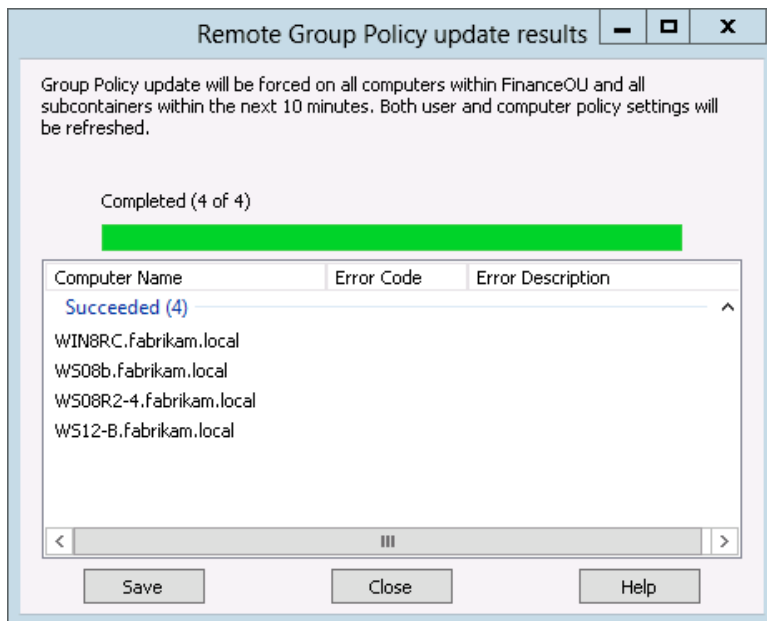


FIGURE 9-3 Updating Group Policy on all computers in an OU

Updating Group Policy with Invoke-GPUdate

You can update Group Policy on computers in a much more flexible way if you use the `Invoke-GPUdate` cmdlet in Windows PowerShell.

Used by itself without any parameters, the cmdlet is similar to `GPUdate.exe`; it updates Group Policy on the local computer only. The difference with `GPUdate.exe` is that, as with the Group Policy Management Console, the task is not performed immediately but is scheduled to be completed within 10 minutes by default.

Used with the `-Computer` parameter however, the `Invoke-GPUdate` cmdlet lets you update a remote computer, as in the following example:

```
Invoke-GPUdate -Computer WS12-B
```

Other options you can use with Invoke-GPUdate include -Force and -RandomDelayInMinutes. The -Force parameter resembles the /Force option with GPUdate.exe: It reapplies all policy settings regardless of whether they have changed. The -RandomDelayInMinutes parameter allows you to specify a random interval in minutes up to the number of minutes specified, before the Group Policy update will be run. The purpose of this option is typically to reduce the network load on domain controllers when many remote computers are updated with a scripted command, but it can also be used with a single computer to reduce or remove the default delay of 10 minutes. A value of 0 will, in fact, cause the Group Policy refresh to run immediately. The following example therefore causes all Group Policy settings to be updated immediately on a computer named WS12-B:

```
Invoke-GPUdate -Computer WS12-B -RandomDelayInMinutes 0 -Force
```

As mentioned, you can also leverage Windows PowerShell to execute the Invoke-GPUdate cmdlet on more than one computer. You can begin with the Get-ADComputer cmdlet to retrieve any group of computer objects and then pipeline the results to a "ForEach" construction that includes Invoke-GPUdate.

For example, the following command displays all of the computers in the container named Computers, in the Fabrikam.local domain:

```
Get-ADComputer -Filter * -Searchbase "CN=Computers,DC=Fabrikam,DC=local"
```

If you pipe the results of this command to a ForEach statement, you can execute the Invoke-GPUdate cmdlet on each computer returned by the command. The net result of the following command, for example, is to schedule GPUdate.exe to run on every computer in the Computers container within 10 minutes.

```
Get-ADComputer -Filter * -Searchbase "CN=Computers,DC=Fabrikam,DC=local" | ForEach {Invoke-GPUdate -Computer $_.name}
```

You don't need to target computers in any specific container or OU. The following example attempts to schedule GPUdate.exe to run on every computer in the domain within 10 minutes:

```
Get-ADComputer -Filter * | ForEach {Invoke-GPUdate -Computer $_.name}
```

This next example schedules GPUdate.exe to run immediately on every computer in the domain with a description that includes the term "finance".

```
Get-ADComputer -Filter 'Description -like "**finance**' | ForEach {Invoke-GPUdate -Computer $_.name -RandomDelayInMinutes 0}
```

One final example: The following schedules GPUdate.exe to run immediately on all computers in the domain with an operating system name that includes the string "Vista":

```
Get-ADComputer -Filter 'OperatingSystem -like "**Vista**' | ForEach {Invoke-GPUdate -Computer $_.name -RandomDelayInMinutes 0}
```




EXAM TIP

You might see questions that use the `Get-ADComputer` or `Get-ADUser` cmdlets with the `-Filter` parameter in the way illustrated above: as a way to either search for objects with specific properties or to perform an operation on objects with specific properties.

You should review the list of user and computer properties that are searchable in Windows PowerShell. Some of these properties include `Description`, `OperatingSystem`, `LastLogonDate`, and `Name`. For a full list of the properties that you can include in such a search, run the following commands at a Windows PowerShell prompt, specifying the name of any domain computer in place of “`ComputerName`” and of any domain user in place of “`UserName`”:

```
Get-ADComputer ComputerName -Properties *
Get-ADUser UserName -Properties *
```

Remote Group Policy update and Task Scheduler

Remote Group Policy update works by remotely creating scheduled tasks for `GPUupdate.exe`. You can see these scheduled tasks for `GPUupdate` if you open Task Scheduler on the target computer and navigate in the console tree to Task Scheduler (Local)\Task Scheduler Library\Microsoft\Windows\Group Policy, as shown in Figure 9-4.

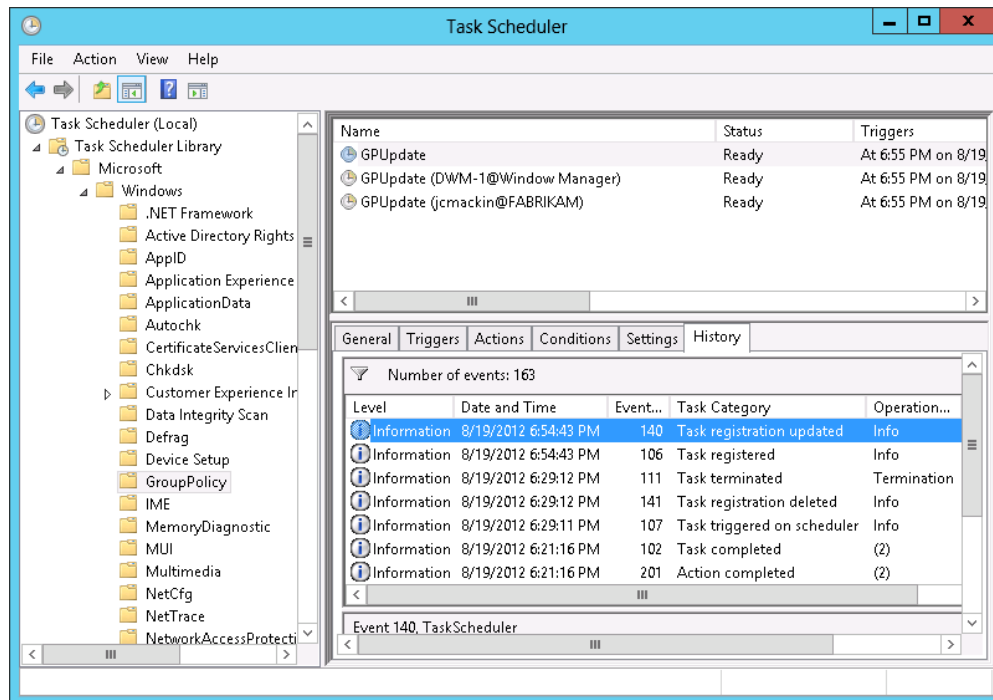


FIGURE 9-4 GPUupdate configured as a scheduled task

The connection between remote Group Policy update and Task Scheduler has implications for troubleshooting. If you are unable to successfully schedule a remote Group Policy update on a remote computer, you should verify that the Task Scheduler service is running on that remote machine. More important, for some computers, remote Group Policy update requires you to enable firewall rules related to remote scheduled tasks, as described in the next section.

Firewall rules for remote Group Policy update

Remote Group Policy update relies on remote management, which is enabled by default in Windows Server 2012 and Windows Server 2012 R2 in a domain environment. Although remote Group Policy update works by default on domain-joined computers that are started and running Windows Server 2012 and later, you might have to enable firewall rules for scheduled tasks on other operating system types, such as Windows clients or earlier versions of Windows Server that do not have Windows Management Framework 3.0 installed.

Fortunately, there's a new starter Group Policy Object (GPO) for remote Group Policy updates that makes the process of enabling the required firewall rules easy. The starter GPO, named Group Policy Remote Update Firewall Ports, is shown in Figure 9-5.

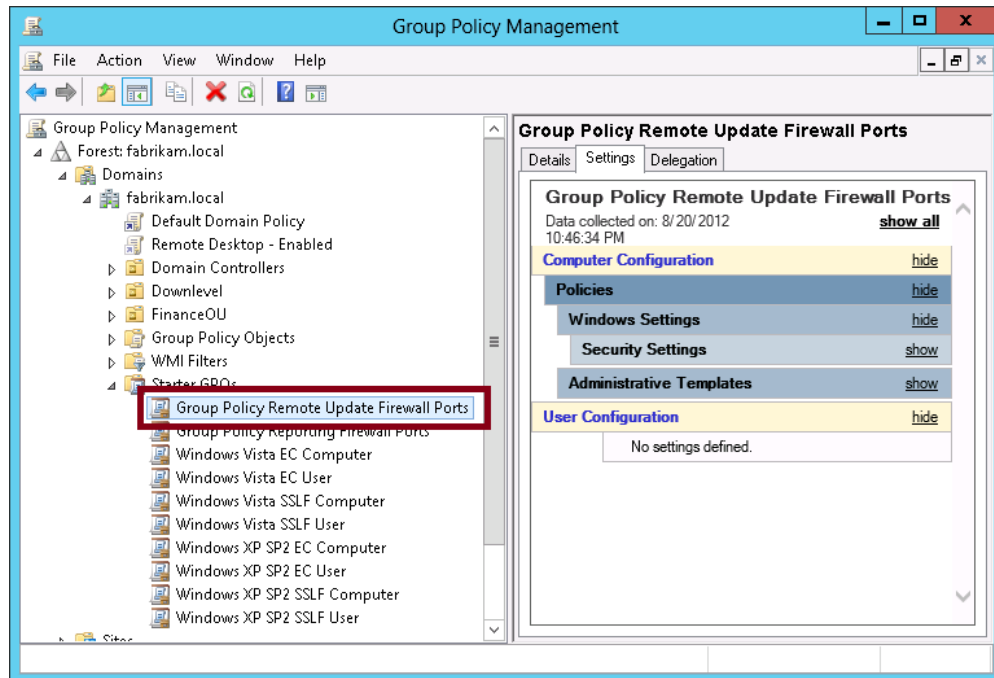


FIGURE 9-5 A starter GPO for remote Group Policy updates

After you create a GPO from the starter GPO and link the new GPO to the domain, you can view the three firewall rules enabled by this GPO, as shown in Figure 9-6:

- Both rules in the Remote Scheduled Tasks Management rule group:
 - Remote Scheduled Tasks Management (RPC)
 - Remote Scheduled Tasks Management (RPC-EPMAP)
- Windows Management Instrumentation (WMI-In)

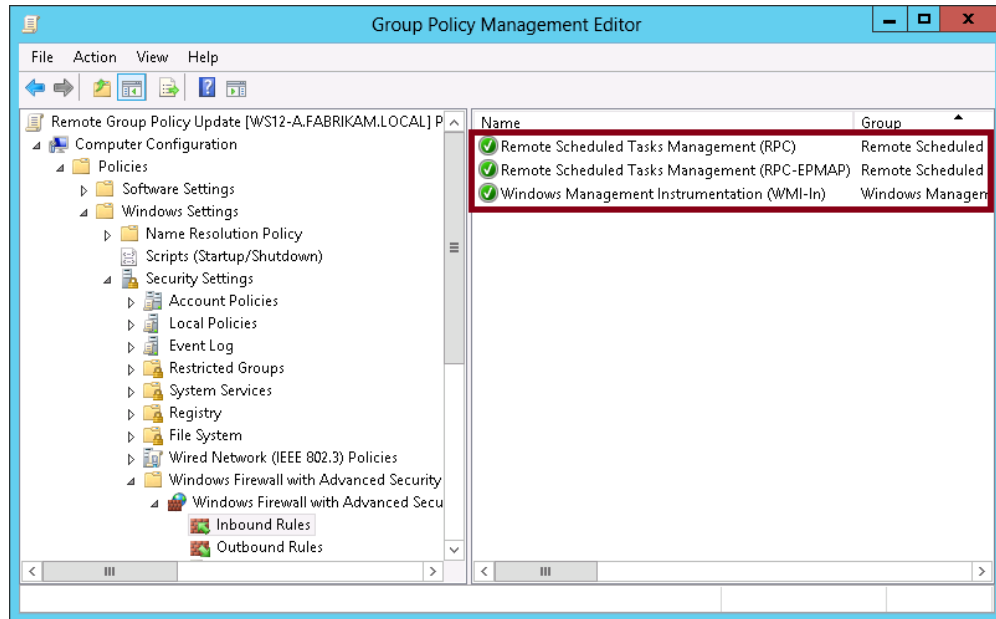


FIGURE 9-6 Inbound firewall rules for remote Group Policy update

Windows PowerShell cmdlets for Group Policy

The Group Policy questions that appear on the 70-417 exam will not be limited to remote Group Policy update. It's likely you will encounter questions that draw upon the same knowledge of Group Policy that you needed to earn your last certification—with one twist: The answer choices provided will refer to Windows PowerShell cmdlets.

To prepare for the 70-417 exam, you must understand the function of each Group Policy cmdlet by name. (You don't need to memorize their names because the test is multiple choice, not fill-in-the-blank.) These cmdlets and their associated functions are shown in Table 9-1.

TABLE 9-1 Group Policy cmdlets in Windows Server 2012 and Windows Server 2012 R2

| Group Policy Cmdlet | Function |
|----------------------------|---|
| Backup-GPO | Backs up one GPO or all the GPOs in a domain |
| Copy-GPO | Copies a GPO |
| Get-GPInheritance | Retrieves Group Policy inheritance information for a specified domain or OU |
| Get-GPO | Gets one GPO or all the GPOs in a domain |
| Get-GPOReport | Generates a report in either XML or HTML format for a specified GPO or for all GPOs in a domain |
| Get-GPPermission | Gets the permission level for one or more security principals on a specified GPO |
| Get-GPPrefRegistryValue | Retrieves one or more registry preference items under either Computer Configuration or User Configuration in a GPO |
| Get-GPRegistryValue | Retrieves one or more registry-based policy settings under either Computer Configuration or User Configuration in a GPO |
| Get-GPResultantSetOfPolicy | Outputs the Resultant Set of Policy (RSOP) information to a file, for a user, a computer, or both |
| Get-GPStarterGPO | Gets one Starter GPO or all Starter GPOs in a domain |
| Import-GPO | Imports the Group Policy settings from a backed-up GPO into a specified GPO |
| Invoke-GPUpdate | Updates Group Policy on a local computer or remote computer |
| New-GPLink | Links a GPO to a site, domain, or OU |
| New-GPO | Creates a new GPO |
| New-GPStarterGPO | Creates a new Starter GPO |
| Remove-GPLink | Removes a GPO link from a site, domain, or OU |
| Remove-GPO | Deletes a GPO |
| Remove-GPPrefRegistryValue | Removes one or more registry preference items from either Computer Configuration or User Configuration in a GPO |
| Remove-GPRegistryValue | Removes one or more registry-based policy settings from either Computer Configuration or User Configuration in a GPO |
| Rename-GPO | Assigns a new display name to a GPO |
| Restore-GPO | Restores one or all GPOs in a domain from one or more GPO backup files |
| Set-GPInheritance | Blocks or unblocks inheritance for a specified domain or OU |
| Set-GPLink | Sets the properties of the specified GPO link, including -Order (precedence), -LinkEnabled, and -Enforced |

| Group Policy Cmdlet | Function |
|-------------------------|--|
| Set-GPPermission | Grants a level of permissions to a security principal for one or all GPOs in a domain |
| Set-GPPrefRegistryValue | Configures a registry preference item under either Computer Configuration or User Configuration in a GPO |
| Set-GPRegistryValue | Configures one or more registry-based policy settings under either Computer Configuration or User Configuration in a GPO |



EXAM TIP

Even though they aren't new to Windows Server 2012 and Windows Server 2012 R2, make sure you know the commands `Dcgpofix` and `Gpfixup` and how to use them. You can use `Dcgpofix` to re-create or restore the original version of the Default Domain Policy GPO, the Default Domain Controllers Policy GPO, or both GPOs at once. `Gpfixup`, for its part, fixes domain links after a domain rename operation. Search for these commands on TechNet to learn about their syntax.

MORE INFO For more information about features related to Group Policy introduced in Windows Server 2012, visit <http://technet.microsoft.com/en-us/library/jj574108>.

Group Policy caching

Windows 8.1 introduces a new feature called Group Policy caching. Computers running Windows 8.1 will—under certain circumstances—cache to a local datastore GPOs that are read from a domain controller. (The datastore is located in `C:\Windows\system32\GroupPolicy\Datastore`.) Later, the Group Policy client will again—only under specific circumstances—read GPOs from this local datastore instead of from the domain controller.

One essential fact to clarify about Group Policy caching is that it is *not* used when the client cannot contact the domain controller. The purpose of Group Policy caching is not to act as a backup source of GPOs but to *speed up synchronous processing* of GPOs. So, to understand Group Policy caching, you first need to understand the difference between synchronous GPO processing and asynchronous GPO processing. If you've forgotten this detail about Group Policy, here's a quick summary: Synchronous processing can occur only upon startup and upon user logon. When processing is synchronous, the user doesn't see the logon screen until computer policy has 100 percent completed processing, and the user doesn't see the desktop until user policy has 100 percent completed processing. That is why synchronous processing can make the startup process and the logon process seem slow. With asynchronous processing, the logon screen and desktop can appear before all GPOs have finished being read and applied. Asynchronous processing generally appears faster as a result.

Here's how the Group Policy caching feature works in Windows 8.1: When Group Policy is processed asynchronously, GPOs are read from the domain controller and cached to the local datastore. Then, if Group Policy is processed synchronously, these cached GPOs are read from the local datastore instead of from the domain controller. (There is one exception to this rule involving Drive Mapping that is mentioned in the following note.)

So when are GPOs processed asynchronously and when are they processed synchronously? In Windows 8.1, almost all GPO processing is asynchronous. In fact, only two Group Policy *policy* settings (as opposed to Group Policy preferences) automatically trigger synchronous processing: Software Installation and Folder Redirection. A third policy setting, Disk Quotas, can trigger synchronous processing when used in conjunction with another policy setting, "Always Wait For The Network At Computer Startup And Logon."

NOTE Drive Mapping, which is a Group Policy preference setting, also triggers synchronous processing. However, this setting is not used with caching.

The only other time synchronous Group Policy processing occurs in Windows 8.1 is when you run the command **Gpupdate /sync** and then restart the computer. In this case, Group Policy processing is synchronous regardless of which settings are configured in the GPOs.

There's one other fact you need to know about Group Policy caching. It is enabled by default in Windows 8.1, but you can disable it by disabling the Group Policy setting named Configure Group Policy Caching. You can also use this policy setting to define slow link and timeout values which, if exceeded, will prevent the Windows 8.1 client from caching GPOs that are read from the domain controller.

What do you need to remember about Group Policy caching for the 70-417 exam? First, remember the purpose of the feature: to speed up synchronous Group Policy processing. Second, remember that Group Policy caching is a feature of Windows 8.1, not of Windows 8. Third, remember conceptually how it works: GPOs are cached to a local datastore during asynchronous processing and read from this cache during synchronous processing. Finally, remember that Group Policy caching is enabled by default but can be disabled through a Group Policy setting.



EXAM TIP

On the 70-417 exam you're likely to see questions about older Group Policy topics that have not changed since Windows Server 2008. As a result, make sure you still have a firm grasp on the foundational Group Policy concepts, such as the order of Group Policy processing, WMI filtering, security filtering, blocking inheritance, slow link detection, and enforcing a GPO.

Objective summary

- In Windows Server 2012 and Windows Server 2012 R2, you can use the Group Policy Management Console to schedule Group Policy to be updated on all computers in a single OU at a random point within 10 minutes. To perform this task, simply right-click an OU and select the Group Policy Update option.
- Windows Server 2012 and Windows Server 2012 R2 introduce the Invoke-GPUUpdate cmdlet in Windows PowerShell. This cmdlet allows you to update Group Policy on remote computers in a flexible way.
- In both Group Policy Management and Windows PowerShell, remote Group Policy updates work through remote task scheduling. The feature schedules GPUUpdate to run on remote computers.
- To receive scheduled tasks from remote computers, all clients and domain-joined servers running an operating system earlier than Windows Server 2012 might need to have certain inbound firewall rules enabled: both rules in the Remote Scheduled Task Management group, and Windows Management Instrumentation (WMI-In).
- You can easily enforce the inbound rules required by using the starter GPO named Group Policy Remote Update Firewall Ports. Use the starter GPO to create a new GPO that enables the required firewall rules, and then link the new GPO to the domain.
- Windows Server 2012 and Windows Server 2012 R2 include a GroupPolicy module for Windows PowerShell that includes 26 cmdlets. You need to be able to recognize the function of these cmdlets by name.
- Group Policy caching is a feature in Windows 8.1 that is enabled by default. When Group Policy is processed asynchronously, GPOs are cached in a local datastore on the client after they are read from the domain controller. Then, if Group Policy is processed synchronously, the client will read GPOs from the local datastore instead of from the domain controller. The purpose of the feature is to speed up the startup and logon processes during synchronous processing.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the “Answers” section at the end of the chapter.

1. You are a network administrator for Cpandl.com. The Cpandl.com network consists of 25 servers running Windows Server 2012 R2 and 300 clients running Windows 8.1. You are administering the network by using Remote Server Administration Tools on a computer running Windows 8.1.

You have recently implemented a change to Group Policy that affects only the servers and computers in the Marketing OU. The Marketing OU includes three servers and 40 clients. You now want to update Group Policy on all computers in the Marketing OU within the next 10 minutes.

All computers in the Marketing OU are capable of having Group Policy updated remotely. Which of the following tools should you use to accomplish this task most efficiently?

- A. Group Policy Management
 - B. Invoke-GPUdate
 - C. GPUdate.exe
 - D. Server Manager
2. You are a network administrator for Adatum.com. The Adatum.com network consists of 20 servers running either Windows Server 2012 or Windows Server 2008 R2, and 250 clients running either Windows 8 or Windows 7. You are administering the network by using Remote Server Administration Tools on a computer running Windows 8.

You have recently implemented a change to Group Policy that affects only the computers running Windows Server 2012 or Windows 8. You now want to update Group Policy on all of these computers over the next hour.

All computers running Windows Server 2012 and Windows 8 in the domain are capable of having Group Policy updated remotely. You want to update Group Policy on these computers without triggering a Group Policy update on computers running Windows Server 2008 R2 or Windows 7. Which of the following tools should you use to accomplish this task most efficiently?

- A. Update Group Policy with Group Policy Management
- B. Windows PowerShell
- C. GPUdate.exe
- D. Server Manager

- 3.** You are a network administrator for Proseware.com. The Proseware.com network consists of 20 servers running either Windows Server 2012 R2 or Windows Server 2008 R2, and 300 clients running either Windows 8.1 or Windows 7. You are administering the network by using Remote Server Administration Tools on a computer running Windows 8.1.

You have recently implemented a change to Group Policy that affects only computers in the Finance OU. When you choose to update Group Policy on all computers in the Finance OU, you receive a message indicating that the update is not successful on a number of computers that you know to be running.

You want to be able to update Group Policy on all running computers in the Finance OU without receiving an error. Which of the following actions should you take?

(Choose all that apply.)

- A.** Enable the inbound firewall rules for Remote Scheduled Tasks Management on all computers in the OU.
 - B.** Enable an inbound firewall rule for Windows Management Instrumentation (WMI) on all computers in the OU.
 - C.** Enable the Remote Registry service on all computers in the OU.
 - D.** Enable the Windows Event Collector service on all computers in the OU.
- 4.** Which of the following statements is true about Group Policy caching?
- A.** It is a feature of Windows 8 and Windows 8.1 only.
 - B.** It must be enabled in Group Policy.
 - C.** It can speed up the computer startup process when Folder Redirection is assigned through Group Policy.
 - D.** It allows a client to apply Group Policy when the connection to a domain controller is unstable.



Thought experiment

Configuring and managing Group Policy at Woodgrove Bank

You are a network administrator for Woodgrove Bank. The woodgrovebank.com private network spans seven branch offices in seven cities throughout New York State. The network includes 50 servers running Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008, and 700 clients running Windows 8, Windows 7, or Windows Vista. Each of the seven offices is assigned its own OU in Active Directory and at least two domain controllers, all of which are running Windows Server 2012.

Your manager has asked you to investigate the requirements for remote updates to Group Policy. He wants to implement this capability within the next few weeks.

With this background information in mind, answer the following questions. You can find the answers to these questions in the “Answers” section.

- 1.** How can you most efficiently create the firewall rules required to allow all of your servers and clients to receive remote Group Policy updates?
- 2.** Besides firewall settings, which other setting could you enforce through Group Policy to ensure that your servers and clients will be able to receive remote Group Policy updates?
- 3.** How can you most efficiently update Group Policy on all computers in one of the seven branch offices?
- 4.** Your manager wants to be able to force all running computers in the domain to update Group Policy at random points over the course of four hours. He wants you to write a Windows PowerShell command for this very purpose. Which Windows PowerShell command would achieve this goal when executed?
- 5.** Your manager wants to be able to force all computers in the domain that are started and running Windows 8 to update Group Policy within 10 minutes. He wants you to write a Windows PowerShell command for this very purpose. Which Windows PowerShell command would achieve this goal when executed?

Answers

This section contains the answers to the Objective Review and the Thought Experiment.

Objective 9.1: Review

1. Correct answer: A

- A. Correct:** The Group Policy Management Console allows you to schedule an update Group Policy on all computers in an OU. To do so, right-click the OU and select Group Policy Update. The update occurs on all computers in the OU within 10 minutes of choosing this option.
- B. Incorrect:** The Invoke-GPUdate cmdlet by itself does not allow you to perform a Group Policy update on all computers in an OU. You can create a scripted command in Windows PowerShell that combines Get-ADComputer and Invoke-GPUdate to achieve this same result, but it is not the most efficient solution if you are managing the network by using Remote Server Administration Tools.
- C. Incorrect:** GPUdate.exe refreshes Group Policy on one computer only. To update Group Policy on all computers in the Marketing OU, you would need to run this command 43 times. This solution is much less efficient than using the Group Policy Management Console.
- D. Incorrect:** Server Manager does not provide an option to refresh Group Policy on multiple computers at once.

2. Correct answer: B

- A. Incorrect:** You can use Group Policy Management to remotely update Group Policy only in a particular OU. The update applies to all computers in the OU, and the update occurs within 10 minutes. You cannot use Group Policy Management to remotely update Group Policy on computers running any particular operating systems, and you cannot use this tool to specify that these updates should occur over the next hour.
- B. Correct:** You can use a single scripted command in Windows PowerShell that will invoke a remote Group Policy update over the next hour only on computers running either Windows Server 2012 or Windows 8. The following command is one way to accomplish this task:

```
Get-ADComputer -Filter {(OperatingSystem -like "*Windows 8*")} -or (OperatingSystem -like "*Windows Server 2012*")} | ForEach {Invoke-GPUdate -Computer $_.name -RandomDelayInMinutes 60}
```

- C. Incorrect:** GPOupdate.exe refreshes Group Policy on one computer only. To update Group Policy on all computers running Windows Server 2012 and Windows 8, you would need to run this command many times either locally on each computer or through a Remote Desktop connection. This solution is much less efficient than using Windows PowerShell.
 - D. Incorrect:** Server Manager does not provide an option to refresh Group Policy on multiple computers at once.
- 3. Correct answers:** A, B
- A. Correct:** Certain operating systems such as clients and older versions of Windows Server without Windows Management Framework 3.0 do not allow you to remotely update Group Policy by default. To allow remote Group Policy updates, you need to enable inbound ports for Remote Scheduled Tasks Management and WMI.
 - B. Correct:** An inbound rule allowing WMI is one of the three firewall rules needed to allow various clients to receive remote Group Policy updates.
 - C. Incorrect:** This service enables remote users to modify registry settings on the local computer. It is not needed to allow a remote computer to schedule GPOupdate.exe to run locally.
 - D. Incorrect:** This service manages persistent subscriptions to events from certain remote sources. It is not needed to allow a remote computer to schedule GPOupdate.exe to run locally.
- 4. Correct answer:** C
- A. Incorrect:** Group Policy caching is a feature that is new to Windows 8.1.
 - B. Incorrect:** Group Policy caching is enabled by default. It doesn't need to be enabled in Group Policy.
 - C. Correct:** The purpose of Group Policy caching is to speed up synchronous processing. Folder Redirection is a policy setting that triggers synchronous processing.
 - D. Incorrect:** Group Policy caching is not used as a secondary source of GPOs when the connection to a domain controller is unstable.

Thought experiment

1. Use the Group Policy Remote Update Firewall Ports starter GPO to create a new GPO. Link the new GPO to the domain.
2. You could use Group Policy to ensure that the Task Scheduler service is set to Automatic on all computers in the domain.
3. Use the Group Policy Management Console to update Group Policy on the OU corresponding to the branch office.
4.

```
Get-ADComputer -Filter * | ForEach {Invoke-GPUdate -Computer $_.name -RandomDelayInMinutes 240}
```
5.

```
Get-ADComputer -Filter 'OperatingSystem -like "*Windows 8*"' | ForEach {Invoke-GPUdate -Computer $_.name}
```

Configure and manage high availability

The Configure and Manage High Availability domain relates to failover clustering and the live migration of virtual machines.

A failover cluster, as you know, is a group of two or more computers that work together to help ensure the availability of a service or application. (In Windows Server 2012 and Windows Server 2012 R2, clustered services and applications are known as roles.) There are a number of improvements in failover clustering in Windows Server 2012 and Windows Server 2012 R2, beginning with scalability: Failover clusters now support up to 64 nodes (as opposed to 16 in Windows Server 2008 R2). Failover clusters in Windows Server 2012 and Windows Server 2012 R2 also support many new enhancements, such as Cluster-Aware Updating, role priority, VM monitoring, and node drain. Windows Server 2012 R2 specifically, for its part, has brought yet another round of new features and improvements, including Active Directory-detached clusters, Dynamic Witness, and virtual machine network health detection.

Live migration of virtual machines (VMs) used to be restricted to failover clusters, but this feature has been expanded to provide uninterrupted availability in all domain contexts.

To learn about the new developments in failover clustering and live migration for the 70-417 exam, it's best (as always) to implement these features in a test environment. For failover clustering, the good news is that you can now perform all of this testing on a single server with VMs running in Hyper-V and using the new built-in iSCSI Target feature for shared storage. For live migration, you will need two physical servers.

Objectives in this chapter:

- Objective 10.1: Configure failover clustering
- Objective 10.2: Manage failover clustering roles
- Objective 10.3: Manage virtual machine (VM) movement

Objective 10.1: Configure failover clustering

Failover clustering in Windows Server 2012 and Windows Server 2012 R2 introduce many improvements. The topics covered here are the ones most likely to appear on the 70-417 exam.

This section covers the following topics:

- Cluster storage pools
- Cluster shared volumes (CSVs)
- Virtual hard disk sharing for guest clusters in Windows Server 2012 R2
- Dynamic quorum
- Dynamic witness in Windows Server 2012 R2
- Node drain
- Cluster-Aware Updating (CAU)
- Active Directory-detached clusters in Windows Server 2012 R2

Cluster storage pools

In Windows Server 2012 and Windows Server 2012 R2, you can now draw from data storage provided by a Serial Attached SCSI (SAS) disk array to create one or more storage pools for a failover cluster. These storage pools are similar to the ones you can create for an individual server by using Storage Spaces. As with the Storage Spaces feature, you can use storage pools in a failover cluster as source from which you can then create virtual disks and finally volumes.

To create a new storage pool, in Failover Cluster Manager, navigate to Failover Cluster Manager*Cluster Name*\Storage\Pools, right-click Pools, and then select New Storage Pool from the shortcut menu, as shown in Figure 10-1. This step starts the same New Storage Pool Wizard used with Storage Spaces. (In fact, if you have a shared SAS disk array, you can use Server Manager to create the pool and use the Add Storage Pool option to add it to the machine.) After you create the pool, you need to create virtual disks from the new pool and virtual volumes from the new disks before you can use the clustered storage space for hosting your clustered workloads.

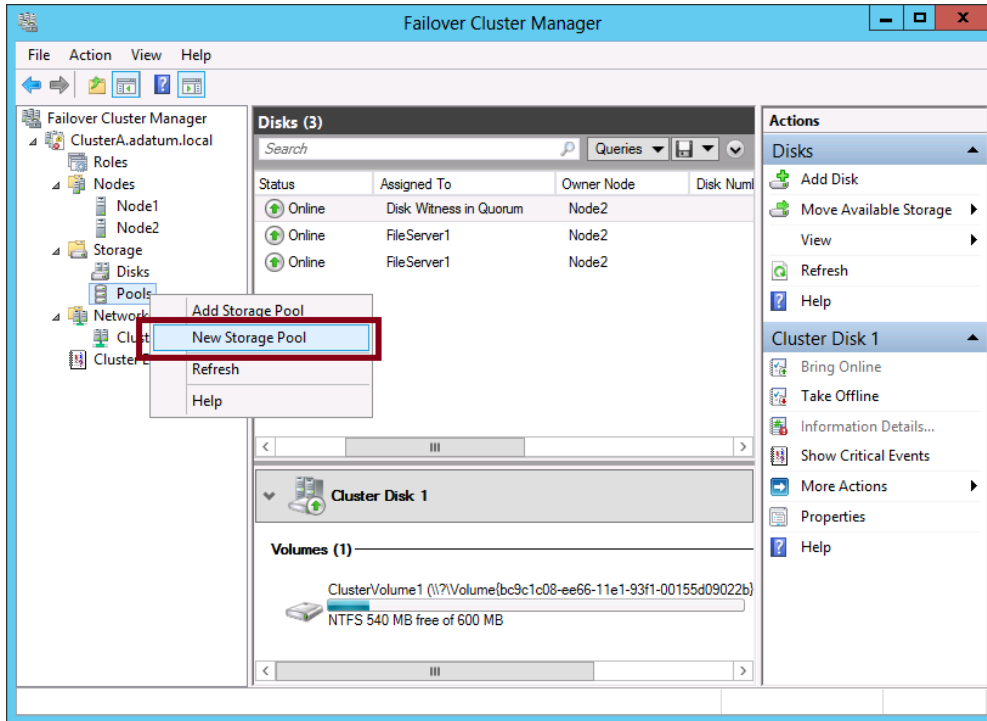


FIGURE 10-1 Creating a new storage pool for a cluster

The availability of storage pools for failover clusters has implications for the 70-417 exam, especially because these storage pools have a number of requirements and restrictions that could easily serve as the basis for a test question. Note the following requirements for failover cluster storage pools:

- A minimum of three physical drives, with at least 4 GB capacity each.
- Only SAS-connected physical disks are allowed. No additional layer of RAID (or any disk subsystem) is supported, whether internal or external.
- Fixed provisioning only for virtual disks; no thin provisioning.
- When creating virtual disks from a clustered storage pool, only simple and mirror storage layouts are supported. Parity layouts are not supported.
- The physical disks used for a clustered pool must be dedicated to that one pool. Boot disks should not be added to a clustered pool.

MORE INFO For more information about cluster storage pools, visit <http://blogs.msdn.com/b/clustering/archive/2012/06/02/10314262.aspx>.

Cluster shared volumes (CSVs)

Cluster shared volumes (CSVs) are a new type of storage used only in failover clusters. CSVs first appeared in Windows Server 2008 R2, so if you earned your last certification before this important feature was introduced, you might have missed CSVs completely. In this case, you need to understand the basics about them before taking the 70-417 exam.

The biggest advantage of CSVs is that they can be shared by multiple cluster nodes at a time. This is not normally possible with shared storage. In fact, even different volumes created on the same logical unit number (LUN) cannot normally be shared by different cluster nodes at the same time.

CSVs achieve this shared access of volumes by separating the data from different nodes into virtual hard disk (VHD) files. Within each shared volume, multiple VHDs are stored, each used as the shared storage for a particular role. The CSVs containing these VHDs are then mapped to a common, integrated namespace on all nodes in the cluster. On every failover cluster configured with CSVs, the CSVs appear on every node as subfolders in the \ClusterStorage folder on the system drive. Example pathnames are C:\ClusterStorage\Volume1, C:\ClusterStorage\Volume2, and so on. The volume objects in the path act as links to remote LUNs, so you are not limited by the size of your local drive.

CSVs are formatted with NTFS (or, optionally, ReFS in Windows Server 2012 R2), but to distinguish them from normal NTFS and ReFS volumes, the Windows Server 2012 and Windows Server 2012 R2 interfaces display these volumes as formatted with CSVFS, or the Cluster Shared Volume File System. An example of a CSV is shown in Figure 10-2.

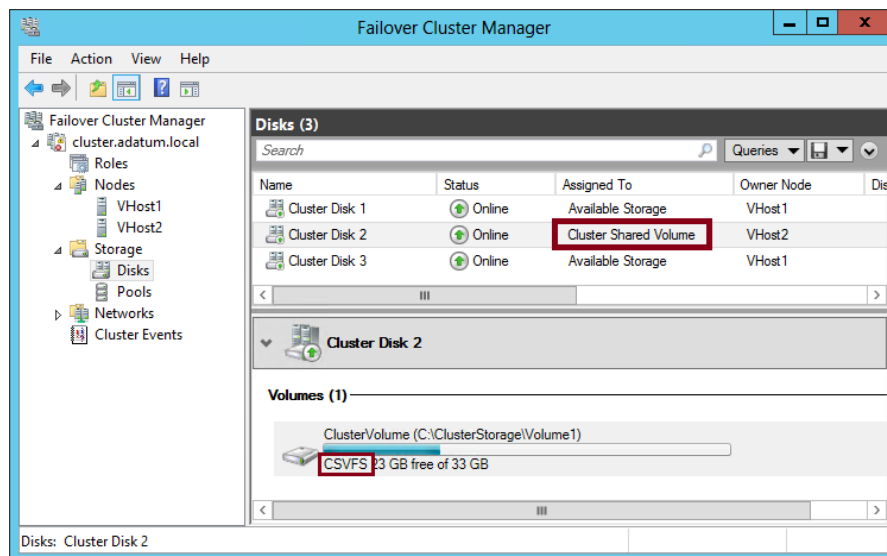


FIGURE 10-2 A cluster shared volume

To create a CSV in Windows Server 2012 or Windows Server 2012 R2, first provision a disk from shared storage, such as from an iSCSI target. Use Server Manager to create a volume from this disk, as shown in Figure 10-3.

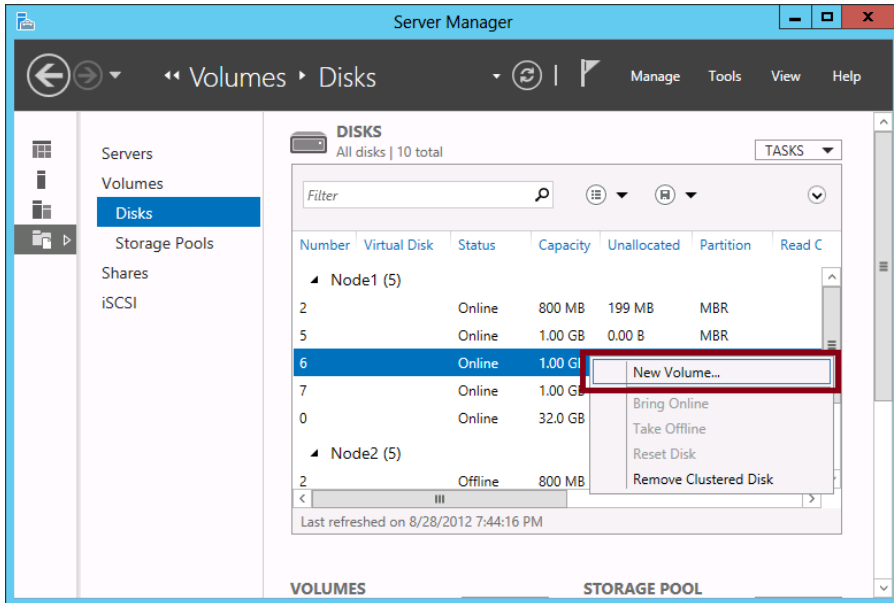


FIGURE 10-3 Creating a new volume in Server Manager

Assign the new volume to the desired failover cluster, as shown in Figure 10-4. (The name of the cluster appears as a server name.)

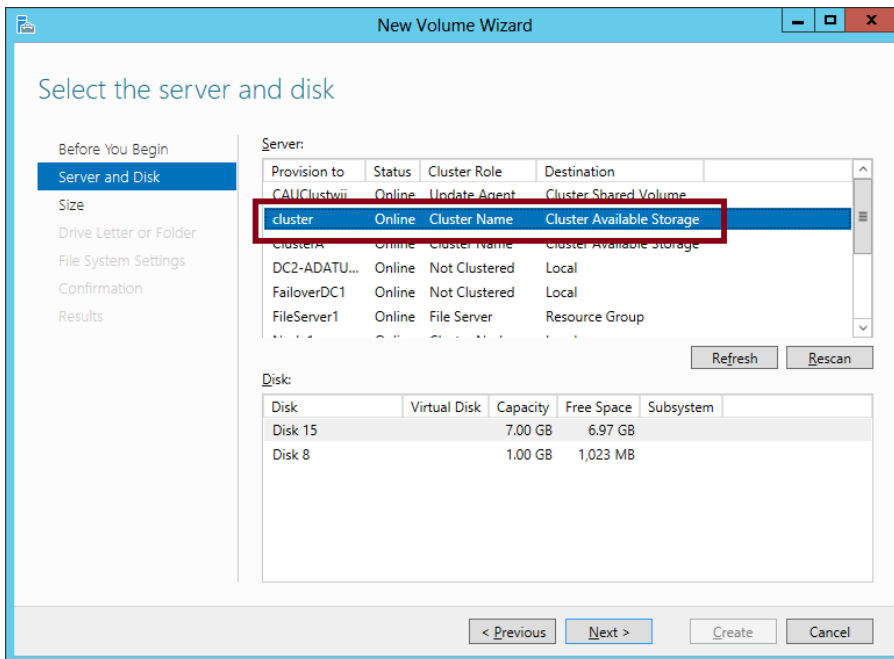


FIGURE 10-4 Assigning a new volume to a cluster

In Failover Cluster Manager, the new volume will appear as a disk. Right-click the disk and select Add To Cluster Shared Volumes from the shortcut menu, as shown in Figure 10-5.

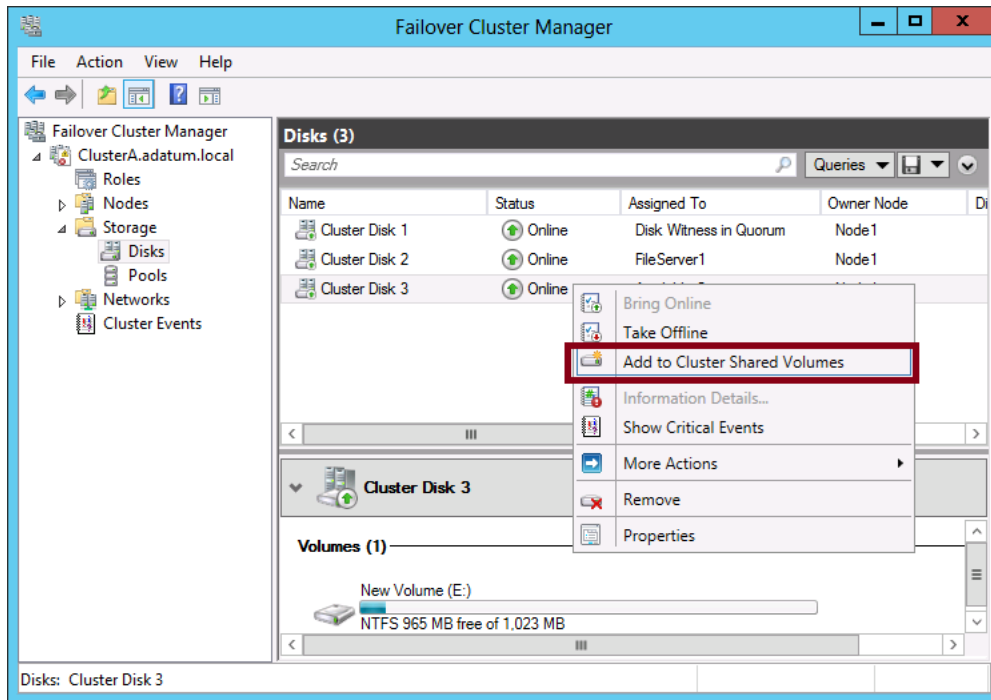


FIGURE 10-5 Adding a volume to cluster shared volumes

In Windows Server 2008 R2, CSVs were used as storage for only one type of workload hosted in a failover cluster: a highly available VM. In Windows Server 2012 and Windows Server 2012 R2, CSVs are now also used as the only storage type for a new role, the Scale-Out File Server, which is described later in this chapter. Another important use for CSVs is with live migration in failover clusters (another feature described later in this chapter). Although CSVs are not required for live migration, they are highly recommended because they optimize the performance of the migration and reduce downtime to almost zero.

How might CSVs appear on the 70-417 exam? If there's a question about CSVs directly, it could come in the form of a requirement that states you want to "minimize administrative overhead" when designing storage for a highly available VM. More likely, you will simply see CSV mentioned in the setup to a question that isn't just about CSVs.

NOTE Why are Cluster shared volumes useful? Here is the original problem CSVs were designed to solve: In Windows Server 2008 and earlier versions of Windows Server, only one cluster node could access a LUN at any given time. If any application, service, or VM connected to a LUN failed and needed to be moved to another node in the failover cluster, every other clustered application or VM on that LUN would also need to be failed over to a new node and potentially experience some downtime. To avoid this problem, each clustered role was typically connected to its own unique LUN as a way to isolate failures. This strategy created another problem, however: a large number of LUNs that complicated setup and administration.

With CSVs, a single LUN can be accessed by different nodes at a time, as long as the different nodes are accessing distinct VHDs or VHDXs on the LUN. You can run these roles on any node in the failover cluster and when the role fails, it can fail over to any other physical node in the cluster without affecting other roles (services or applications) hosted on the original node. CSVs thus add flexibility and simplify management.

Virtual hard disk sharing for guest clusters in Windows Server 2012 R2

Before Windows Server 2012 R2, when you wanted to provision storage for a guest cluster (that is, a failover cluster in which the nodes are virtualized), you needed to provision storage from within the operating system of these Hyper-V guests by connecting to an underlying storage infrastructure, such as an iSCSI or Fibre Channel SAN.

However, allowing visibility of an underlying storage infrastructure from a guest OS in Hyper-V is not always ideal. For example, if you are a cloud hosting company you might want to allow customers to create a failover cluster from multiple VMs hosted in your cloud. In this case, you want to be able to provide shared storage to these VMs on demand without allowing customers to see your underlying SAN fabric.

IMPORTANT A *guest cluster* is a virtualized cluster in which the nodes are VMs and the failover clustering feature is installed in the guest operating system of each of those VMs. In contrast, the terms *physical host cluster*, *Hyper-V host cluster*, or even just *host cluster* can be used to refer to a failover cluster in which the nodes are physical computers and the failover clustering feature is installed in the host operating system of each node.

With VHDX sharing, Windows Server 2012 R2 introduces a way to provide shared storage to guest clusters in a way that hides the underlying storage infrastructure. VHDX sharing requires the Hyper-V hosts themselves that are supporting the guest cluster also to be configured in their own failover cluster, so as a prerequisite you must have a (virtualized) guest cluster on top of a (physical) host cluster. After you configure the shared VHDX from the physical host, the virtual hard disk appears as a raw SAS disk that is eligible to be added to the guest cluster from within the guest operating system.

Here's how to configure the shared VHDX: First, on one of the nodes of the physical host cluster, provision a new disk from shared storage and add it to CSVs as described in the preceding section, "Cluster Shared Volumes." (You can also provide a new volume to the host cluster by means of an SMB 3.0 share on a Scale-Out File Server.) Next, in Hyper-V Manager, open the settings of the guest VM that is acting as the management node in your guest cluster. Select the SCSI Controller settings and choose to add a new hard disk. The new SCSI disk must be a VHDX file and must be saved in the path to the new CSV or SMB 3.0 share that has been configured for the physical host cluster. Finally, in the settings of the guest VM in Hyper-V Manager, expand the new SCSI disk, select Advanced Features and then click Enable Virtual Hard Disk Sharing.

When you start the VM next, a new raw SAS disk will appear in Disk Management and Server Manager. You can use Failover Cluster Manager in the guest to add this disk to cluster storage in the guest cluster.

Remember the following points about VHDX sharing:

- The shared virtual disk can act as a data disk or witness disk for the guest cluster. It cannot be used as the operating system disk.
- The shared virtual disk must be a VHDX file, not a VHD file.
- The shared virtual disk must be attached to the virtual machine's SCSI controller, not its IDE controller.
- When you attach the new virtual hard disk to the SCSI controller, don't click Apply until you have selected the Enable Virtual Hard Disk Sharing check box in Advanced Features.

Dynamic quorum

Dynamic quorum configuration is a new feature in Windows Server 2012 and Windows Server 2012 R2 in which the number of votes required to reach quorum automatically adjusts to the number of active nodes in the failover cluster. If one or more nodes shuts down, the number of votes to reach quorum in Windows Server 2012 and Windows Server 2012 R2 changes to reflect the new number of nodes. With dynamic quorum, a failover cluster can remain functional even after half of its nodes fail simultaneously. In addition, it's possible with dynamic quorum for the cluster to remain running with only one node remaining.

The applicability of dynamic quorum to the 70-417 exam is uncertain, mainly because by definition this new feature doesn't require you to remember any configuration settings. However, if you see a question in which nodes in the cluster are running an earlier version of Windows and "you want to have the number of votes automatically adjust based on the number of available nodes at any one time," you know you need to upgrade all nodes to Windows Server 2012 or later. In addition, remember that dynamic quorum works only with the following quorum configurations and not with the Disk Only quorum configuration:

- Node Majority
- Node and Disk Majority
- Node and File Share Majority



EXAM TIP

Remember that the cluster quorum settings determine the number of elements in a failover cluster that must remain online for it to continue running. You access these settings by right-clicking a cluster in the Failover Cluster Manager console tree, selecting More Actions and then clicking Configure Cluster Quorum Settings.

Dynamic witness in Windows Server 2012 R2

A witness in a failover cluster, as you remember, is a disk or file share that holds information about the cluster configuration. The witness can act as a tiebreaker vote that helps to achieve quorum when only half of the nodes in a cluster can communicate with each other. Because of this tiebreaker role for the witness, Microsoft has recommended—until Windows Server 2012 R2—that you configure a witness disk or file share only when your cluster contained an even number of nodes.

However, the extra vote provided by the witness could occasionally cause a problem in versions of Windows Server before Windows Server 2012 R2. The witness did ensure that a majority of votes was possible when exactly 50 percent of the nodes were online and could communicate with each other. However, if one of the original nodes failed, the witness raised the possibility that the remaining votes could be split evenly between two groups, preventing either the majority needed to achieve quorum.

This potential problem is solved in Windows Server 2012 R2 through *dynamic witness*. With dynamic witness, the witness is active only when the number of online nodes is even. When the number of active nodes is odd, the witness loses its vote.

Because of the new dynamic nature of witness disks and witness file shares in Windows Server 2012 R2, Microsoft recommends that you always configure a witness for failover clusters in Windows Server 2012 R2. Unlike in previous versions of Windows Server, a witness is now always included as a step when you run the Configure Cluster Quorum Wizard and choose the default quorum configuration. The witness is then active only when needed.

Node drain

Node drain is a feature new to Windows Server 2012 and Windows Server 2012 R2 that simplifies the process of shutting a node down for maintenance. In previous versions of Windows, if you wanted to bring a node down for maintenance, you first needed to pause the node and then move all hosted applications and services (now called roles) over to other nodes. With node drain, these two steps are combined into one.

To prepare a node to be shut down for maintenance in this way, first navigate to the Nodes container in the Failover Cluster Manager console tree. Then right-click the node you want to shut down in the details pane, point to Pause and then select Drain Roles, as shown in Figure 10-6.

To achieve this same result by using Windows PowerShell, use the `Suspend-ClusterNode` cmdlet.

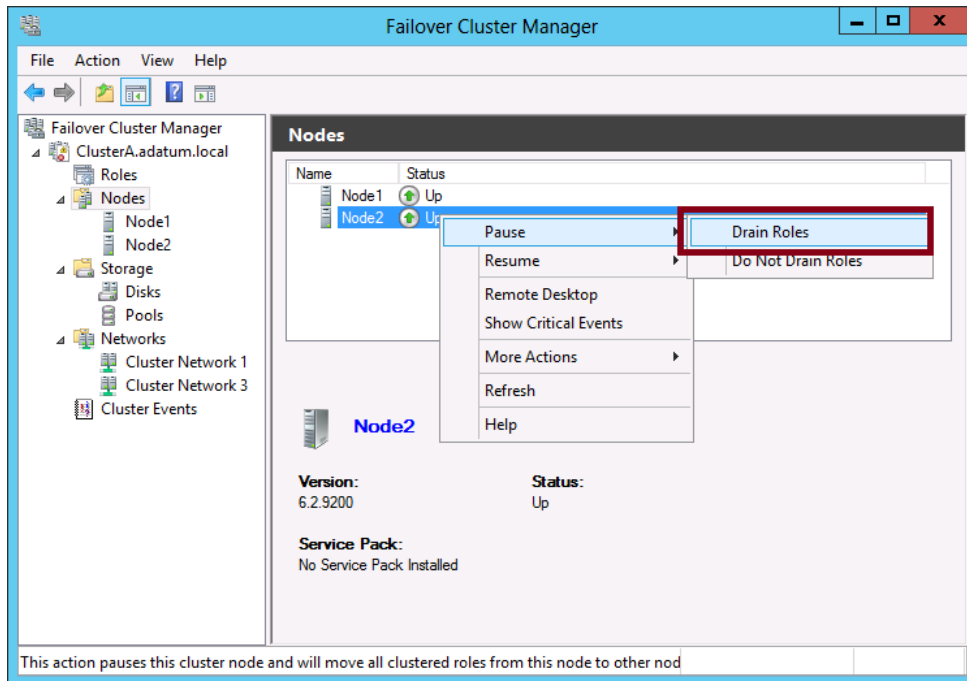


FIGURE 10-6 Draining roles from a node

Automatic node drain on shutdown in Windows Server 2012 R2

In Windows Server 2012 R2, if you choose to shut down a node before draining the roles, the roles will automatically be live-migrated to other nodes before the shutdown is performed. However, this feature is designed mainly as a backup safety mechanism. It is preferable to drain roles manually to make sure these roles have migrated before shutting down a node.

Cluster-aware updating (CAU)

Cluster-aware updating (CAU) is a new feature in Windows Server 2012 and Windows Server 2012 R2 that addresses the difficulty of performing software updates on failover cluster nodes. This difficulty stems from the fact that updating software normally requires a system restart. To maintain the availability of services hosted on failover clusters in previous versions of Windows, you needed to move all roles off one node, update the software on that node, restart the node and then repeat the process on every other node, one at a time. Windows Server 2008 R2 failover clusters could include up to 16 nodes, so this process sometimes had to be repeated as many times. In Windows Server 2012 and Windows Server 2012 R2, failover clusters can scale up to 64 nodes. At this point, the older, manual method of updating software on failover clusters is simply no longer practical.

Instead, Windows Server 2012 and Windows Server 2012 R2 automate the process of updating software for you. To initiate the process of updating a failover cluster, simply right-click the cluster in the list of servers in Server Manager and then select Update Cluster from the shortcut menu, as shown in Figure 10-7.

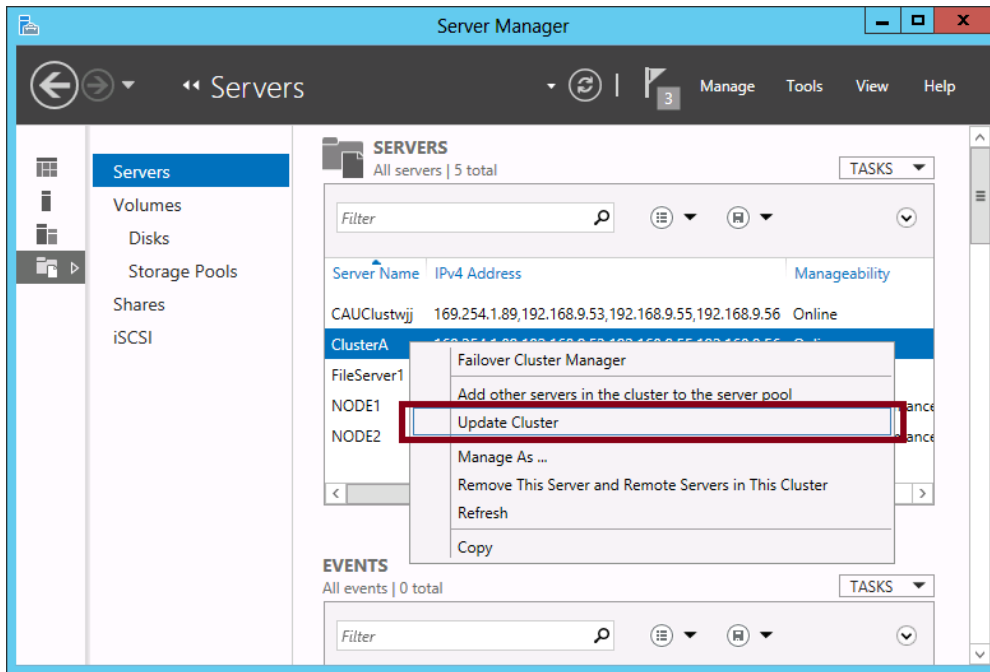


FIGURE 10-7 Manually updating a cluster

By default, only updates configured through Windows Update are performed. Updates are received as they normally would be, either directly from the Microsoft Update servers or through Windows Software Update Services (WSUS), depending on how Windows Update is configured. Beyond this default functionality, CAU can be extended through third-party plug-ins so that other software updates can also be performed.

MORE INFO For more information about CAU plugins work, visit <http://technet.microsoft.com/en-us/library/jj134213>.

The preceding step shows how to trigger an update to a cluster manually. Triggering updates manually might be too straightforward a task to appear on the 70-417 exam. More likely, you could see a question about configuring self-updates. You can access these self-update configuration settings in Failover Cluster Manager by right-clicking the cluster name in the console tree, pointing to More Actions, and then selecting Cluster-Aware Updating, as shown in Figure 10-8.

NOTE Manual updating is called *remote updating mode* when you use this method to update a failover cluster from a remote machine on which the failover cluster management tools have been installed.

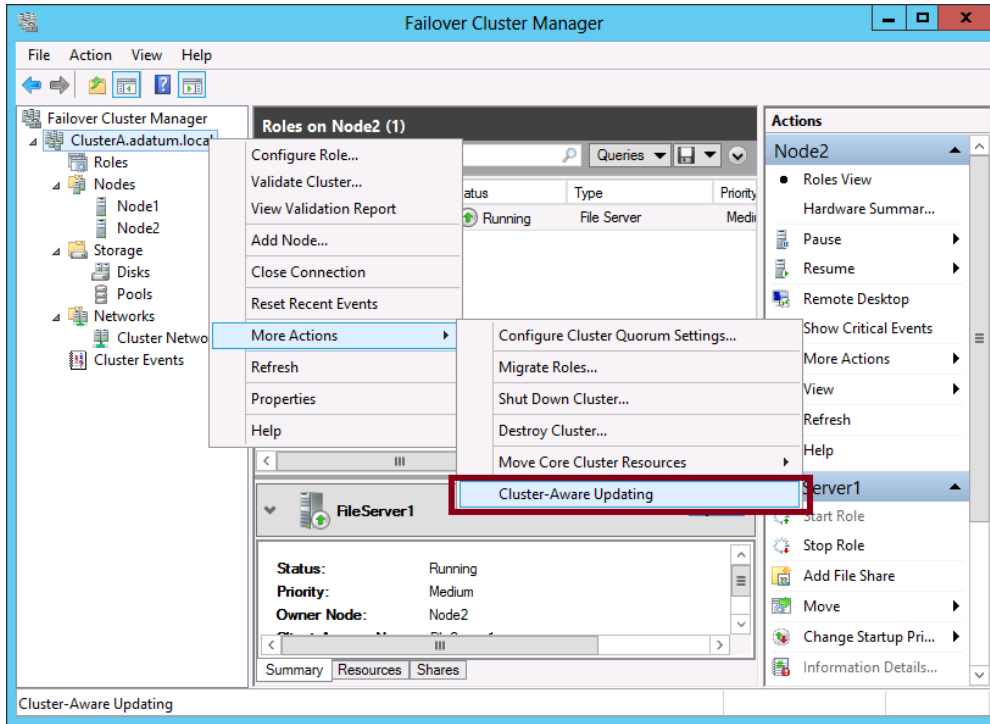


FIGURE 10-8 Opening CAU actions

This step opens the Cluster-Aware Updating dialog box, shown in Figure 10-9.

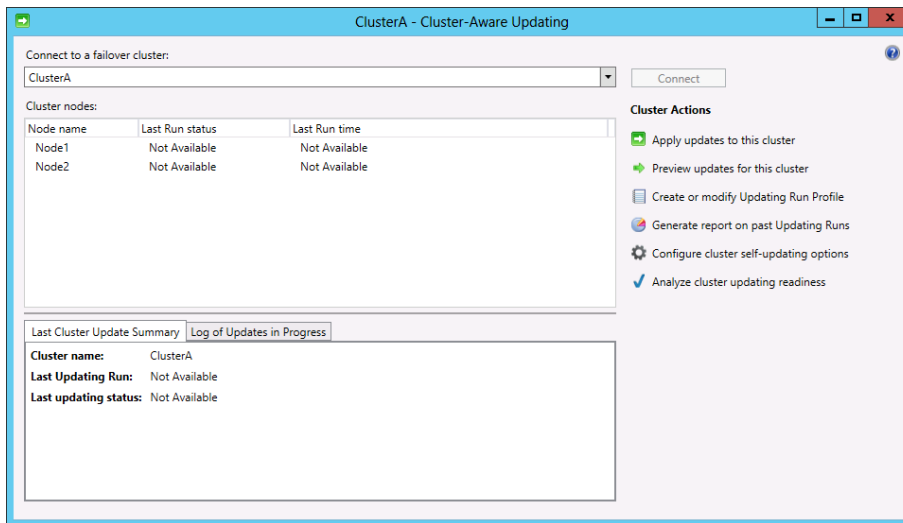


FIGURE 10-9 CAU actions

To configure self-updating for the cluster, click Configure Self-Updating Options beneath Cluster Actions. This step will open the Configure Self-Updating Options Wizard. You can enable self-updating on the cluster on the second (Add Clustered Role) page of the wizard by selecting the option to add the CAU clustered role, with self-updating mode enabled (shown in Figure 10-10).

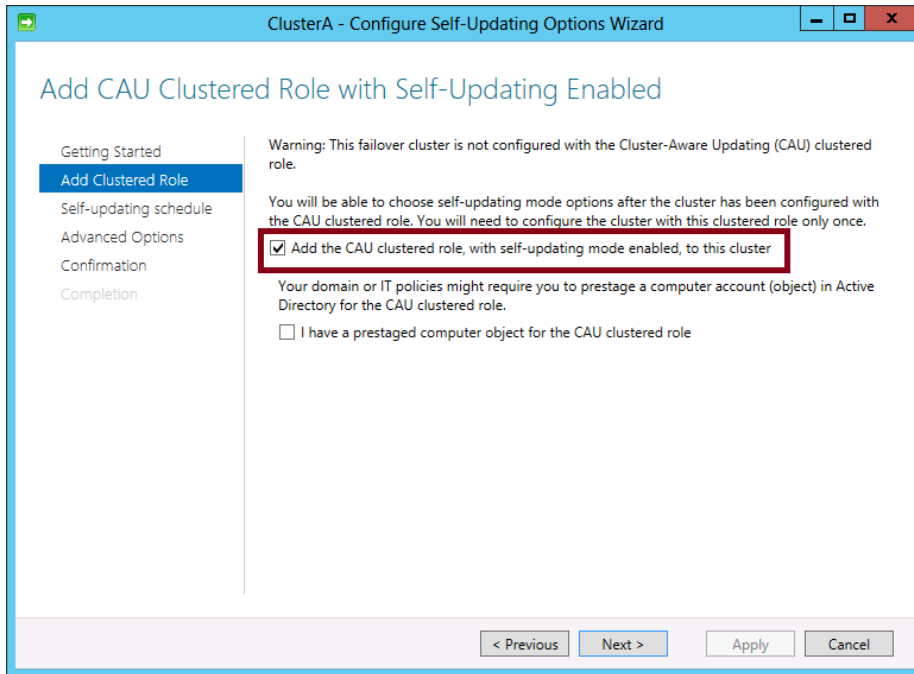


FIGURE 10-10 Enabling self-updating mode for CAU

NOTE The Cluster-Aware Updating dialog box also allows you to perform the following actions:

- Apply updates to the cluster
- Preview updates for the cluster
- Create or modify Updating Run Profile directly (without running a wizard)
- Generate a report on past Updating Runs
- Analyze cluster updating readiness

The third (Self-Updating Schedule) page of the wizard lets you specify a schedule for updating. The fourth (Advanced Options) page lets you change profile options, as shown in Figure 10-11. These profile options let you set time boundaries for the update process and other advanced parameters.

MORE INFO For more information about profile settings for CAU, visit <http://technet.microsoft.com/en-us/library/jj134224.aspx>.

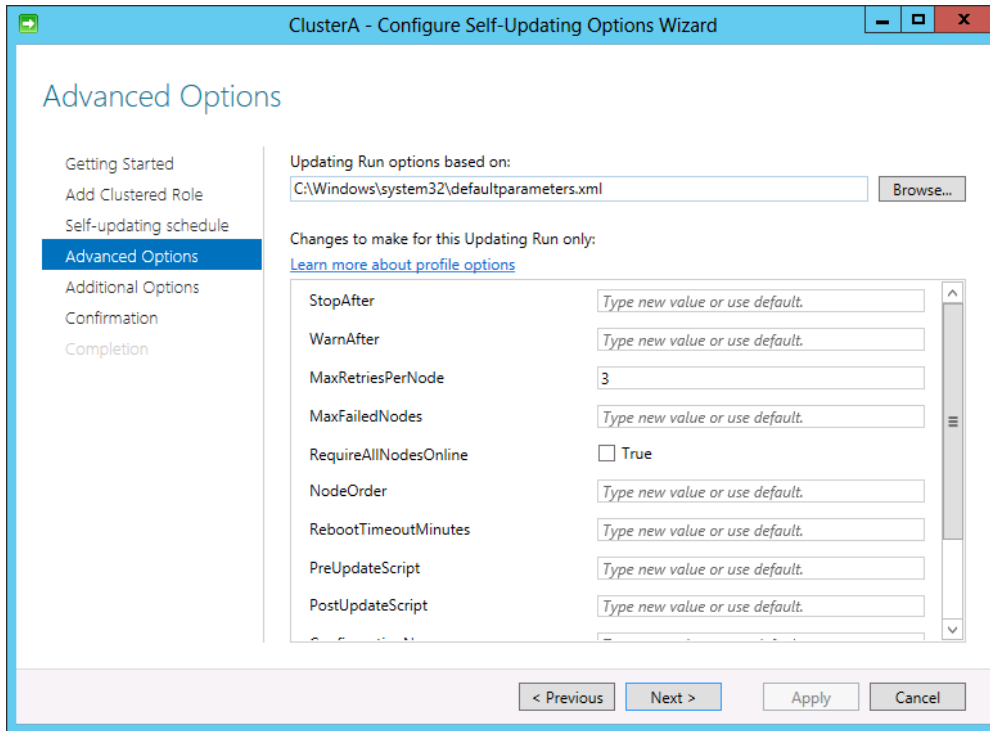


FIGURE 10-11 Configuring advanced options for Cluster-Aware Self-Updating

MORE INFO For a more detailed description of CAU, visit <http://blogs.technet.com/b/filecab/archive/2012/05/17/starting-with-cluster-aware-updating-self-updating.aspx>.



EXAM TIP

Expect to see a question or two that tests your knowledge of Windows PowerShell cmdlets for failover clusters. For a complete list of those cmdlets, visit "Failover Clusters Cmdlets in Windows PowerShell" at <http://technet.microsoft.com/en-us/library/hh847239.aspx>.

Active Directory-Detached Clusters in Windows Server 2012 R2

In the first release of Windows Server 2012 (not Windows Server 2012 R2) and earlier versions of Windows Server, every failover cluster requires in Active Directory its own computer object, with a name corresponding to the name of the cluster. This computer object is created auto-

matically when the cluster is created. However, in Windows Server 2012 R2, you can now also create a failover cluster that doesn't require a computer account in AD DS.

The purpose of avoiding this requirement is to reduce the complexity of deploying, managing, and maintaining the cluster. For example, you will no longer require elevated privileges in the domain to create the cluster. You also avoid the cluster failure that would arise from an accidental deletion of the cluster's computer account.

This new type of failover cluster that doesn't require a computer account is known as an *Active Directory-detached cluster*. Note that even though these failover clusters do not require computer objects in Active Directory, they still depend on Active Directory for other services such as authentication. Cluster nodes therefore must still be members of the same Active Directory domain. All nodes, in addition, must be running Windows Server 2012 R2.



EXAM TIP

Active Directory-detached clusters are not recommended for scenarios that require Kerberos authentication. Authentication of the cluster network name uses NTLM, not Kerberos.

NOTE The only workload that is currently supported without qualification for Active Directory-detached clusters is SQL Server. (File server and Hyper-V workloads are currently supported but not recommended.)

To create an Active Directory-detached cluster, you use Windows PowerShell, not Failover Cluster Manager. Use the `New-Cluster` cmdlet to create the new cluster, along with the `-AdministrativeAccessPoint` parameter set to a value of `Dns`. (The `Dns` value ensures that the cluster name will be registered in DNS. To avoid registering these DNS names, you would use the `None` value)

For example, the following command creates a two-node Active Directory-detached cluster with the name `Cluster1`.

```
New-Cluster Cluster1 -Node Node1,Node2 -StaticAddress 192.168.2.10 -NoStorage  
-AdministrativeAccessPoint Dns
```

After you create the cluster, you can configure and manage it in Failover Cluster Manager. (The cluster may take several minutes to appear in the interface.)

To determine whether a cluster is Active Directory-detached, type the following command at a Windows PowerShell prompt on one of the cluster nodes:

```
(Get-Cluster).AdministrativeAccessPoint
```

If the output is `Dns`, then the cluster is Active Directory-detached. If the output is `ActiveDirectoryAndDns`, then the cluster is a normal cluster integrated with Active Directory.

Configuring Cluster Properties in Windows PowerShell

You can use Windows PowerShell to view and set properties on entire clusters or on individual nodes. Many of these cluster and node properties, in fact, can only be viewed and configured through Windows PowerShell. To view or configure a cluster property or node property, you can expose it through the `Get-Cluster` cmdlet.

WitnessDynamicWeight

For example, the `WitnessDynamicWeight` cluster property relates to the Dynamic Witness feature and determines whether the witness assigned to the cluster currently has a quorum vote. To determine the current status of the quorum witness vote, type the following command at an elevated Windows PowerShell prompt:

```
(Get-Cluster).WitnessDynamicWeight
```

If the command returns a value of 1, it indicates the witness currently has a vote. A value of 0 indicates the witness does not currently have a vote.

DynamicQuorum

The `DynamicQuorum` cluster property determines whether the Dynamic Quorum feature is enabled or disabled. (It is enabled by default in Windows Server 2012 R2.) To disable Dynamic Quorum, type the following:

```
(Get-Cluster).DynamicQuorum = 0
```

To re-enable Dynamic Quorum, type:

```
(Get-Cluster).DynamicQuorum = 1
```

DatabaseReadWriteMode

The `DatabaseReadWriteMode` cluster property is used to configure the Global Update Manager. The Global Update Manager is the software component within the Failover Clustering feature that is used to manage cluster database updates, such as the type that occurs when a node goes offline. Before Windows Server 2012 R2, the Global Update Manager always ensured that *all* nodes received the update before the update was committed to the cluster database. The Cluster Service therefore always read the database from the local node because it was known to be up-to-date.

Beginning in Windows Server 2012 R2, however, you can now alter the mode in which the Global Update Manager handles database updates and reads. The way that cluster database updates were handled in previous versions of Windows Server is now known as “All (Write) And Local (Read)” mode, and this mode is associated with a `DatabaseReadWriteMode` property value of 0. The All (Write) And Local (Read) mode is still the default in Windows Server 2012 R2 except when the clustered role is a VM.

The default mode when the clustered role is a VM is “Majority (Read And Write)”, which is associated with a `DatabaseReadWriteMode` property value of 1. In this latter mode, the

Global Update Manager commits an update to the cluster database when the update is received by a majority of nodes. Instead of reading automatically from the local node, the Cluster Service in Majority (Read And Write) mode checks a majority of the running nodes and reads the data with the latest timestamp.

To view the currently active Global Update Manager mode, type the following:

```
(Get-Cluster).DatabaseReadWriteMode
```

To set the Global Update Manager mode to All (Write) And Local (Read), type the following:

```
(Get-Cluster).DatabaseReadWriteMode = 0
```

To set the Global Update Manager mode to Majority (Read And Write), type the following:

```
(Get-Cluster).DatabaseReadWriteMode = 1
```

NodeWeight

Finally, in some cases, you might want to remove the quorum vote from a particular node or restore that vote after it has been removed. The ability to cast a vote for quorum is governed by the NodeWeight node property. A value of 1 (the default) indicates that the node has a right to vote for quorum. A value of 0 indicates that it does not.

To view the vote status of a node named Node 2, you would type the following:

```
(Get-ClusterNode Node2).NodeWeight
```

To remove the quorum vote from that node, type the following:

```
(Get-ClusterNode Node2).NodeWeight = 0
```

To give the quorum vote back to the same node, type the following:

```
(Get-ClusterNode Node2).NodeWeight = 1
```



EXAM TIP

Even though Network Load Balancing (NLB) hasn't changed significantly since Windows Server 2008 and isn't mentioned in this chapter, be sure to review the feature and its configurable options. For example, remember that in port rules for Network Load Balancing clusters, the Affinity setting determines how you want multiple connections from the same client handled by the NLB cluster. "Affinity: Single" redirects clients back to the same cluster host. "Affinity: Network" redirects clients from the local subnet to the cluster host. "Affinity: None" doesn't redirect multiple connections from the same client back to the same cluster host.



EXAM TIP

Make sure you know how to set the cluster properties described in the previous section.

Objective summary

- In Windows Server 2012 and Windows Server 2012 R2, you can create storage pools for failover clusters. Cluster storage pools are compatible only with SAS disks.
- CSVs are a new type of storage used only in some failover clusters. With CSVs, each node on a failover cluster creates its own virtual disk on the volume. The storage is then accessed through pathnames that are common to every node on the cluster.
- In Windows Server 2012 R2, you can under certain circumstances enable sharing on a VHDX file that is attached to a guest VM in Hyper-V Manager. The VHDX when shared in this way is then exposed to the guest operating system as a raw SAS disk. The SAS disk is eligible to be added to a cluster configured in the guest OS.
- With cluster-aware updating in Windows Server 2012, you can automate the process of updating Windows in a failover cluster.
- In Windows Server 2012 R2, you can create an Active Directory-detached failover cluster, which doesn't create a computer object for the cluster in Active Directory.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the "Answers" section at the end of the chapter.

1. You are designing storage for a failover cluster on two servers running Windows Server 2012 R2. You want to provision disks for the cluster that will enable you to create a storage pool for it. Which of the following sets of physical disks could you use to create a storage pools for the failover cluster?
 - A. Three individual disks in an iSCSI storage array without any RAID configuration
 - B. Four disks in an iSCSI storage array configured as a RAID 5
 - C. Three individual disks in a SAS storage array without any RAID configuration
 - D. Four disks in a SAS storage array configured as a RAID 5
2. You are an IT administrator for Adatum.com. The Adatum.com network includes 50 servers and 750 clients. Forty of the servers are virtualized. To provide storage for all servers, the Adatum.com network uses an iSCSI-based storage area network (SAN). You are designing storage for a new VM hosted in a failover cluster. Your priorities for the storage are to simplify management of SAN storage and to minimize downtime in case of node failure.

What should you do?

 - A. Use Server Manager to create a storage pool.
 - B. Keep VM storage on a CSV.
 - C. Provision volumes from an external SAS disk array instead of the iSCSI SAN.
 - D. Assign a mirrored volume to the cluster.

3. You have configured high availability for a cluster-aware application named ProseWareApp in a two-node failover cluster named Cluster1. The physical nodes in Cluster1 are named Node1 and Node2 and they are both running Hyper-V in Windows Server 2012 R2. Node1 is currently the active node for ProseWareApp.

You want to configure Cluster1 to perform critical Windows Updates with a minimum of administrative effort and a minimum of downtime for ProseWareApp users. What should you do?

- A. Drain the roles on Node1 and then start Windows Update on Node1.
- B. In Server Manager on Node1, right-click Cluster1 and select Update Cluster.
- C. Configure cluster-aware updating to add the CAU clustered role to Cluster1 with self-updating mode enabled.
- D. Configure Task Scheduler to run Windows Update daily on Node1 outside of business hours.

Objective 10.2: Manage failover clustering roles

This objective covers the configuration of roles in failover clusters. Within this area, there are three new features that you are likely to be tested on: the Scale-Out File Server role, role priority, and VM monitoring.

This section covers the following topics:

- Create a Scale-Out File Server (SoFS)
- Assign role priority
- Configure VM monitoring

Creating a Scale-Out File Server (SoFS)

Windows Server 2012 and Windows Server 2012 R2 let you configure two different types of file server roles for high availability: a file server for general use (which is the same option available in previous versions of Windows Server) and a new Scale-Out File Server For Application Data alternative. Both of these options are provided on the File Server Type page of the High Availability Wizard, as shown in Figure 10-12. Each of these clustered file server types is used for different purposes and they can both be hosted on the same node at the same time.

MORE INFO For more information about SoFS, visit <http://technet.microsoft.com/en-us/library/hh831349>.

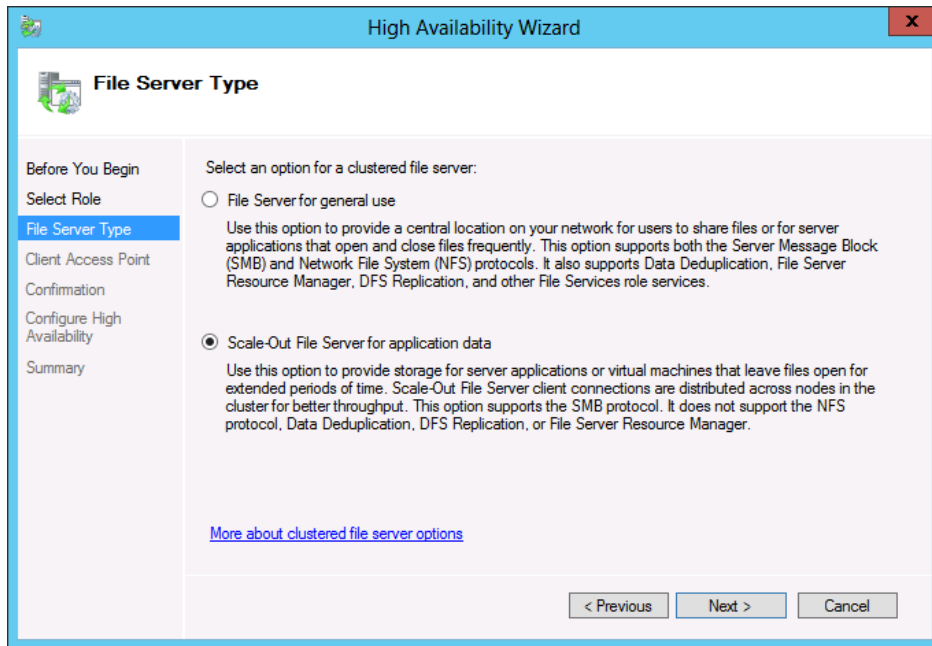


FIGURE 10-12 Selecting a Scale-Out File Server for the File Server role type

You will likely see a question or two on the 70-417 exam that tests basic knowledge about Scale-Out File Servers (SoFS). Here's what you need to remember:

- SoFS clusters are not designed for everyday user storage but for applications, such as SQL database applications, that store data on file shares and keep files open for extended periods of time.
- Client requests to connect to an SoFS cluster are distributed among all nodes in the cluster. For this reason, SoFS clusters can handle heavy workloads that increase proportionally to the number of nodes in the cluster.
- SoFS clusters use *only* CSVs for storage.
- SoFS clusters are not compatible with BranchCache, Data Deduplication, DFS Namespace servers, DFS Replication, or File Server Resource Manager.



EXAM TIP

Learn both the benefits and limitations of SoFS well. If a scenario requires a highly available file server for application data and all added nodes must remain online and able to respond to client requests, an SoFS is a good fit. But don't be tricked into selecting SoFS as the file server type for a new clustered file server just because the question states it will host application data. If the file server is also used with incompatible features (such as BranchCache, DFS, or File Server Resource Manager), or if no CSVs are available, you must choose File Server For General Use as the file server type.

Assign role startup priority

Unlike previous versions of Windows Server, Windows Server 2012 and Windows Server 2012 R2 let you assign one of four startup priorities to clustered roles: High, Medium, Low, or No Auto Start. Medium is the default priority. In the case of node failure, this priority setting determines the order in which roles are failed over and started on another node. A higher priority role both fails over and starts before the role of the next highest priority. If you assign the No Auto Start priority to a role, the role is failed over after the other roles but is not started on the new node. The purpose of startup priority is to ensure that the most critical roles have prioritized access to resources when they fail over to another node.

To change the startup priority of a role, right-click the role in Failover Cluster Manager, point to Change Startup Priority and select the desired priority, as shown in Figure 10-13.

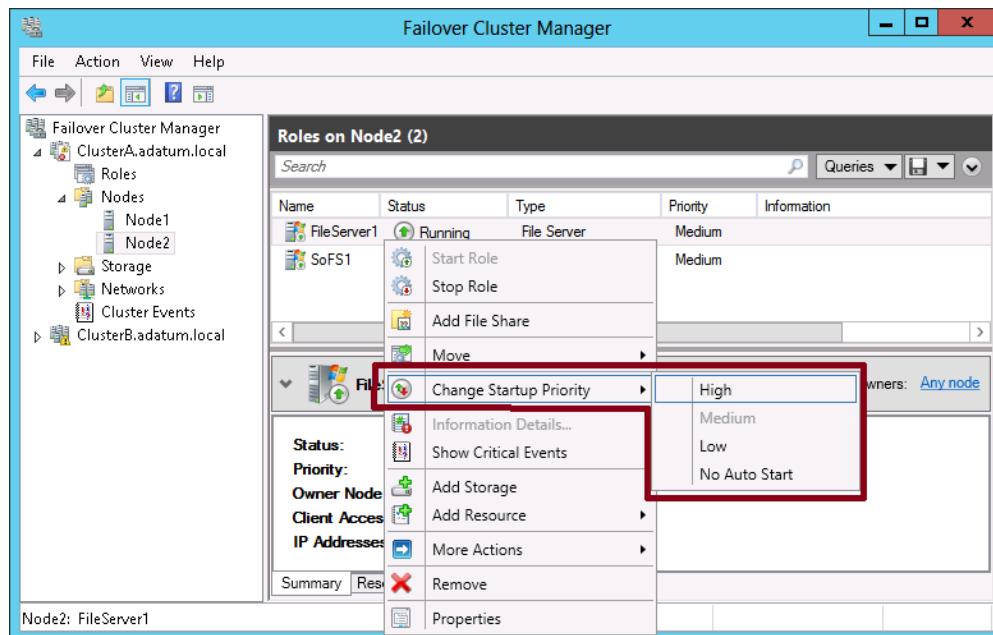


FIGURE 10-13 Setting the startup priority of a role

Role startup priority is a fairly easy feature to understand that is also likely to appear on the 70-417 exam. Be sure you remember the No Auto Start priority especially because that's the only priority setting with a meaning that isn't made obvious by its name.



EXAM TIP

You need to understand the difference between startup priority settings and preferred owner settings. Startup priority settings determine the order in which roles should be failed over and started after node failure. Preferred owner settings determine which *node*, if available, should handle the client requests for a role both before and after a node failure.

Virtual machine application monitoring

Windows Server 2012 and Windows Server 2012 R2 introduce the ability for a Hyper-V host to monitor the health of chosen services running on a clustered VM. If the Hyper-V host determines that a monitored service in a guest VM is in a critical state after normal attempts to restart the service within the guest OS have failed, the host is able to trigger a recovery. The Cluster service first attempts to recover the VM by restarting it gracefully. Then, if the monitored service is still in a critical state after the VM has restarted, the Cluster service fails the VM over to another node.

To monitor VM services with the VM Monitoring feature in Windows Server 2012 and Windows Server 2012 R2, the following requirements must be met:

- Both the Hyper-V host and its guest VM must be running Windows Server 2012 or later.
- The guest VM must belong to a domain that trusts the host's domain.
- The Failover Clustering feature must be installed on the Hyper-V host. The guest VM must also be configured as a role in a failover cluster on the Hyper-V host.
- The administrator connecting to the guest through Failover Cluster Manager must be a member of the local administrators group on that guest.
- All firewall rules in the Virtual Machine Monitoring group must be enabled on the guest, as shown in Figure 10-14.

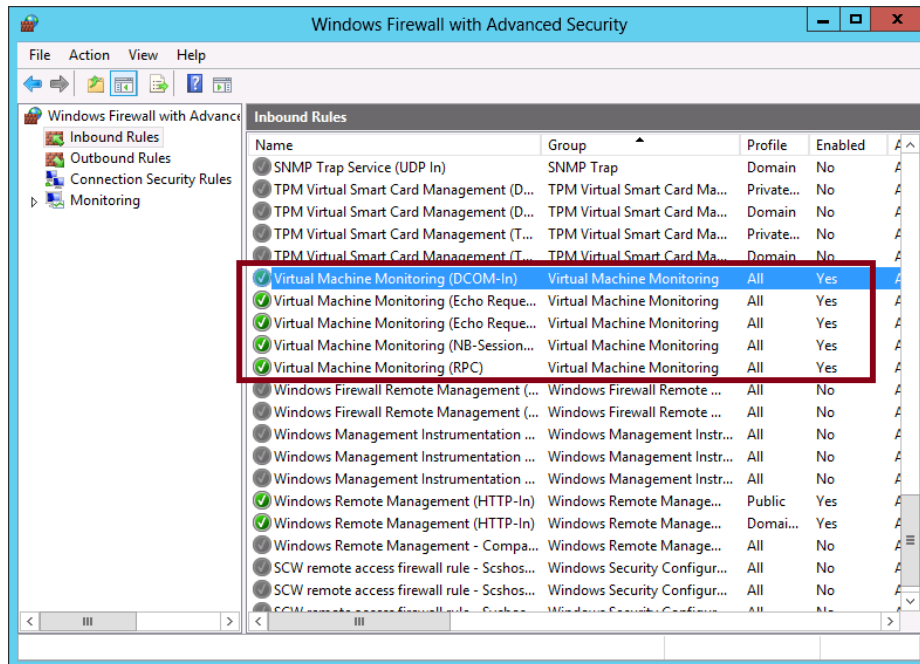


FIGURE 10-14 Enabling firewall rules for VM monitoring

To configure VM monitoring, right-click the VM in Failover Cluster Manager, point to More Actions, and then select Configure Monitoring, as shown in Figure 10-15.

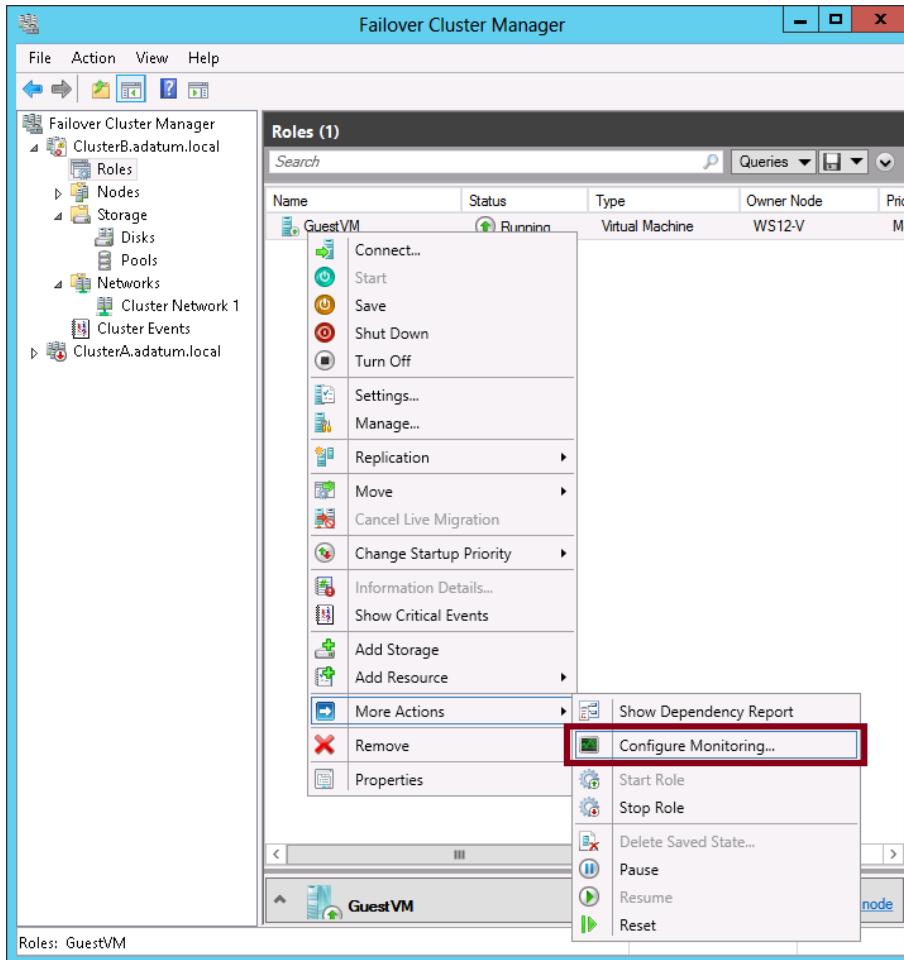


FIGURE 10-15 Configuring the monitoring of a VM application

In the Select Services dialog box that opens, select the services that you want to monitor, as shown in Figure 10-16.

When you configure a service to be monitored as shown in Figure 10-16, the failure of the service is still primarily handled by the guest OS by default, not by the Cluster service on the Hyper-V host. The general response to a service failure is determined by the service properties accessible through the Services console within the guest. As you might remember, the Recovery tab in a service's properties allows you to specify the local computer's response if that service fails. You can specify a First Failure response, a Second Failure response, and the Subsequent Failures response. The default settings are shown in Figure 10-17.

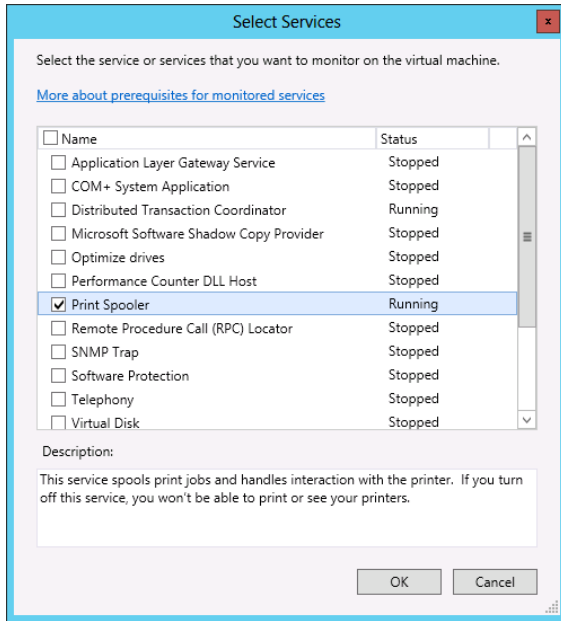


FIGURE 10-16 Selecting services to be monitored in a VM

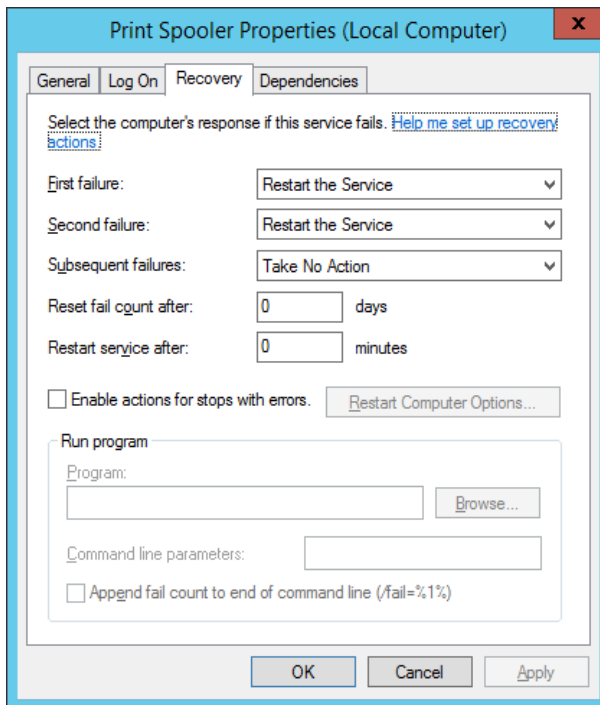


FIGURE 10-17 Default recovery properties for a service

As shown in Figure 10-17, the default recovery properties for a service are configured with the Restart The Service setting for First Failure and Second Failure. These attempts to restart the service are handled by the local operating system and not the Cluster service on the Hyper-V host.

The default setting for Subsequent Failures, also shown in Figure 10-17, is Take No Action. The Take No Action setting is where VM application monitoring comes in. If the service has been configured for monitoring, the Cluster service on the Hyper-V host will move in to take over service recovery at whichever stage the Take No Action setting is configured.

In some circumstances, you might want to redirect the Cluster service recovery to a third-party application that allows you more control over the recovery process. In this case, you can disable the default behavior to restart and fail over the VM. You can achieve this first by opening the resource properties of the guest VM in Failover Cluster Manager, as shown in Figure 10-18.

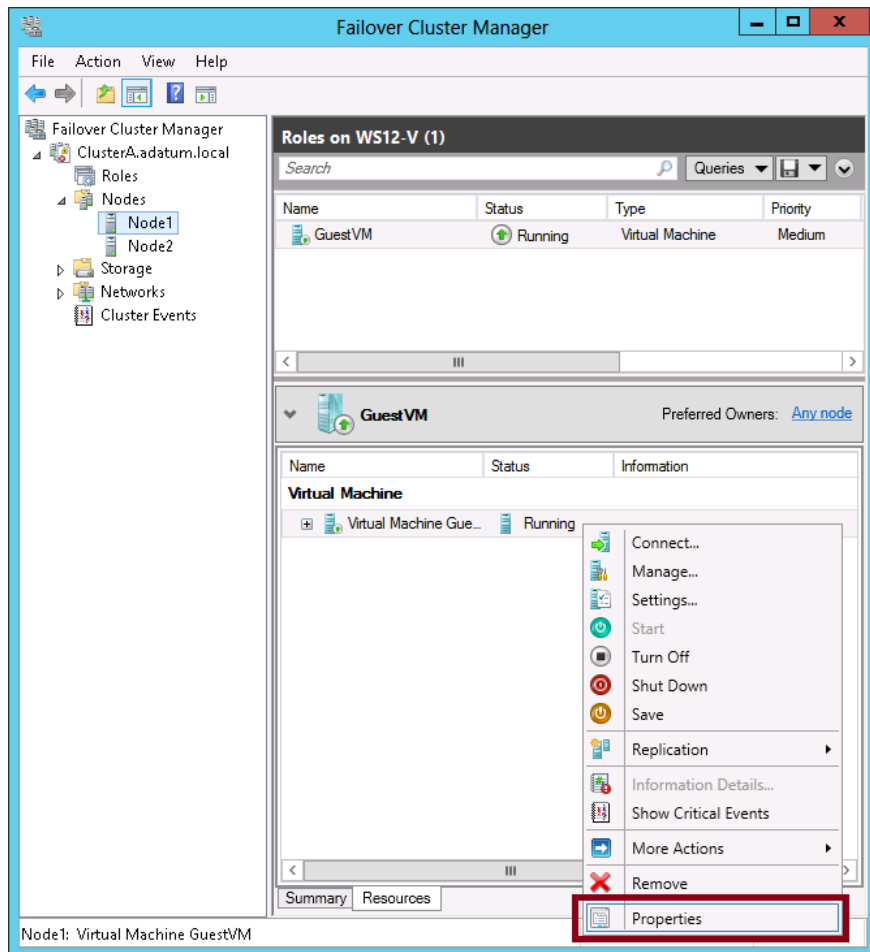


FIGURE 10-18 Modifying properties of a clustered VM



EXAM TIP

Remember the significance of the Take No Action setting. For example, if you want a clustered VM to fail over to another node when a monitored service in that VM fails the *first* time (as opposed to the third time, which is the default), you need to configure the recovery settings for the service so that First Failure is set to Take No Action.

Then on the Settings tab of the Properties dialog box shown in Figure 10-19, clear the Enable Automatic Recovery For Application Health Monitoring check box. The Cluster service will still log an error when a monitored service is in a critical state, but it will no longer attempt to restart or fail over the VM.

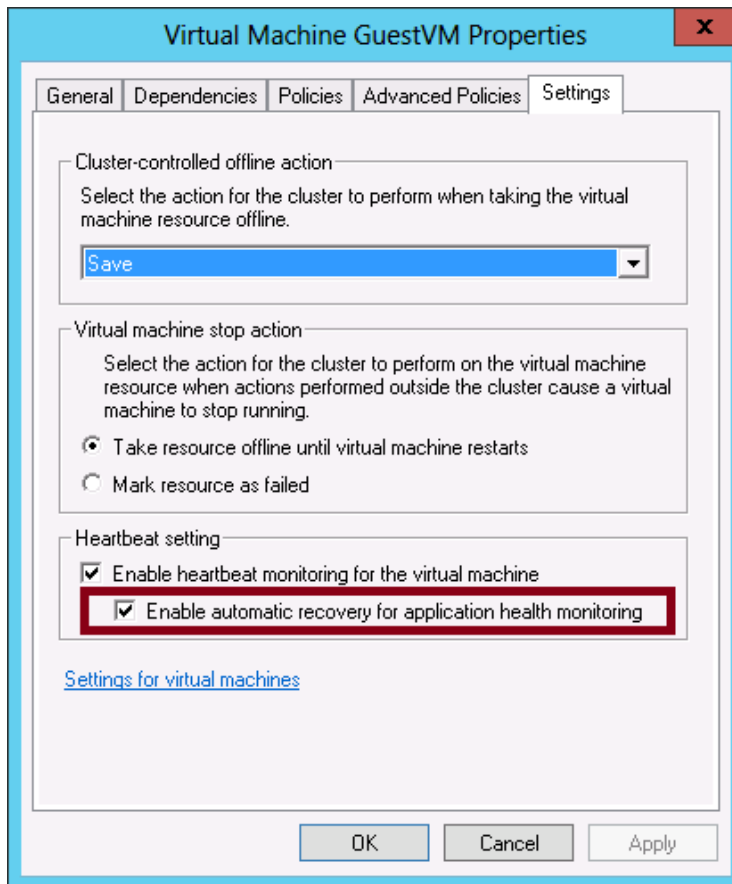


FIGURE 10-19 The setting to enable automatic recovery for a monitored VM application



EXAM TIP

To migrate all clustered roles from a cluster running on Windows Server 2008 R2 to a new cluster running on Windows Server 2012 or Windows Server 2012 R2, right-click the cluster in Failover Cluster Manager on the source cluster and then select **Migrate Roles**. This step will start the **Migrate A Cluster Wizard** and guide you through the migration process.

Objective summary

- A Scale-Out File Server (SoFS) is a new type of role for which you can configure high availability in a failover cluster. An SoFS can be used to ensure that an application that connects to a file share doesn't generate errors during failover. In addition, an SoFS works on many live nodes at a time, so every additional node you add enables the cluster to handle more requests. An SoFS is not well-suited for use with file storage for users.
- With startup priority, you can determine the order in which roles should be failed over from one node to the next in case of node failure.
- Windows Server 2012 and Windows Server 2012 R2 let you monitor the health of an application running in a guest VM. When a monitored application reaches a critical state, the host computer can trigger the VM to restart. If the application remains in a critical state after the system restart, the host computer will trigger failover to another node.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the "Answers" section at the end of the chapter.

1. You work as a network administrator for Adatum.com. The Adatum.com network includes 25 servers running Windows Server 2012 R2 and 400 clients running Windows 8.1. You want to create a failover cluster to support a file share used by a resource-intensive application. Your priorities for the failover cluster are to prevent file handling errors in the event of failover and to maintain high performance as the usage of the application grows. Which role and storage type should you configure for the failover cluster? (Choose two. Each answer represents part of the solution.)
 - A. Configure as the role for the failover cluster a file server for general use.
 - B. Configure as the role for the failover cluster an SoFS.
 - C. Store the share on an NTFS volume provisioned from shared storage. Do not add the volume to CSVs.
 - D. Store the share on a CSV.

2. You work as a network administrator for Fourth Coffee, Inc. The Fourthcoffee.com network spans offices in five cities in North America. All servers in the network are running Windows Server 2012 R2 and all clients are running Windows 8.1.
- You want to create a failover cluster to support a new file share that will be used by members of the marketing team in all branch offices. Your requirements for the failover cluster and the file share in general are to minimize downtime if a node fails, to minimize storage space needed for the share, to reduce or eliminate the possibility of file conflicts, and to minimize the amount of data transferred over wide area network (WAN) links.
- How should you configure the failover cluster and file server? (Choose all that apply.)
- A. Configure as the role for the failover cluster a file server for general use.
 - B. Configure as the role for the failover cluster an SoFS.
 - C. Enable Data Deduplication on the file share.
 - D. Enable BranchCache on the file share.
3. You want to create a two-node failover cluster to provide high availability for a virtual machine. The VM will host an important line-of-business (LOB) application used often by members of your organization throughout the day. You want to configure VM monitoring of the application so that the VM will restart if the application is found to be in a critical state and fail over to the other node if the application still is in a critical state after the system restart.
- Which of the following is *not* a requirement of meeting this goal?
- A. The host Hyper-V server needs to be running Windows Server 2012 or later.
 - B. The guest VM needs to be running Windows Server 2012 or later.
 - C. The host and the guest need to be members of the same domain.
 - D. The guest VM needs to have enabled the firewall rules in the Virtual Machine Monitoring group.

Objective 10.3: Manage virtual machine (VM) movement

Windows Server 2012 and Windows Server 2012 R2 both expand and improve on VM migration features that were introduced in Windows Server 2008 R2. The two biggest improvements are the addition of live migration in a nonclustered environment and storage migration.

This section covers the following topics:

- Configuring and performing live migration in a failover cluster
- Configuring and performing live migration outside of a failover cluster
- Performing storage migration
- VM network health protection in Windows Server 2012 R2

Live migration

Live migration is a feature that first appeared in Windows Server 2008 R2. Live migration lets you move a running VM from one Hyper-V host to another without any downtime. Originally, this feature was available only for VMs hosted in failover clusters, but in Windows Server 2012 and Windows Server 2012 R2, you can now perform live migration of a VM outside of a clustered environment. However, the process of performing live migration is different inside and outside of clusters, and each of these live migration types has slightly different requirements.

Live migration requires a few configuration steps that you need to understand for the exam. To start configuring this feature, open Hyper-V Settings for each Hyper-V host, as shown in Figure 10-20.

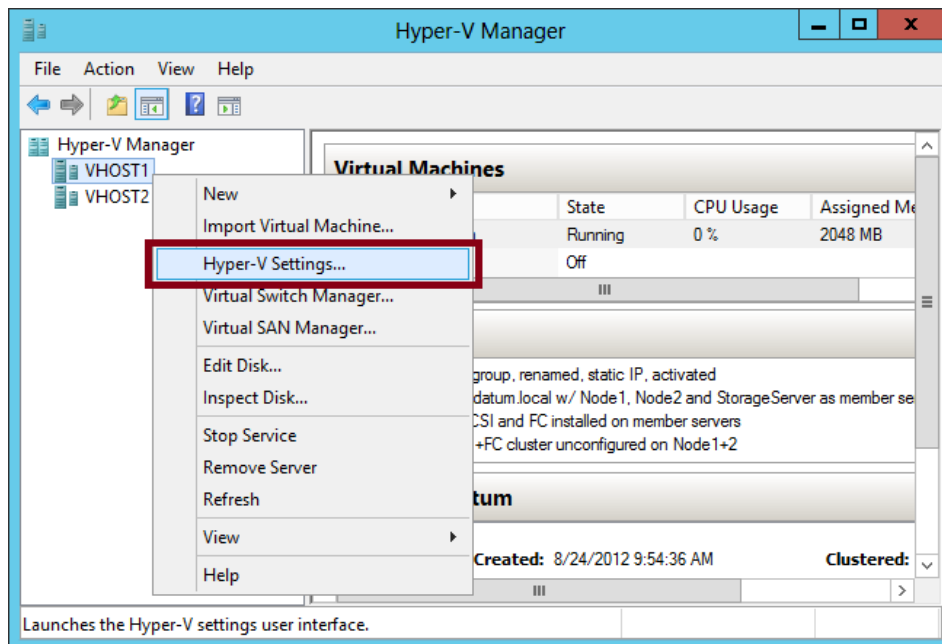


FIGURE 10-20 Configuring Hyper-V settings

In the Hyper-V Settings dialog box that opens, click Live Migrations on the menu on the left. The associated live migration settings are shown in Figure 10-21.




FIGURE 10-21 Live migration settings

Live migrations are not enabled by default. To enable this feature, perform the following four configuration steps in this dialog box:

1. Select the Enable Incoming And Outgoing Live Migrations check box for both the source and destination hosts.
2. For live migrations outside of clustered environments, you need to choose an authentication protocol on both host servers: either Credential Security Support Provider (CredSSP) or Kerberos.
 - **CredSSP** The advantage of this choice is that it requires no configuration. The limitation of choosing CredSSP as the authentication protocol, however, is that you need to be logged on to the source computer when you perform a live migration. You can't perform live migrations through Hyper-V Manager on a remote computer.
 - **Kerberos** The advantage of choosing Kerberos as the authentication protocol is that you don't need to be logged on to a source computer to perform the live

migration. The disadvantage of using Kerberos as the authentication protocol for live migrations is that it requires configuration. Specifically, aside from selecting Kerberos in Hyper-V Settings, you need to adjust the properties of the source and destination computer accounts in Active Directory Users and Computers, on the Delegation tab. On each computer account, select the option to trust this computer for delegation to specified services only, and then add the following two services from the other computer: CIFS and Microsoft Virtual System Migration Service. The actual configuration required is shown in Figure 10-22. Note that this configuration step is also known as configuring *constrained delegation*. Expect to see a question about configuring constrained delegation on the 70-417 exam.

3. Set a value for the maximum number of simultaneous live migrations you want to allow on the network. This is a new feature of Windows Server 2012 and Windows Server 2012 R2. In Windows Server 2008 R2, you were limited to one live migration at a time. (In the real world, you can estimate 500 Mbps of network bandwidth required per individual live migration. In a Gigabit Ethernet network, you can safely leave the default value of 2.)
4. Add a list of subnets or individual IP addresses from which you want to allow live migrations. Live migration does not provide data encryption of VMs and storage as they are moved across the network, so security is an important consideration. Do not leave the default selection to use any available network for live migration unless you are in a testing environment.

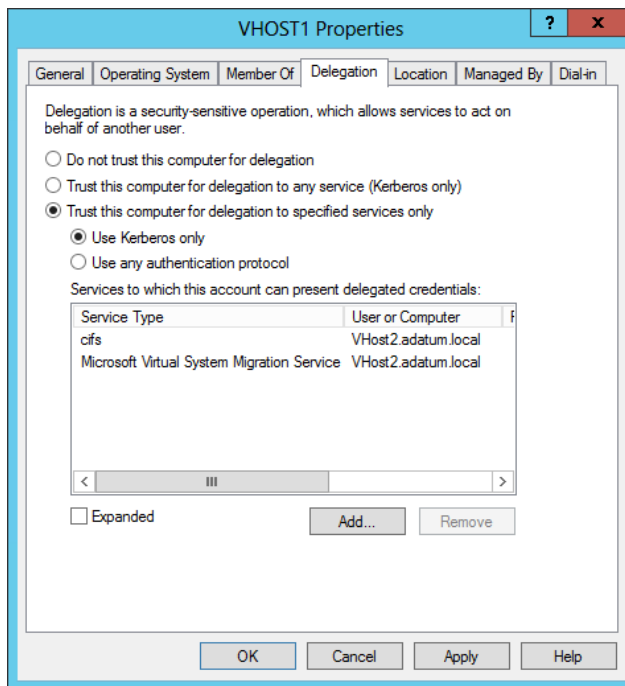


FIGURE 10-22 Configuring constrained delegation

Live migration in a failover cluster

Although CSVs are not required for VM storage when you perform live migration in a failover cluster, CSVs are nonetheless highly recommended. If the VM is not already stored in a CSV, you should move it there to prepare for clustered live migration.

MOVING VM STORAGE TO A CSV

To move VM storage to a CSV, right-click the VM in Failover Cluster Manager, point to Move, and then click Virtual Machine Storage, as shown in Figure 10-23.

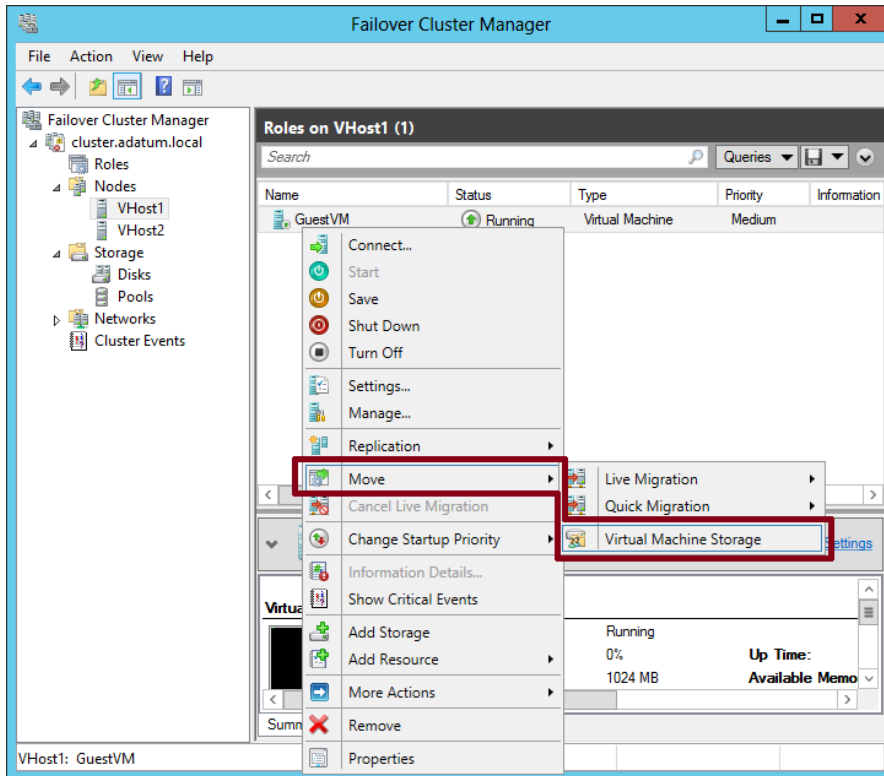


FIGURE 10-23 Moving virtual machine storage

Then, in the Move Virtual Machine Storage dialog box that opens, shown in Figure 10-24, select the VM in the top pane and then drag it to a CSV folder in the bottom left pane. Click Start to begin the copy operation.

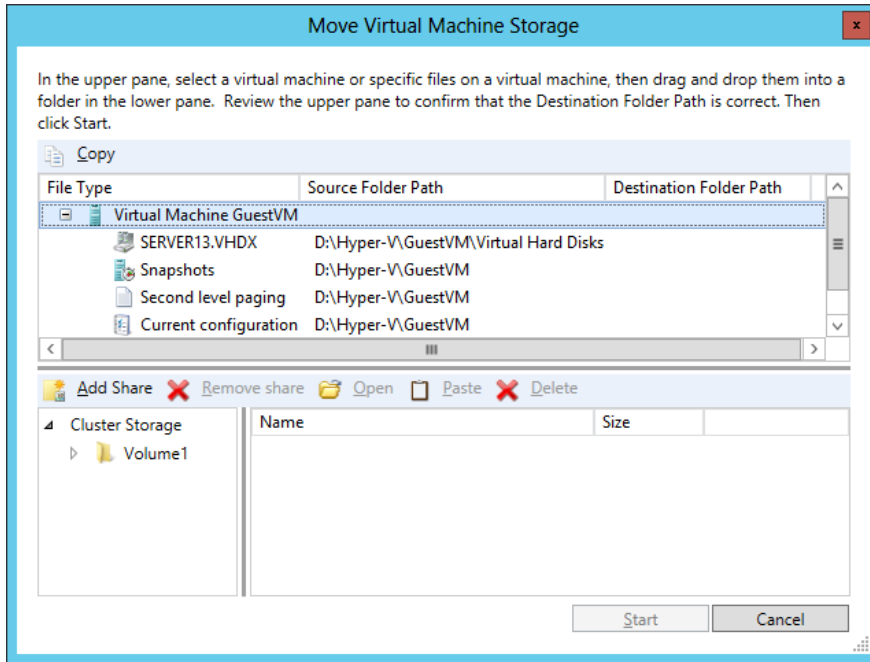


FIGURE 10-24 Moving VM storage to a CSV

PERFORMING LIVE MIGRATION

After the transfer is complete, you can perform a live migration as long as the Hyper-V environments are the same on the source and destination nodes, including the names of virtual switches in both locations. To perform the live migration, in Failover Cluster Manager, right-click the clustered VM, point to Move, point to Live Migration, and then click Select Node from the shortcut menu, as shown in Figure 10-25. (Note also the Best Possible Node option, which selects a destination node for you.)

In the Move Virtual Machine Dialog box that opens, shown in Figure 10-26, select the destination node in the failover cluster to which you want to transfer the running VM and then click OK to start the process.

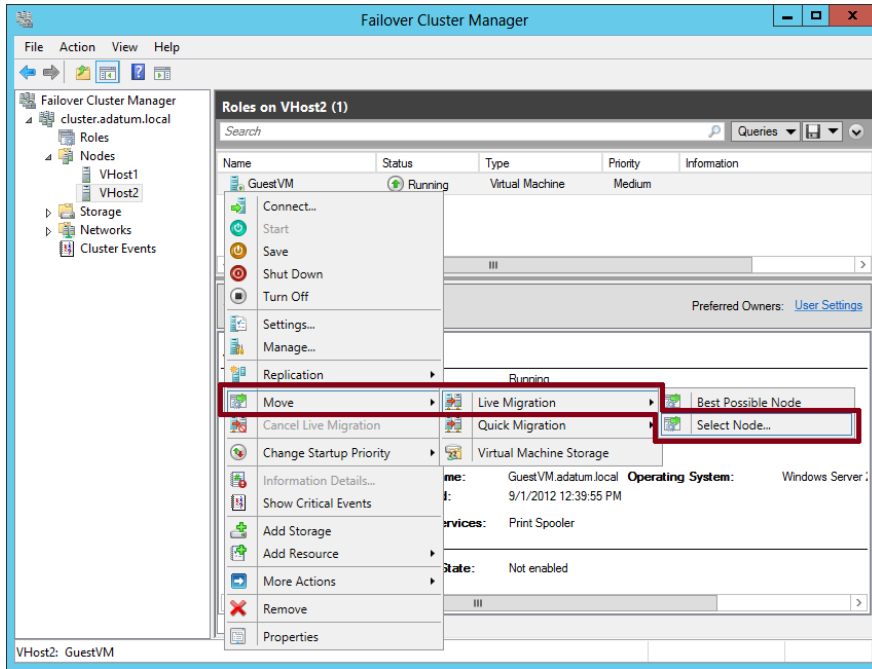


FIGURE 10-25 Performing a live migration in a failover cluster

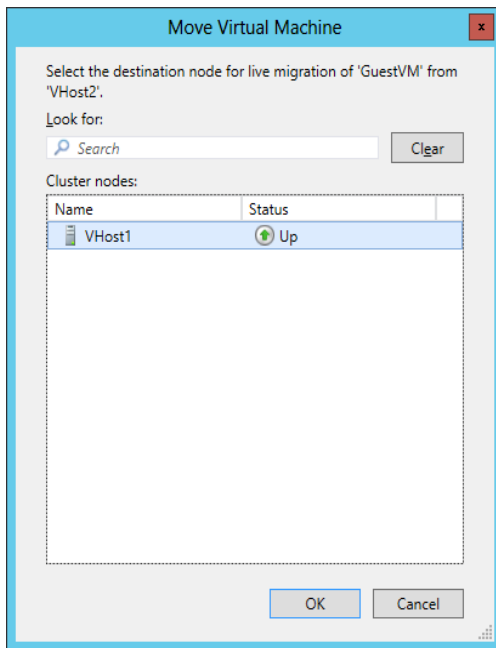


FIGURE 10-26 Selecting a destination node for live migration

You can keep track of the migration status in Failover Cluster Manager, as shown in Figure 10-27.




FIGURE 10-27 Viewing live migration status.

During and after migration, the VM continues to operate without interruption.

NOTE First introduced in Windows Server 2008, Quick Migration is still an available option when you choose to move a clustered VM in Windows Server 2012 or Windows Server 2012 R2 from one node to another. Quick Migration saves the VM state and resumes the machine on the destination node. The advantage of Quick Migration is that it is a faster process from start to finish for the VM you are migrating, and it requires less network bandwidth. The disadvantage of Quick Migration is that the VM is briefly brought offline during the migration process. If minimizing downtime is not a priority and you want to transfer a VM as quickly as possible, then Quick Migration is the best option.

Live migration outside of a clustered environment

Nonclustered live migration is a new feature in Windows Server 2012 and Windows Server 2012 R2 in which you can move a running VM from one Hyper-V host to another, with no downtime, outside of a clustered environment. The feature does require that the source and destination Hyper-V hosts belong to domains that trust each other. However, it doesn't require SAN storage or a clustered environment. It's also worth noting that a disadvantage of nonclustered live migration, compared to clustered live migration, is that the process takes much longer because all files are copied from the source to the destination host. (An exception to this rule is if the VM and its storage are kept on a file share and do not need to be copied from one host to the other during the migration process.)

Once you have configured live migration settings in Hyper-V Manager on the source and destination computers, you can perform the live migration. It's a simple procedure. In Hyper-V Manager, right-click the running VM you want to live migrate and select Move from the shortcut menu, as shown in Figure 10-28.




FIGURE 10-28 Initiating a live migration outside of a clustered environment

In the wizard that opens, select the Move The Virtual Machine option, as shown in Figure 10-29.

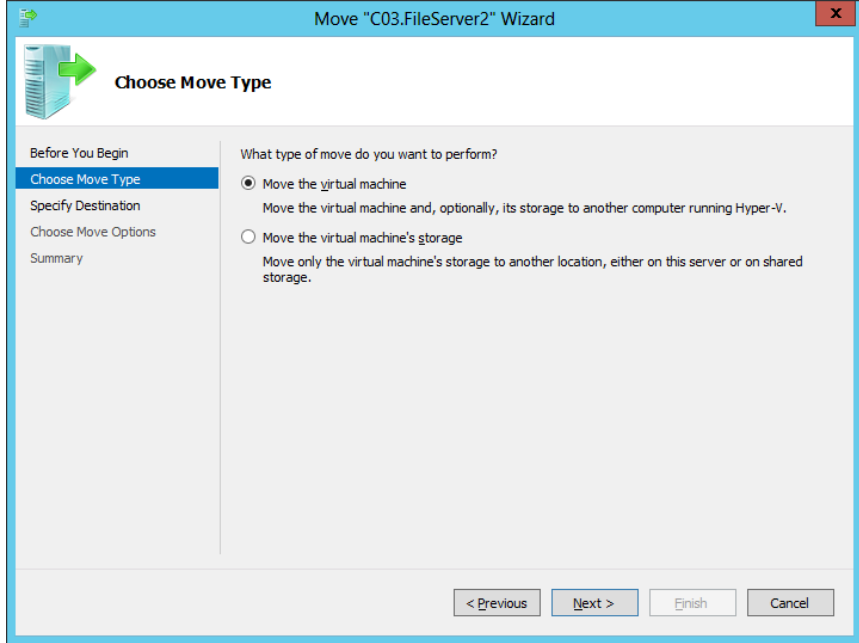


FIGURE 10-29 Live migrating a VM in a nonclustered environment

Then, specify the destination server, as shown in Figure 10-30.

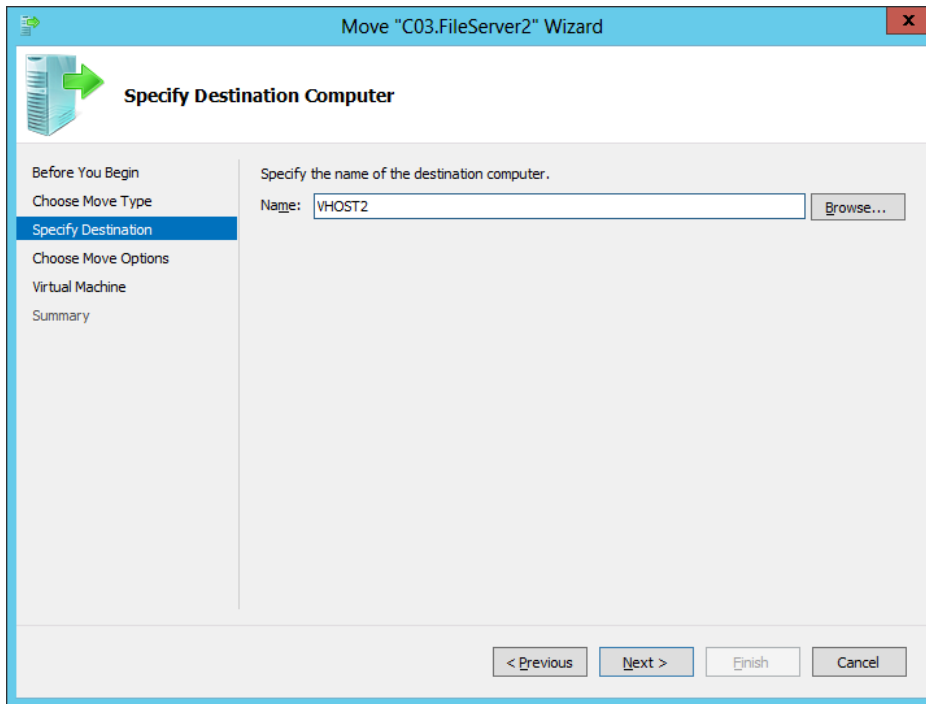


FIGURE 10-30 Choosing a destination host for live migration in a nonclustered environment

You have three options for how you want to move the VM's items when you perform the live migration, as shown in Figure 10-31. First, you can move all of the VM's files and storage to a single folder on the destination computer. Next, you can choose to move different items to different folders in a particular way that you specify. Finally, you can migrate just the VM while leaving the storage in place. Note that this option requires the VM storage to reside on shared storage such as an iSCSI target, a fact that could easily serve as the basis for an incorrect answer choice on a test question.

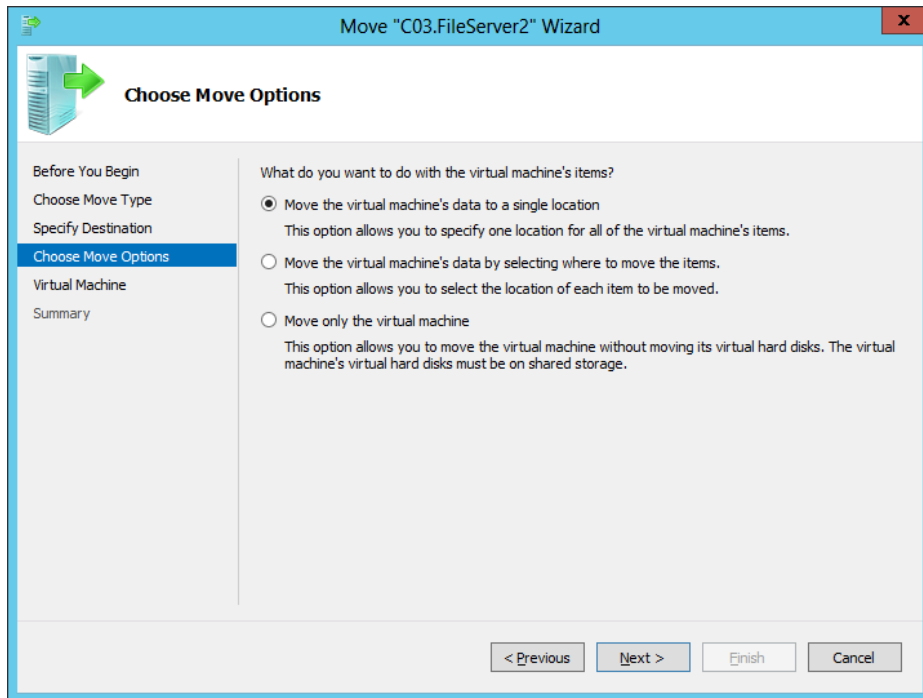


FIGURE 10-31 Moving a VM to a single destination folder

Processor compatibility

One potential problem that can arise when you perform a live or quick migration is that the processor on the destination Hyper-V host supports different features than does the processor on the source host. In this case, you might receive the error message shown in Figure 10-32, and the migration fails. (A failed migration does not negatively affect the running VM. The result of the failure is simply that the VM is left running on the source computer.)

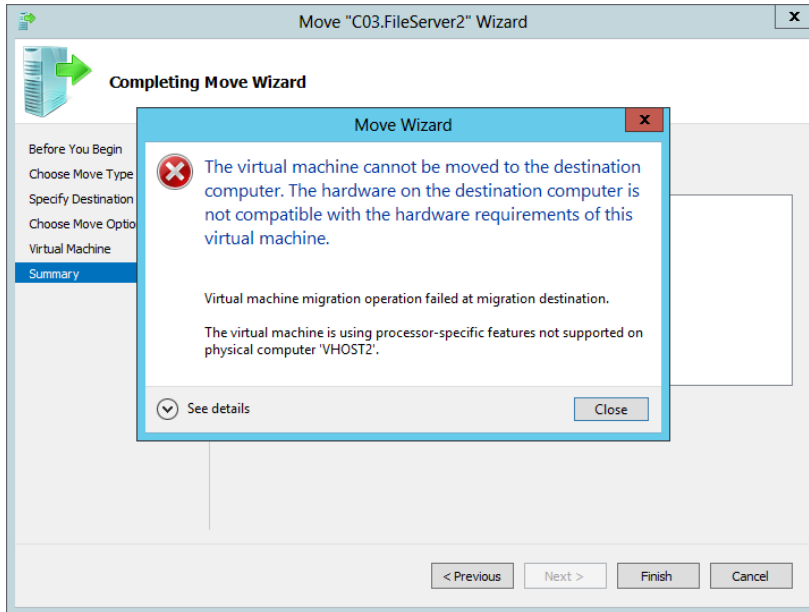


FIGURE 10-32 An error indicating processor feature incompatibility

Neither live migration nor quick migration is supported between hosts with processors from different manufacturers. However, if the processors on the source and destination computers are from the same manufacturer and are found to support incompatible features, you have the option of limiting the virtual processor features on the VM as a way to maximize compatibility and improve the chances that the migration will succeed. (Again: This setting does not provide compatibility between different processor manufacturers.)

To enable processor compatibility, expand the Processor settings in the VM's settings dialog box and select the Compatibility node, as shown in Figure 10-33. Then, select the Migrate To A Physical Computer With A Different Processor Version check box. Alternatively, you can run the following command at a Windows PowerShell prompt:

```
Set-VMProcessor VMname -CompatibilityForMigrationEnabled $true
```

If you see a question about live migration in which you receive an error indicating trouble related to processor-specific features, you now know how to handle it: Just enable the processor compatibility setting.



EXAM TIP

You should expect to see questions about VM migration in which the fact that the source and destination hosts have different processor manufacturers is located somewhere in a table or list. When the processor manufacturers are different, your best option for migration is to manually export it from the source machine and import it on the destination machine. You can't use live migration or quick migration. In Windows PowerShell, you perform these manual export and import operations with the Export-VM and Import-VM cmdlets.




FIGURE 10-33 Enabling processor compatibility for migration

Virtual switch name matching

Another common problem that can occur when you attempt to perform a live migration is that the destination Hyper-V environment doesn't provide virtual switches with names that exactly match those in the Hyper-V environment on the source computer. This problem is detected as you complete the Move Wizard. For each snapshot of the source VM that defines a virtual switch without an exact equivalent on the destination Hyper-V host, you are given an opportunity to choose another virtual switch on the destination that the VM should use in place of the original. This step is shown in Figure 10-34.




FIGURE 10-34 Matching a virtual switch on a destination host for live migration

After you make the required substitutions, the wizard begins the live migration when you click Finish on the final page, as shown in Figure 10-35.




FIGURE 10-35 A live migration in progress in a nonclustered environment

MORE INFO For more information about live migration in Windows Server 2012 and Windows Server 2012 R2, visit <http://technet.microsoft.com/en-us/library/jj134199.aspx>.

Storage migration

Another useful new feature in Windows Server 2012 and Windows Server 2012 R2 is the live migration of VM storage. With this option, you can move the data associated with a VM from one volume to another while the VM remains running. This option is useful if storage space is scarce on one volume or storage array and is more plentiful on another source of storage. An important advantage of storage-only live migration to remember for the exam is that unlike live migration, it can be performed in a workgroup environment because the source and destination servers are the same.

To perform storage migration, use the Move option to open the Move Wizard, as you would do to begin the process of live-migrating the VM. Then, on the Choose Move Type page of the wizard, select Move The Virtual Machine's Storage, as shown in Figure 10-36.

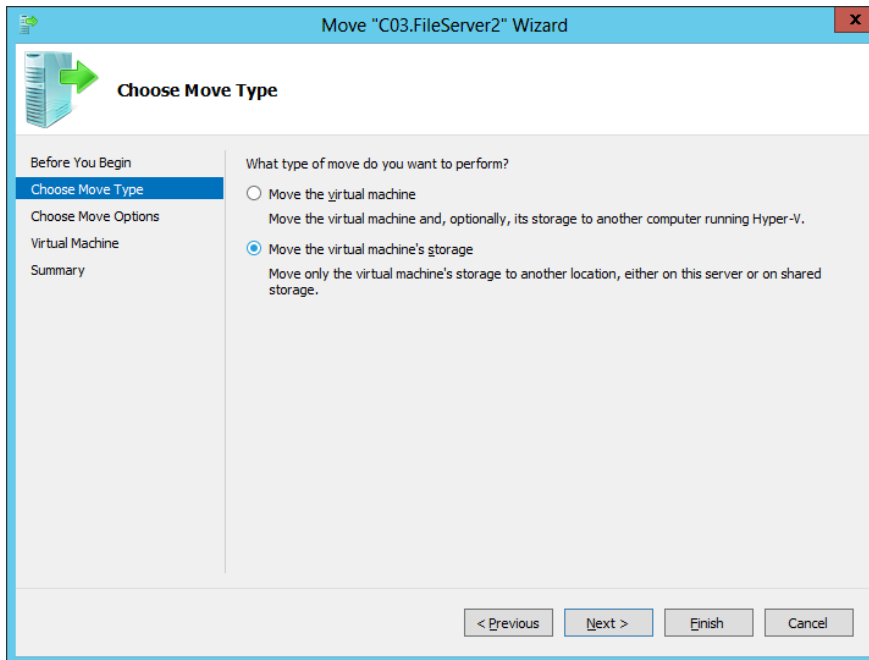


FIGURE 10-36 Choosing a migration type

You have three options for how to migrate the storage of the VM, as shown in Figure 10-37. The first option is to move all storage to a single folder on the destination volume.

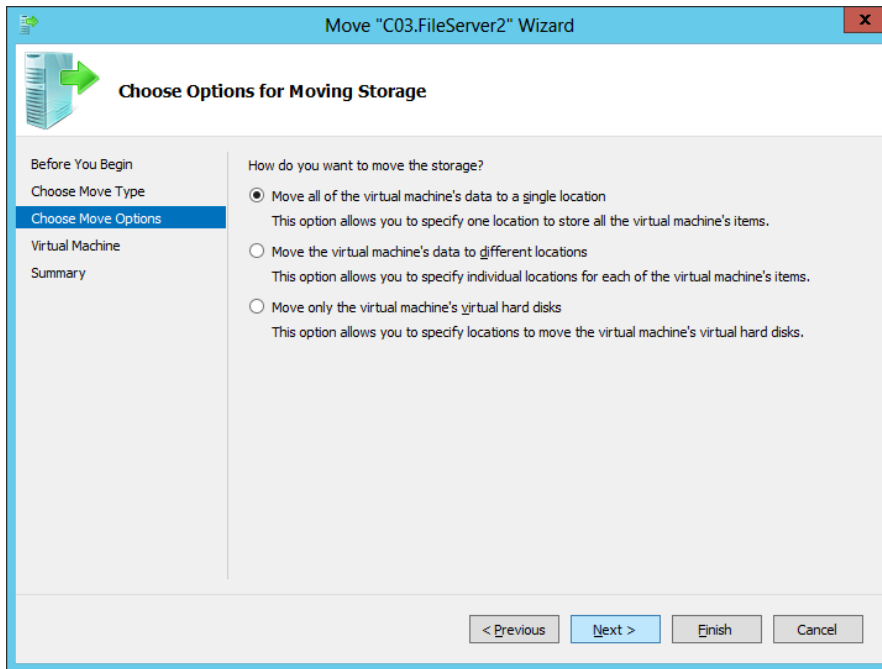


FIGURE 10-37 Moving storage to a single folder

The second option allows you to select which particular storage items you want to migrate, such as snapshot data or the smart paging folder, as shown in Figure 10-38. The third and final option allows you to specify particular VHDs to migrate only.

For the exam, what's most important to remember about storage migration is that this feature provides an option that is often the best way to solve a problem. If space runs out for a running VM, it's not necessarily a good idea to migrate that VM to another server. No other server might be available, for example, and you might want to spare your organization the unnecessary expense of buying a new one. In this case, it's often more prudent simply to attach a new disk array to the server and move the VM storage to this newly available space.

MORE INFO For a good review of Live Migration and Storage Migration, see the TechEd Australia session on the topic at: <http://channel9.msdn.com/Events/TechEd/Australia/2012/VIR314>.




FIGURE 10-38 Selecting VM items to migrate to a new volume

VM network health protection in Windows Server 2012 R2

In Windows Server 2012 R2, there's a new option named Protected Network in the advanced Network Adapter settings of a VM in Hyper-V Manager, as shown in Figure 10-39. You use this setting to protect a highly available VM from the failure of the associated network connection.

Before Windows Server 2012 R2, if a clustered VM were to lose network connectivity in a way that didn't affect its heartbeat connection, the VM would remain on the same physical node and not trigger a failover. Such a network disruption could prevent clients from connecting to the VM.

NOTE Failover cluster nodes exchange heartbeats once per second by default. The number of heartbeats that can be missed before a failover is triggered is called the heartbeat threshold.

When you select the Protected Network option on the settings of a virtual network adapter, the physical node monitors that network for disruptions. If that network connection is broken, the VM is automatically live-migrated to another available node.

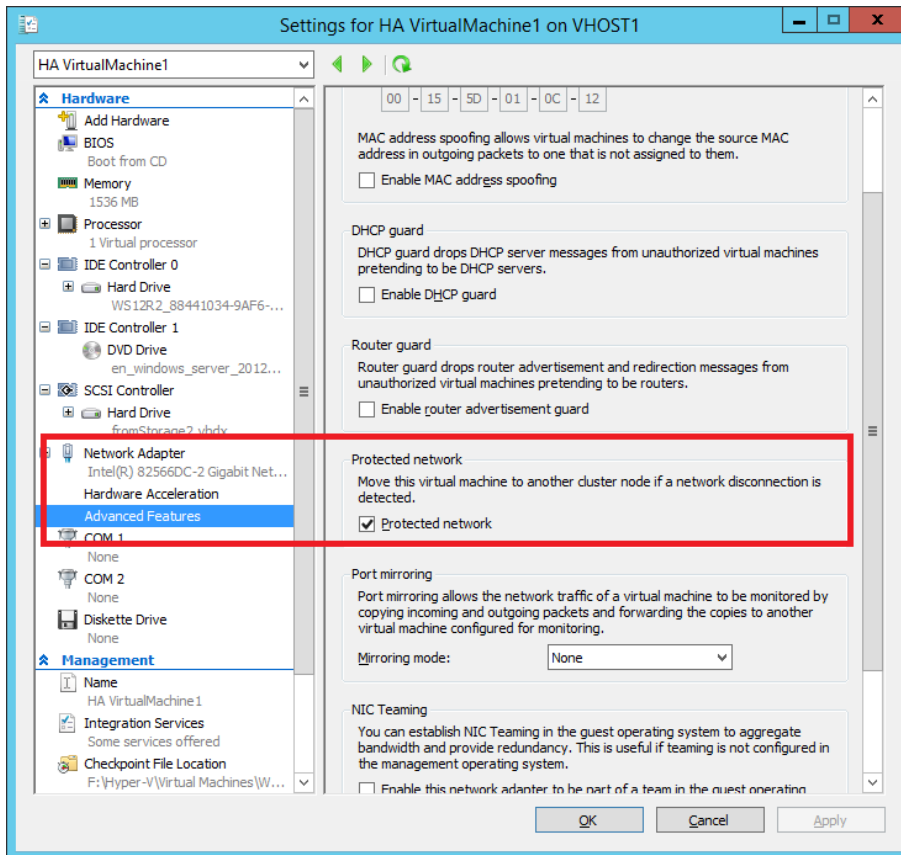


FIGURE 10-39 Configuring a highly available VM for network protection

Objective summary

- Live migration is a feature in which a running VM is transferred from one host computer to another without any downtime. In Windows Server 2012 and Windows Server 2012 R2, live migration can be performed inside or outside a failover cluster, but live migration within a failover cluster is much faster.
- When configuring servers for live migration outside of a failover cluster, you must choose an authentication protocol, CredSSP or Kerberos. CredSSP needs no configuration, but it requires you to trigger the live migration while logged in to the source host. Kerberos allows you to trigger the live migration from a remote host, but it requires you to configure constrained delegation for the source and destination hosts.
- Windows Server 2012 and Windows Server 2012 R2 introduce storage migration for VMs. With storage migration, you can move all of the storage associated with a running VM from one disk to another without any downtime.
- In Windows Server 2012 R2, you can enable network protection on a network adapter in a VM's settings. If the VM is clustered and a disconnection is then detected on the selected network, a live migration of the VM is triggered to another cluster node.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the "Answers" section at the end of the chapter.

1. You are a network administrator for Contoso.com. You have recently upgraded all of your servers to Windows Server 2012 R2. Your manager has indicated that he wants to start testing the live migration feature in a nonclustered environment so that you can eventually take advantage of this functionality in production.

You create a small test network consisting of two Hyper-V servers running Windows Server 2012 named Host1 and Host2. The hardware and software settings on these two physical servers exactly match those of two physical servers in your production network. Host1 is currently hosting a guest VM named VM1.

You enable live migration on both servers and configure CredSSP as the authentication protocol. You then log on locally to Host1 and initiate a live migration of VM1 from Host1 to Host2. You receive an error message indicating that the VM is using processor-specific features not supported on the destination physical computer.

You want to perform a live migration successfully in your test network so that you will know what is required to use this feature successfully in production. What should you do?

- A. Configure constrained delegation for Host1 and Host2.
- B. Disable VM monitoring on VM1.
- C. Configure Kerberos as the authentication protocol on Host1 and Host2.
- D. On Host1, run the following command:

```
Set-VMProcessor VM1 -CompatibilityForMigrationEnabled $true
```

- 2.** You are a network administrator for Adatum.com. You have recently upgraded all of your servers to Windows Server 2012 R2. Your manager has indicated that she wants to start testing the live migration feature so that you can eventually take advantage of this functionality in production.

You create a small test network consisting of two Hyper-V servers running Windows Server 2012 R2 named VHost1 and VHost2. The hardware and software settings on these two physical servers exactly match those of two physical servers in your production network. VHost2 is currently hosting a guest VM named VM2.

You enable live migration on both servers and configure Kerberos as the authentication protocol. You then log on locally to Host1 and initiate a live migration of VM1 from VHost2 to VHost1. The live migration fails and you receive an error indicating “No credentials are available in the security package.”

You want to perform a live migration successfully in your test network so that you will know what is required to use this feature successfully in production. You also want to initiate live migrations when you are not logged on to the source host server. What should you do next?

- A.** Configure constrained delegation for VHost1 and VHost2.
- B.** Disable VM monitoring on VM2.
- C.** Configure CredSSP as the authentication protocol on VHost1 and VHost2
- D.** On VHost1, run the following command:

```
Set-VMProcessor VM2 -CompatibilityForMigrationEnabled $true
```

- 3.** You are a network administrator for Proseware.com. One of your servers is named HV1 and is running Windows Server 2012 with the Hyper-V role. HV1 is hosting 10 virtual machines on locally attached storage. It is not a member of any domain.

The available storage used by the 10 guest VMs on HV1 is close to being depleted. At the current rate of growth, the current physical disks attached to HV1 will run out of space in three months.

You want to provide more space to your guest VMs. How can you solve the storage problem with a minimum financial expense and minimum impact on users?

- A.** Perform a quick migration of the VMs on HV1 to a new server with more space.
- B.** Perform a live migration of the VMs on HV1 to a new server with more space.
- C.** Perform a storage migration of the VMs on HV1 to a new storage array with ample storage space.
- D.** Attach a new storage array with ample storage space to HV1 and expand the VHD files used by the guest VMs.



Thought experiment

Configuring and managing high availability at Proseware

You are a network administrator for Proseware.com, a software company with offices in several cities. You are designing high availability for certain applications and services at the Philadelphia branch office. You have the following goals:

- You want to ensure that two domain controllers from the Proseware.com domain remain online with high availability in the Philadelphia branch office, even if one server experiences a catastrophic failure or is brought down for maintenance. (The domain controllers will not host any operations master roles.)
- You want to ensure that a heavily used LOB application can withstand the failure of one server without experiencing any downtime or file handling errors, even during failover. The LOB application is not cluster-aware. It also frequently reads and writes data stored on a network share.

With these details in mind, answer the following questions. You can find the answers to these questions in the “Answers” section.

1. How many physical servers will you need to support your requirements, at a minimum?
2. How can you best provide high availability for file sharing?
3. How can you best provide high availability for the LOB application?
4. Which of your goals require Windows Server 2012 or Windows Server 2012 R2, as opposed to an earlier version of Windows Server?

Answers

This section contains the answers to the Objective Reviews and the Thought Experiment.

Objective 10.1: Review

1. Correct answer: C

- A. Incorrect:** A cluster storage pool can only be created from SAS disks.
- B. Incorrect:** A cluster storage pool can only be created from SAS disks. In addition, a cluster storage pool is incompatible with external RAIDs.
- C. Correct:** To create a cluster storage pool, you need three independent SAS disks that are not configured with any RAID or governed by any disk subsystem.
- D. Incorrect:** You cannot create a cluster storage pool from disks that are configured as part of a RAID or governed by any disk subsystem.

2. Correct answer: B

- A. Incorrect:** Creating a storage pool by itself might simplify management of SAN storage, but it won't minimize downtime in case of node failure. In addition, the SAN storage cannot be configured as a storage pool for the cluster because it is iSCSI based. Only SAS storage can be used for a cluster storage pool.
- B. Correct:** Keeping VM storage on a CSV will optimize live migration of the VM in case of node failure and minimize downtime. CSVs will also simplify management of SAN storage by allowing multiple failover cluster nodes to share LUNs.
- C. Incorrect:** If you provision volumes from a SAS array, you will later be able to create a storage pool for the cluster, which might simplify management of storage. However, using a SAS array will not minimize downtime in case of node failure.
- D. Incorrect:** Assigning a mirrored volume to the cluster might prevent node failure if one disk fails, but it will not minimize downtime if a node does fail. In addition, it will not simplify management of SAN storage.

3. Correct answer: C

- A. Incorrect:** This solution only performs updates once on Node1 only, not the entire cluster.
- B. Incorrect:** This solution only updates the cluster once. It doesn't minimize administrative effort because you would need to do it repeatedly.
- C. Correct:** This solution configures Cluster1 to perform Windows Updates automatically and regularly on both nodes in the cluster.
- D. Incorrect:** This solution performs updates only on Node1, not the entire cluster.

Objective 10.2: Review

1. **Correct answers:** B, D
 - A. **Incorrect:** A traditional file server for general use is best suited for users, not resource-intensive applications. In addition, a traditional file server would not easily allow you to handle an increased load as the usage of the file share increased.
 - B. **Correct:** A Scale-Out File Server allows an application to maintain file handles even during failover, which minimizes application errors. In addition, an SoFS allows you to keep all nodes active and to add additional nodes as needed to handle an increased load.
 - C. **Incorrect:** A Scale-Out File Server requires CSV storage. Choosing this storage type would not allow you to meet your requirements or reducing errors and maintaining high performance.
 - D. **Correct:** A Scale-Out File Server requires CSV storage.
2. **Correct answers:** A, C, D
 - A. **Correct:** A File Server for general use is the more suitable role to provide high availability for a file share that users (as opposed to applications) will use for file storage. In addition, only the File Server role is compatible with Data Deduplication and BranchCache.
 - B. **Incorrect:** An SoFS is not compatible with Data Deduplication or BranchCache, two features that will help you meet your requirements for the share.
 - C. **Correct:** Data Deduplication will help minimize storage space requirements.
 - D. **Correct:** BranchCache will minimize the amount of data transferred over WAN links and prevent file conflicts.
3. **Correct answer:** C
 - A. **Incorrect:** VM monitoring does indeed require Windows Server 2012 or later to be running on the host Hyper-V server and failover cluster node.
 - B. **Incorrect:** VM monitoring requires Windows Server 2012 or later to be running on the clustered VM.
 - C. **Correct:** The host and guest do not need to be members of the same domain. However, the two domains need to trust each other.
 - D. **Incorrect:** The firewall rules in the Virtual Machine Monitoring group do need to be enabled on the clustered VM.

Objective 10.3: Review

1. Correct answer: D

- A. Incorrect:** Constrained delegation is required for Kerberos authentication. You have configured CredSSP as the authentication protocol. In addition, you have received an error related to processor compatibility, not authentication.
- B. Incorrect:** VM monitoring isn't incompatible with live migration, so it wouldn't generate an error such as this one.
- C. Incorrect:** There is no reason to change the authentication protocol to Kerberos under these circumstances. CredSSP allows you to initiate a live migration when you are logged on locally to the source host.
- D. Correct:** If you enabled processor compatibility on the VM, the virtual processor will use only the features of the processor that are available on all versions of a virtualization-capable processor by the same processor manufacturer. You would see the error described if each host server used a different processor from the same manufacturer.

2. Correct answer: A

- A. Correct:** When you choose Kerberos as the authentication protocol, you need to configure constrained delegation on the computer accounts for the source and destination computers.
- B. Incorrect:** VM monitoring is not incompatible with live migration and would not generate an error such as the one described.
- C. Incorrect:** CredSSP as an authentication protocol would not enable you to initiate live migrations when you are not logged on to the source host server.
- D. Incorrect:** The error received was not related to processor compatibility, so this step would not fix the problem.

3. Correct answer: C

- A. Incorrect:** A quick migration is possible only in a failover cluster environment. In addition, purchasing a new server with ample new storage is unnecessarily costly compared to purchasing only new storage.
- B. Incorrect:** You cannot perform a live migration from a computer outside of a domain environment. In addition, purchasing a new server with ample new storage is unnecessarily costly compared to purchasing only new storage.
- C. Correct:** This option avoids the unnecessary expense of purchasing a new server and lets you transfer storage to the new storage array live, without taking your VMs offline.
- D. Incorrect:** This option will not solve your problem. If you purchase a new disk array, you need to find a way to move the VMs onto the new storage. You will be able to expand the size of the VHD files only to the point that they will use up the space on the old disks.

Thought experiment

1. Three. You need to have two highly available domain controllers even after one server is brought down. You can provide high availability for all workloads on those three servers.
2. Use an SoFS with CSVs so that the LOB application can remain connected to files even during failover.
3. You should host the LOB application in a highly available VM because the application itself isn't cluster-aware.
4. A virtualized domain controller is not recommended in older versions of Windows Server. In addition, an SoFS is available only in Windows Server 2012 and later.

Configure file and storage solutions

The Configure File and Storage Solutions domain includes just a single objective, but it's a big one: Implement Dynamic Access Control.

Access permissions for files have traditionally been controlled by a combination of share permissions and NTFS permissions. Dynamic Access Control is a set of features new to Windows Server 2012 and Windows Server 2012 R2 that adds to these two traditional access controls an optional third security gate for file access. Dynamic Access Control controls access in a way that is dependent not on user groups or file location, but on object attributes cited in access rules. The advantage of Dynamic Access Control is that it can provide highly customized and simplified management of file security, especially in a large organization.

Objectives in this chapter:

- Objective 11.1: Implement Dynamic Access Control

Objective 11.1: Implement Dynamic Access Control

Dynamic Access Control relies on file classifications, on user and device attributes called claims, and on rules and policies built from all of these elements. Dynamic Access Control, admittedly, can be very complex. On the 70-417 exam, it's likely you will need to understand only the fundamentals about this feature. However, these concepts are new and require some effort to learn even just the fundamentals well enough for the exam.

This section covers the following topics:

- Configuring claims-based authentication
- Configuring file classification
- Configuring access policies

Introduction to Dynamic Access Control

Dynamic Access Control is a new way to control access to files. It doesn't replace NTFS and share permissions but is sometimes combined with them. When Dynamic Access Control permissions are combined with the NTFS and share permissions, the most restrictive permissions always apply to the account requesting access.

You can think of Dynamic Access Control as being based on *access rules*. These rules are if-then statements built on the attributes of files, users, and devices. An example expression to serve as the basis for an access rule could be "If a user is a member of the finance department with an office on Floor 10 and is connecting from a device that is located in the company HQ, then that user can access finance files and folders designated as having a high business impact." Before you can even create such an access rule, you need to create and assign the needed attributes to all the objects mentioned in that rule. The user and device attributes are called *claims*. The file attributes are called *classifications* (or *resource properties*).

The way these three attribute types relate to an access rule is illustrated in Figure 11-1.

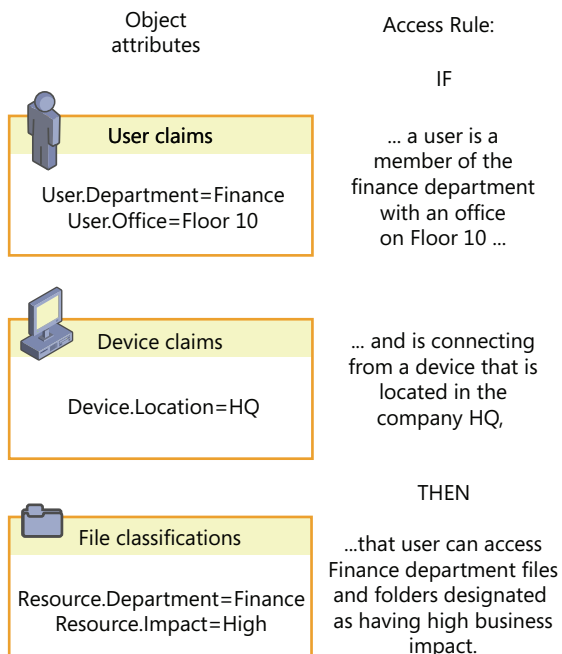


FIGURE 11-1 Access rules refer to the attributes of users, devices, and files

Dynamic Access Control is advantageous for a number of reasons. First, it allows administrators to manage file access centrally, in a way that impacts all file servers in the organization. (It should be noted, however, that you cannot enforce access rules centrally through Dynamic Access Control; you can only make access rules available for enforcement.)

Another advantage of Dynamic Access Control is that it allows you to dramatically reduce the number of user groups you would otherwise need to create and manage to implement a particular access policy. A third advantage of Dynamic Access Control is that it allows you to construct access rules in a way that is much more flexible and much more likely to correspond to the needs of your organization. Instead of access control lists (ACLs) based only on user and group accounts, you can create rules based on location, office, country, telephone number, or any other parameter that is most useful to you.

To implement Dynamic Access Control, you need at least one file server running Windows Server 2012 or later, at least one domain controller running Windows Server 2012 or later (one recommended at each site), and clients running Windows 7 or later. In addition, specific features (such as access-denied assistance) require Windows 8 or later. The domain functional level must also be set to Windows Server 2012 or later.

Even more than with most new features that you will be tested on for the 70-417 exam, Dynamic Access Control is best understood by working with it hands-on. It isn't the easiest of features to master, but without implementing it on a test network, it can seem more complicated than it really is. To prepare for the exam, then, use the following sections as a general walk-through for hands-on configuration, if at all possible. Plan to walk through these steps at least twice, and it will all begin to make sense.



EXAM TIP

Expect to see questions on the 70-417 exam for which you need to remember the exact order of steps necessary to configure Dynamic Access Control. Use this chapter to help understand the concepts in Dynamic Access Control and to learn the main steps as they are presented in each section.

Configuring claims-based authentication

In the Microsoft implementation of the Kerberos authentication protocol, the Kerberos ticket-granting ticket (TGT) includes a special access token that is used to help authorize users for access to domain resources. This access token always contains the user's ID and his or her group memberships.

Dynamic Access Control, however, relies on an expanded Kerberos access token that includes more than the user's ID and group memberships. Besides that usual information, the expanded token used in Dynamic Access Control includes certain attribute values—called *claims*—about the user, about the device to which the user is signed on, and about the same device's own group memberships. The expanded access token used in Dynamic Access Control is illustrated in Figure 11-2.

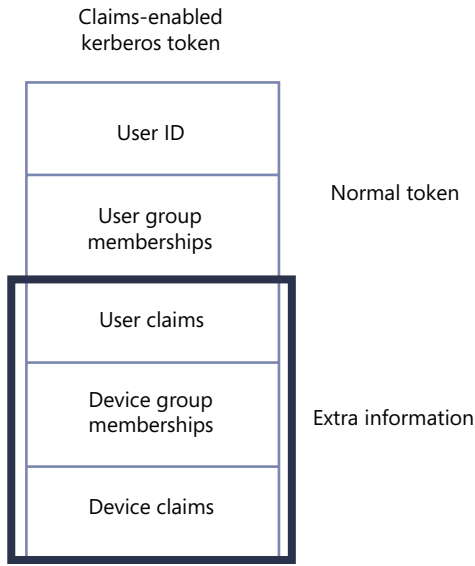


FIGURE 11-2 The Kerberos access token used in Dynamic Access Control

To configure a Dynamic Access Control policy, you need to perform the following steps:

1. Define the types of claims about users and devices you want to include in Kerberos access tokens.
2. Configure Active Directory Domain Services to use the expanded Kerberos access tokens that include these claims.

Step 1: Define user and device claims types

In this step, you choose the specific user and device properties that will be presented as claims in the Kerberos access token whenever access permissions are evaluated. User and device claim types correspond to names of Active Directory attributes (such as “Department” or “City”) for user and computer account objects. The actual claim values included in a token are copied from the corresponding Active Directory attribute values. Because access rules refer to these claim types in their specifications about who is allowed or denied access to a given resource, you want to define claims types you will need later when you create access control rules.

You can use Active Directory Administrative Center to configure the user and device claim types. In the console tree, select tree view and then navigate to Dynamic Access Control\Claim Types. Right-click Claim Types, click New, and then select Claim Type, as shown in Figure 11-3.

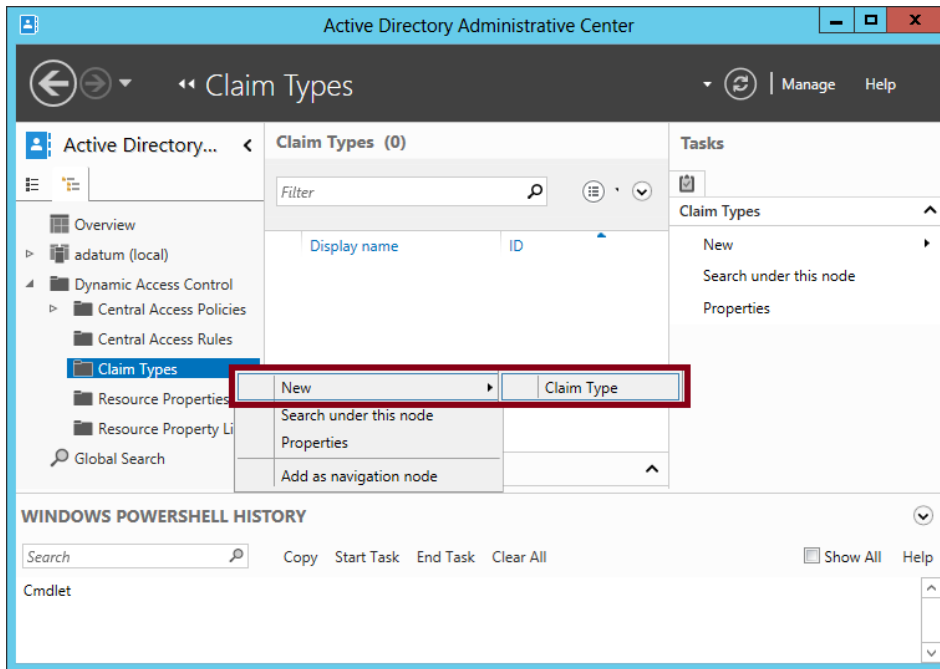


FIGURE 11-3 Creating a new claim type for a user or device

The Create Claim Type page that opens is shown in Figure 11-4.




FIGURE 11-4 Creating a claim type for a user object

In the Source Attribute section, click the Active Directory object attribute name that you want to use as the basis of the claim type. You can also specify whether you want this claim type to be issued for users, for computers (devices), or both. For example, if you plan to define rules that include references to the department of either the user or the device to which a user is signed on, you should select both User and Computer when you create the Department claim type.

In the Suggested Values section, you can provide a list of suggested matching values that you will later be able to supply in access rules. For example, if you plan to create access rules that specify user or device Department values such as "Finance," "Engineering," "Operations," "Marketing," and "Sales," you can precreate those same strings as suggested values now when you are creating the claim type. Note that if you define any suggested values, those values you supply will be the only ones available to select when you create rules that refer to the claim type.

Step 2: Enable Kerberos support for claims-based access control

In this step, you use Group Policy to enable Kerberos support for claims on domain controllers. This step ensures that Kerberos tokens include claims information and that this information can then be evaluated by domain controllers for access authorization.

In the Group Policy Management Console, create or edit a Group Policy Object (GPO) linked to the domain controllers organizational unit (OU), and then enable the following setting: Computer Configuration/Policies/Administrative Templates/System/KDC/KDC Support For Claims, Compound Authentication, And Kerberos Armoring. (Within the Policy Setting dialog box, leave selected the default option of Supported.)

The requirement that you set this policy for claims-based authorization, and that you should do so at the domain controllers OU level, is one of the most likely aspects about Dynamic Access Control that you'll be tested on during the 70-417 exam. Learn to recognize not only the full name of this policy, but also the possible ways its name might be shortened. (For example, an answer choice might say simply "Use Group Policy to enable Kerberos armoring on the domain controllers OU.") The location of this policy setting within a GPO is shown in Figure 11-5.

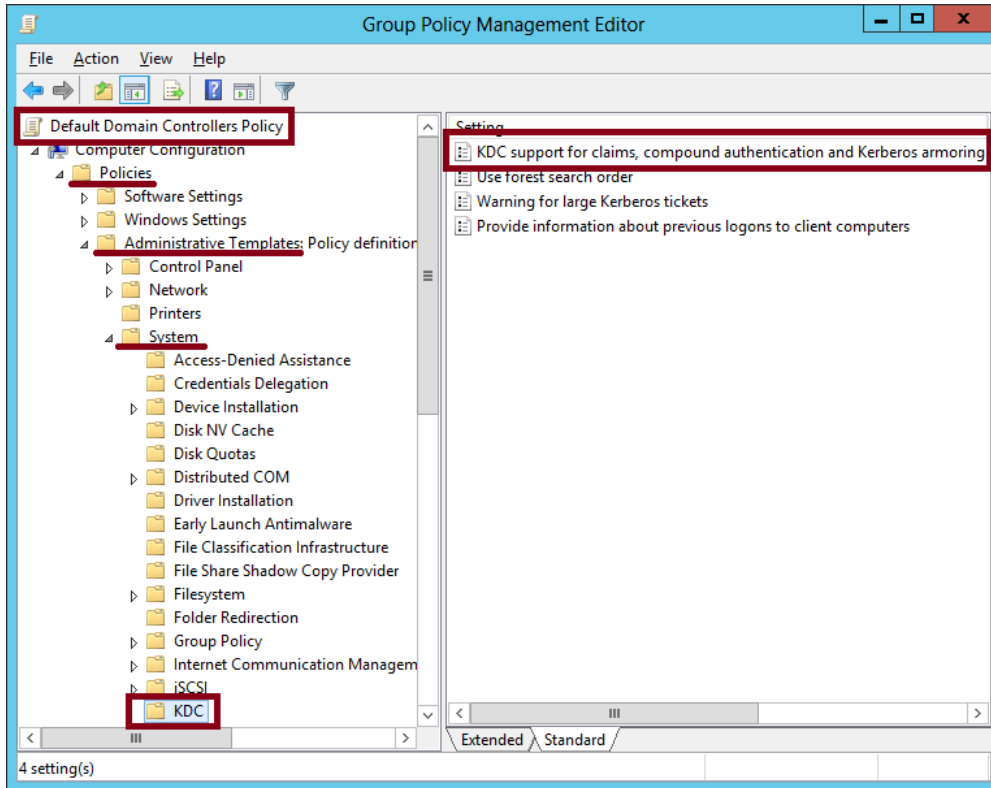


FIGURE 11-5 Enabling Kerberos support for claims

Configuring file classification

File classification refers to the process of adding attributes to the properties of files and folders. These attributes allow you to construct access rules that apply to these resources. Configuring file classification can be broken down into the following four steps:

1. Enable or create selected resource properties.
2. Add resource properties to a resource property list.
3. Update Active Directory file and folder objects.
4. Classify files and folders.

Step 1: Enable or create selected resource properties

You perform this step on a domain controller running Windows Server 2012 or Windows Server 2012 R2, in Active Directory Administrative Center. In the console tree, select tree view (the right tab in the navigation pane) and then the Resource Properties container, as shown in Figure 11-6. To enable a property, select it in the center pane and then click Enable on the Tasks (right) pane.

Resource properties correspond to attribute categories, such as Department, that you can make appear on the Classification tab of the Properties dialog box of files and folders. You make a resource property appear on this Classification tab by first enabling the property and then performing steps 2 and 3 described later in this section. Generally, you should enable only the resource properties you plan to use later in access rules. For example, if your eventual goal is to create and apply the access rule shown in Figure 11-1, you should enable the Department and Impact resource properties.

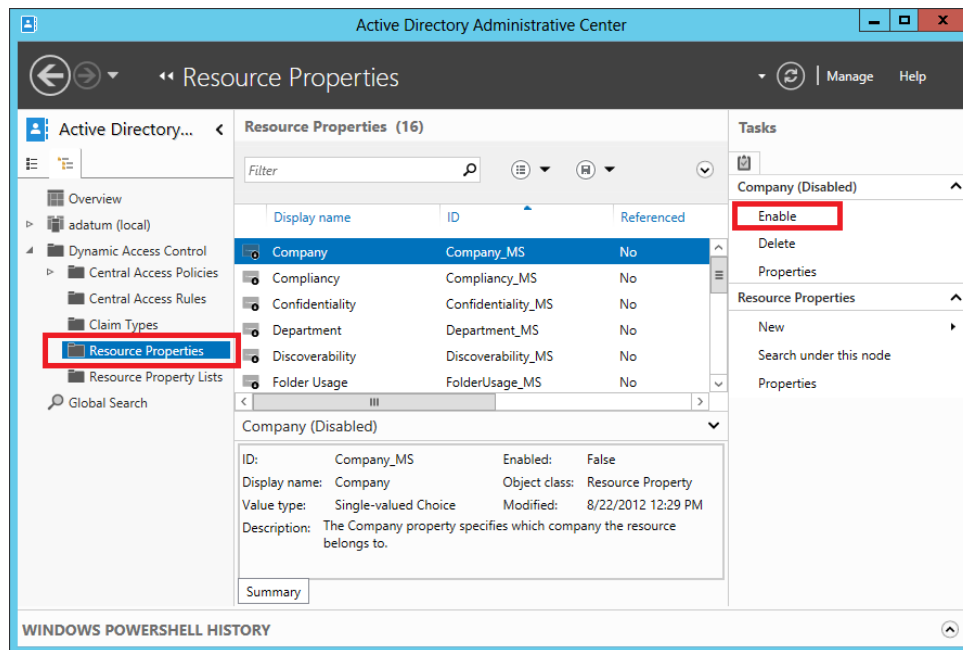


FIGURE 11-6 Enabling a resource property

Windows Server 2012 and Windows Server 2012 R2 include 16 predefined resource properties, including Department, Impact, Compliancy, Intellectual Property, and Confidentiality. These resource properties include predefined suggested values you can eventually assign to objects, values such as the specific names of departments; High, Medium, or Low; and Yes or No. However, if a resource property you need isn't predefined (such as City or Country), you can create it and define suggested values you need, such as London, New York, UK, US, and so on. Any new resource properties you create are automatically enabled.

Step 2: Add resource properties to a resource property list

After you enable your desired resource properties, you have to add them to a resource property list before they can be applied to objects. Begin by selecting the Resource Property Lists container in Active Directory Administrative Center. One predefined list is available, named Global Resource Property List. If you want the same classifications to be available for all objects, use this list. To add the resource properties you have enabled, right-click the list and select Add Resource Properties, as shown in Figure 11-7. In the Select Resource Properties dialog box that opens, add the desired resource properties that you have enabled, and click OK.

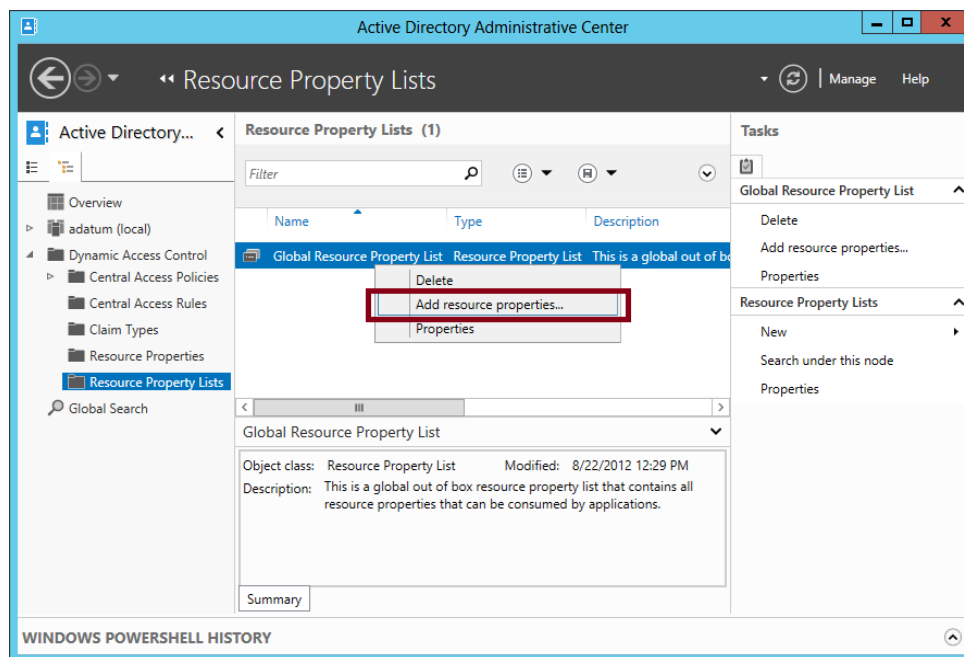


FIGURE 11-7 Adding resource properties to a resource property list



EXAM TIP

Beware of incorrect answer choices that suggest you need to *create* a resource property list when you configure file classification. You don't need to create a resource property list. You just need to add the resource properties to a list (usually the built-in Global Resource Property List).

Step 3: Update Active Directory file and folder objects

To update Active Directory Domain Services with the new classifiable properties, you now need to run the following cmdlet on a file server on which the File Server Resource Manager (FSRM) component of the File Server role has been installed:

```
Update-FSRMClassificationPropertyDefinition
```

After you perform this step, the resource properties you chose in step 1 appear on the Classification tab of every file and folder on that file server. The Classification tab is shown in Figure 11-8.

Note that this cmdlet is one of the most likely items related to Dynamic Access Control to appear on the 70-417 exam. Make sure you understand its function.

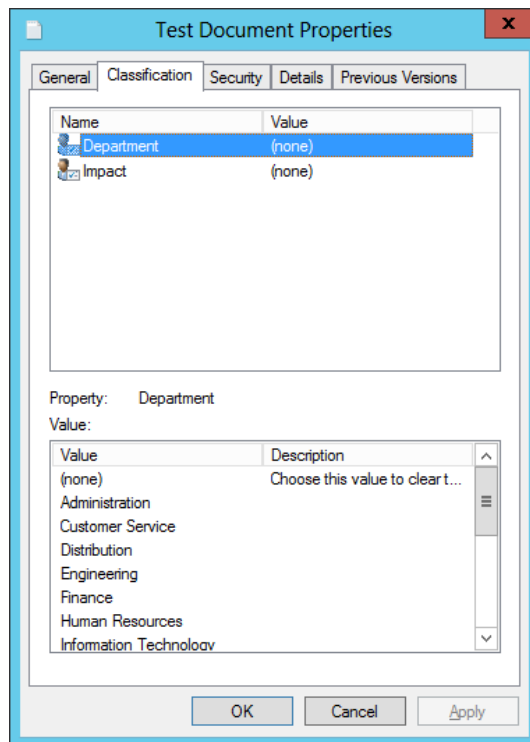


FIGURE 11-8 Resource properties on the Classification tab



EXAM TIP

Remember that you need to install File Server Resource Manager in order for the Classification tab to appear in the properties of files and folders.

Step 4: Classify files and folders

Objects in the file structure can be classified manually or automatically. The following sections provide instructions about how to classify files by using both of these strategies.

MANUAL CLASSIFICATION

To classify file objects manually, you can select and apply a resource property value on the Classification tab directly on selected files or on their parent folder. For example, for the folder shown in Figure 11-9, the Finance value and the High value have been selected for the Department and Impact properties, respectively. When you click Apply, these classifications will automatically be applied to all child objects within the folder.

Note that child objects keep these classification settings until they are reapplied. Files do not automatically inherit the values of other parent folders if they are moved into those other folders. In fact, the classifications remain applied to those objects even when you copy them from computer to computer. However, you can only see and read these classifications that have been applied to objects after you install FSRM and run the Update-FSRMClassificationPropertyDefinition cmdlet.

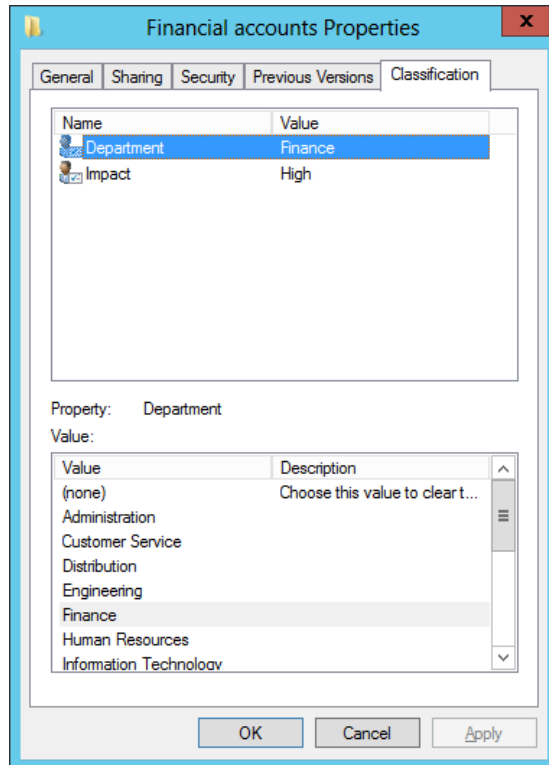


FIGURE 11-9 Classification values set on a parent folder

AUTOMATIC CLASSIFICATION

Windows Server 2012 and Windows Server 2012 R2 include a built-in file classifier that can be configured to automatically classify files within targeted folders. You can automatically classify all files within targeted folders, or you can restrict this function to a subset of the files, limiting classification to those Microsoft documents with contents that include a match of a specified expression. You can also restrict classification to the files selected by a custom Windows PowerShell script. Besides this built-in functionality, automatic classification (and Dynamic Access Control in general) can be greatly extended through third-party applications.

To start configuring automatic file classification, you first need to install the FSRM component of the File and Storage Services role. Then, in the File Server Resource Manager console tree, navigate to Classification Management\Classification Rules. In the Actions pane, click Create Classification Rule, as shown in Figure 11-10.

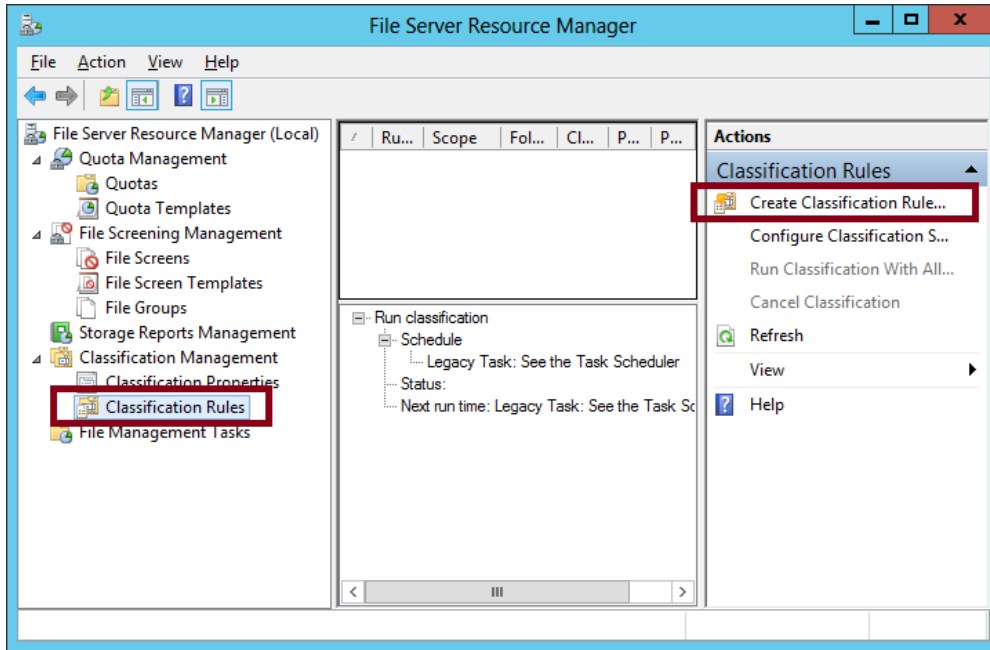


FIGURE 11-10 Creating a classification rule

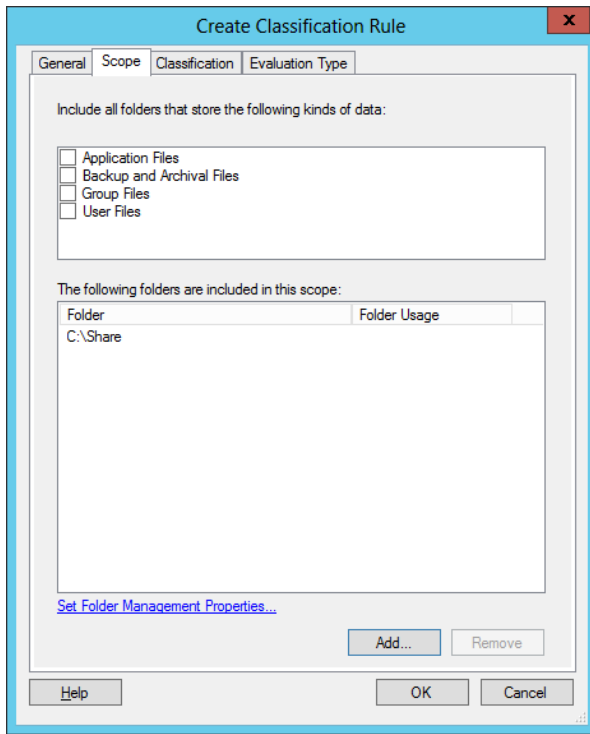


FIGURE 11-11 Setting the scope for a classification rule

This step opens the Create Classification Rule dialog box. On the General tab, enter a name and description for the new rule. The General tab also includes an Enabled check box, which is selected by default.

On the Scope tab, shown in Figure 11-11, click Add to select the folders where this rule will apply. The classification rule applies to all folders and their subfolders in the list. Alternatively, you can target *all* folders that store *any* of the following selected classifications of data: Application Files, Backup And Archival Files, Group Files, or User Files.

On the Classification tab, shown in Figure 11-12, choose a classification method along with the classification value for one selected property that the classification rule will assign.

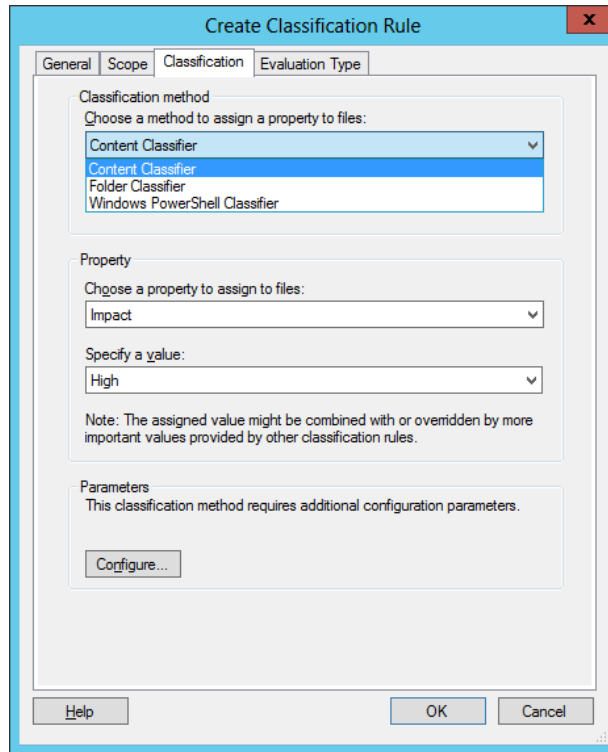


FIGURE 11-12 Configuring a classification method and property value

For a classification method, there are three options:

- **Content Classifier** Searches Microsoft documents for a text or regular expression string.
- **Folder Classifier** Assigns the property value to all files that fall within the scope of the classification rule.
- **Windows PowerShell Classifier** Prompts you to specify a script to determine the target files within the scope of the classification rule.

Click Configure to further configure this option by using the Classification Parameters dialog box, shown in Figure 11-13.

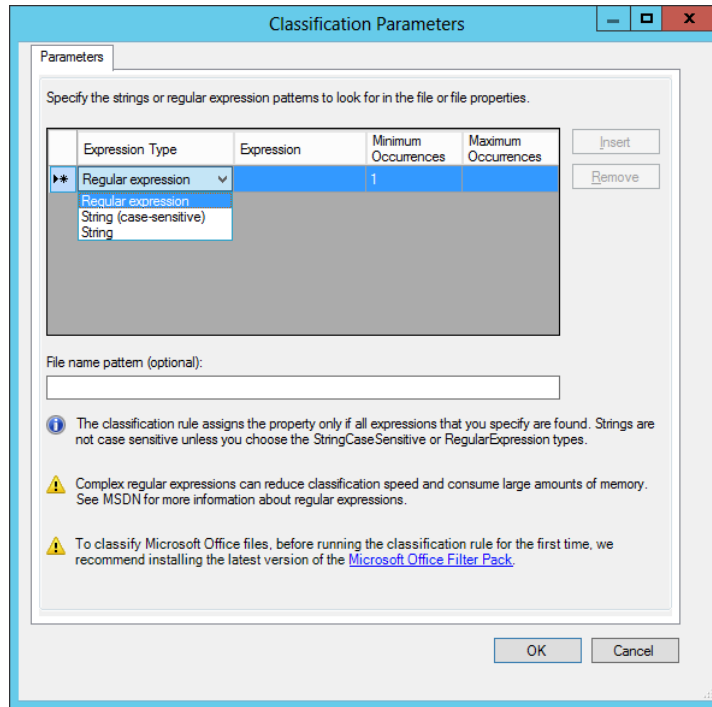


FIGURE 11-13 Configuring a content search for automatic classification

This dialog box lets you specify an expression that will be searched for in the content of Microsoft documents that fall within the scope of the classification rule. If the content search results in a match for the specified expression, the file is tagged with the property value specified on the Classification tab of the Create Classification Rule dialog box.

You can choose one of three expression types to search for: string, case-sensitive string, or *regular expression*. A regular expression, sometimes called a *regex*, is used in programming to match *patterns* of text strings as opposed to exact sequences of specific numbers or letters. A regular expression is often a useful matching mechanism for classifying files that include sensitive numbers such as credit card numbers.

The following is an example of a regular expression. It matches credit card numbers from most vendors:

```
^(4\d{3})|(5[1-5]\d{2})|(6011)|(34\d{1})|(37\d{1})-?\d{4}-?\d{4}-?\d{4}|3[4,7]
[\d\s-]{15}$
```

The Evaluation Type tab is the final tab of the Create Classification Rule dialog box. On this tab, you choose how to handle files that already exist within the scope of the classification rule. By default, the classification rule does not apply to preexisting files. You can choose,

however, to run the rule against existing files. If matches are found, you can either overwrite any existing classification that conflicts with the new value, or attempt to aggregate them if possible.

After you create the desired classification rule, click Configure Classification Schedule in File Server Resource Manager to determine how often you want the rule to run. This step opens the File Server Resource Manager Options dialog box. This is the same dialog box that opens when you select Configure Options from the shortcut menu of the main (parent) node in the File Server Resource Manager console tree. On the Automatic Classification tab, shown in Figure 11-14, select the Enable Fixed Schedule check box. You must then specify days and times at which you want the rule to run. In addition, you can select the Allow Continuous Classification For New Files check box to run the rule on newly created or edited files that fall within the scope of the rule, and on existing files that are moved to a new location that falls within the scope of the rule. (For the exam, be sure to remember the option for continuous classification.)

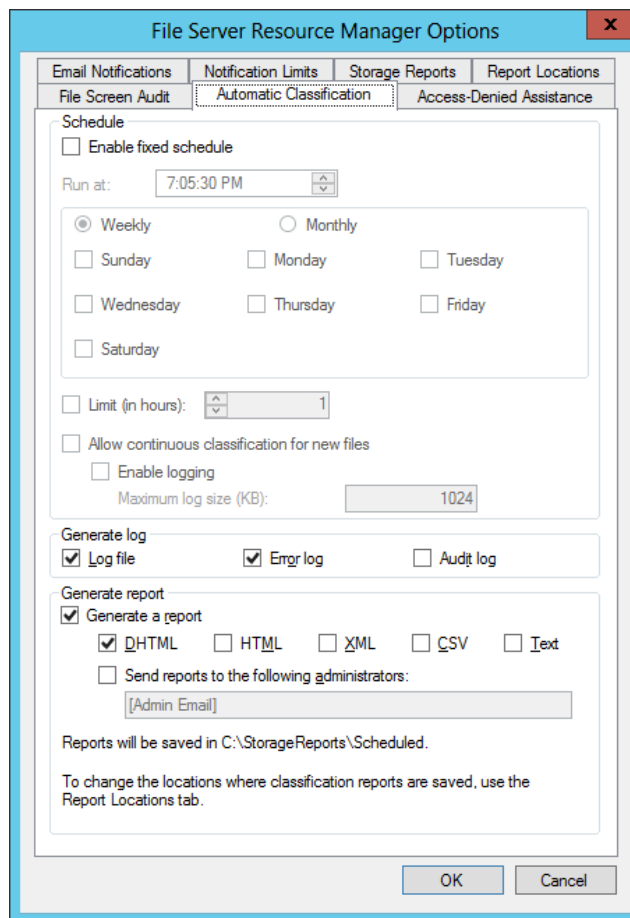


FIGURE 11-14 Configuring a schedule for a classification rule

After configuring the schedule, you can click Run Classification With All Rules Now in the Actions pane of File Server Resource Manager. This step will run all rules immediately and classify the targeted files.

ACCESS-DENIED ASSISTANCE

In Windows Server 2012 and Windows Server 2012 R2, you can enable the local file server to provide helpful information to a user whose access to a file or folder has been denied.

To enable this functionality, open the dialog box by right-clicking the parent node of the File Server Resource Manager console tree and then clicking Configure Options. On the Access-Denied Assistance tab of the File Server Resource Manager Options dialog box that opens, select the Enable Access-Denied Assistance check box, as shown in Figure 11-15.

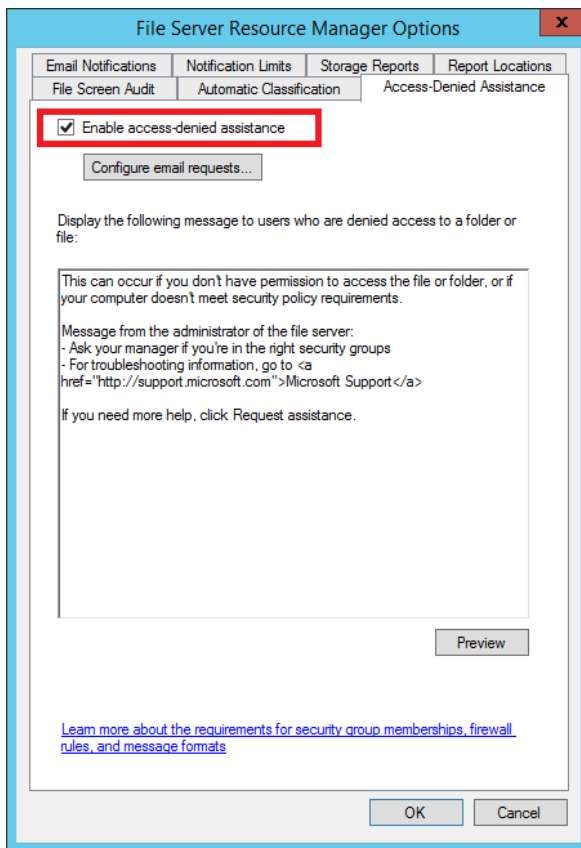


FIGURE 11-15 Enabling access-denied assistance

In the Display The Following Message text box, you can type a custom message that users will see when they are denied access to a file or folder. You can also add certain variables in brackets that will insert customized text, such as:

- **[Original File Path]** The original file path that was accessed by the user.

- **[Original File Path Folder]** The parent folder of the original file path that was accessed by the user.
- **[Admin Email]** The administrator email recipient list.
- **[Data Owner Email]** The data owner email recipient list.

You can also configure the file server to provide in access-denied messages a Request Assistance button, which allows the user who was denied access to send an email to a pre-defined user. To configure this option, click Configure Email Requests, select the Enable Users To Request Assistance check box, and then click OK.

Finally, you can configure access-denied assistance on a per-folder basis. To do that, select Classification Properties in the console tree and then click Set Folder Management Properties in the Actions pane, as shown in Figure 11-16. In the dialog box that opens, use the Add button to specify the folder for which you want to configure access-denied assistance.

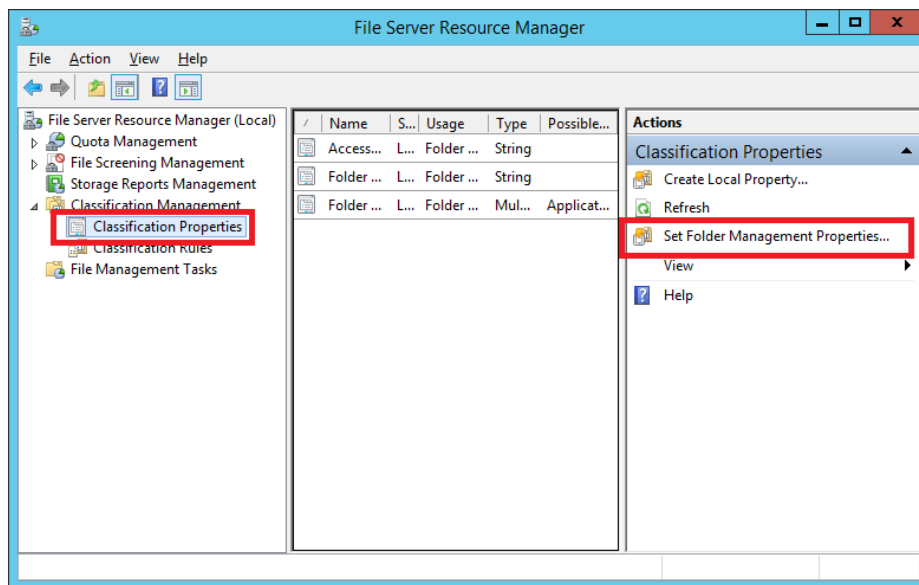


FIGURE 11-16 Configuring access-denied assistance for a specific folder



EXAM TIP

You need to remember how to configure access-denied assistance, including how to configure access-denied assistance for a specific folder.

RMS ENCRYPTION If your network environment includes an Active Directory Rights Management Service server, you can use FSRM to automatically apply RMS encryption to files in designated folders. This feature is configured through the File Management tasks node. For more information, visit <http://technet.microsoft.com/en-us/library/hh831572.aspx>.

Configuring access policies

Finally, you are ready to create access policies after you have assigned attributes to users, devices, and files. To configure access policies, you need to perform the following steps:

1. Create a claims-based central access policy.
2. Use Group Policy to deploy this central access policy to your file servers.

Step 1: Create a central access policy that includes claims

This step consists of two parts, both of which you can perform in Active Directory Administrative Center. First, you create one or more central access rules that include claims. Then, you add those rules to a central access policy.



EXAM TIP

Normally you'd want to create access rules and then create the central access policy to add them to.

CREATE A NEW CENTRAL ACCESS RULE

A central access rule is similar to an ACL in that it describes which conditions must be met for access to be granted to a resource.

To create a new central access rule, in Active Directory Administrative Center, select tree view in the navigation pane and then select Central Access Rules. In the Tasks pane, click New, and then click Central Access Rule. This step opens the Create Central Access Rule page, shown in Figure 11-17.

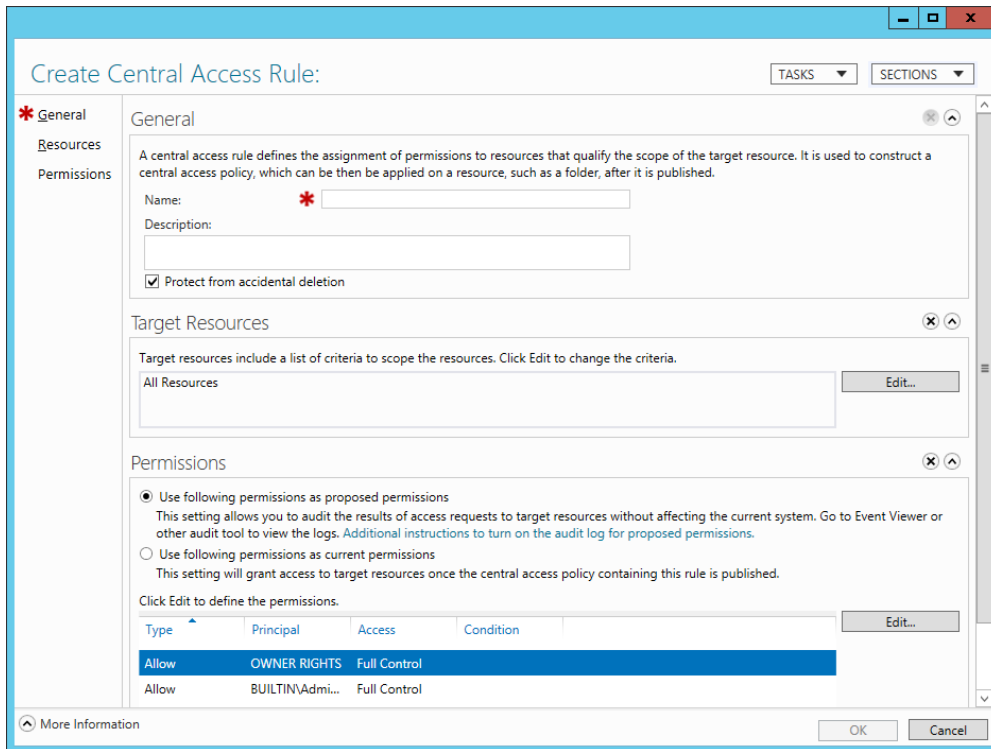


FIGURE 11-17 Creating a new central access rule

Use the following instructions to complete the page:

1. In the Name text box, type the name you want to give to the rule.
2. In the Target Resources section, click Edit, and in the Central Access Rule dialog box, add the conditions that match the target resources for which you want to define access. For example, if your goal is to define access permissions to resources that have been configured both with a Department classification property of Finance and with an Impact classification property of High, then you want to add the two conditions configured as shown in Figure 11-18.

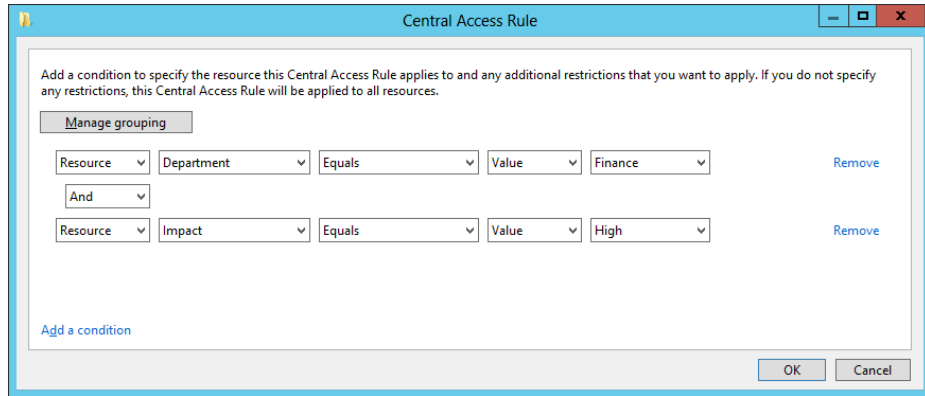


FIGURE 11-18 Configuring matching conditions for target resources



EXAM TIP

You need to remember the general syntax of central access rule conditions, such as those shown in Figure 11-18.

3. In the Permissions section of the Create Central Access Rule page, select Use Following Permissions As Current Permissions and then click Edit. In the Advanced Security Settings For Permissions dialog box, click Add to open the Permission Entry For Permissions dialog box, shown in Figure 11-19. In this dialog box, do the following:
 - A. Near the top of the dialog box, click Select A Principal. A principal is another name for a user or group account. To configure Dynamic Access Control, you normally want to select Authenticated Users as the principal. (Remember this point both for the real world and the exam.)
 - B. In the middle of the dialog box, beneath Basic Permissions, select the permissions that you want to assign to users who match the conditions in your rule.
 - C. Near the bottom of the dialog box, add conditions that match the users for whom you want to define access. For example, if you want to provide access only to users whose accounts in Active Directory have defined a Department value of Finance and an Office value of Floor 10, and who are signed on to computers whose accounts in Active Directory have defined a Location value of HQ, then you want to add the three conditions configured as shown in Figure 11-19. Remember that if Authenticated Users attempt to access the target resource and do *not* match these conditions, the users will be completely denied access (with the exception of the file owner).



EXAM TIP

As an alternative to step 3, you can leave selected the “Use Following Permissions as Proposed Permissions” option, which you can see in Figure 11-17. This option is used to stage a policy rule. Staging policies can be used to monitor the effects of a new policy entry before you enable it. You can use this option with the Group Policy setting name Audit Central Access Policy Staging. For more information, see the procedure described at http://technet.microsoft.com/en-us/library/hh846167.aspx#BKMK_1_2.

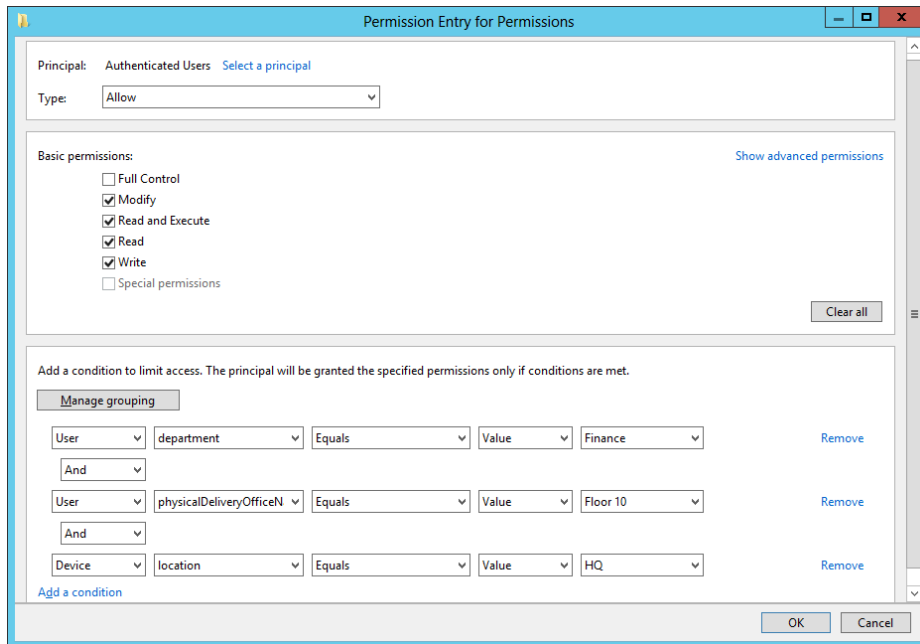


FIGURE 11-19 Configuring permissions and matching conditions for users and devices

4. Click OK three times to finish and return to Active Directory Administrative Center.

ADD CENTRAL ACCESS RULES TO A CENTRAL ACCESS POLICY

In the navigation pane of Active Directory Administrative Center, select tree view and then click Central Access Policies. In the Tasks pane, click New, and then click Central Access Policy.

On the Create Central Access Policy page that opens, do the following:

1. In the Name text box, type the name you want to assign to the policy.
2. In Member Central Access Rules, click Add and then add the desired central access rules you have created. Click OK twice to return to Active Directory Administrative Center.

MULTIPLE CENTRAL ACCESS RULES When you include multiple access rules in a policy, all the rules will be applied along with that policy when the policy is applied. The most restrictive access permissions always take effect when two rules provide different levels of access to the same user.

Step 2: Deploy central access policy to file servers

In this step, you configure a policy setting at the domain level that will deliver chosen central access policies to your file servers. Note that you can't actually *enforce a central access policy* by using Group Policy. You use Group Policy only to make desired central access policies available for selection in the Advanced Security Settings dialog box of all objects within the folder structure on file servers. The policy must then be applied to the object (usually a folder) manually.

To make your central access policies available to objects on file servers, in a GPO linked to the domain, navigate to Computer Configuration/Policies/Windows Settings/Security Settings/File System, and then click Central Access Policy. On the Action menu, select Manage Central Access Policies. In the Central Access Policies Configuration dialog box, add the central access policies that you want to make available to file servers, and then click OK.

When this Group Policy policy setting is enforced, the central access policies appear on a new Central Policy tab of this dialog box, shown in Figure 11-20. A particular central access policy applies to a folder or file object only when an administrator selects and applies it manually in these advanced security settings.

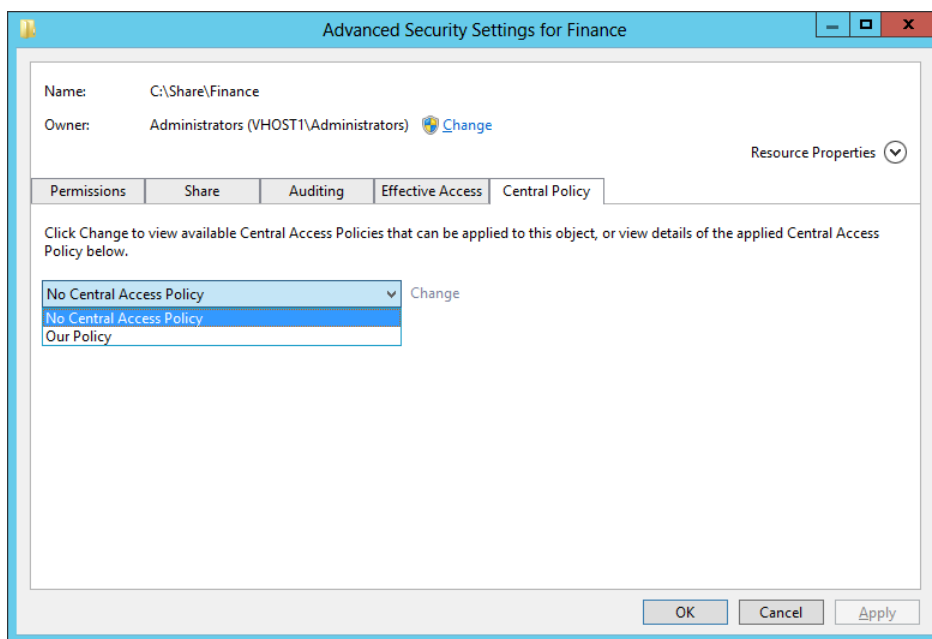


FIGURE 11-20 The Central Policy tab of the Advanced Security Settings dialog box

MORE INFO For more practice implementing Dynamic Access Control, you can complete the Windows Server 2012 virtual lab “Using Dynamic Access Control to Automatically and Centrally Secure Data” at <http://go.microsoft.com/?linkid=9838457> or perform the Dynamic Access Control walkthrough named “Deploy a Central Access Policy (Demonstration Steps)” at <http://technet.microsoft.com/en-us/library/hh846167.aspx>. (Microsoft virtual labs require Internet Explorer.)



EXAM TIP

The Configure and Optimize Storage objective is one of the original objectives in the File and Storage Solutions domain on the 70-412 exam. The Configure and Optimize Storage objective mostly covers iSCSI support in Windows Server 2012 and Windows Server 2012 R2. Although this objective hasn't been identified as a topic on the 70-417 exam, it's recommended that you learn the basics about configuring iSCSI in these operating systems anyway. For example, be sure you know the steps required to configure the iSCSI Target and iSCSI Initiator as a way to provide and provision storage in Windows Server 2012 and Windows Server 2012 R2. You should also understand the function of iSNS, a DNS-like feature used for locating iSCSI resources. To practice implementing iSCSI on Windows Server 2012, complete the virtual lab at <http://go.microsoft.com/?linkid=9838443>. (Microsoft virtual labs require Internet Explorer.)

Objective summary

- Dynamic Access Control is a new option for setting access permissions to file and folder objects in Windows Server 2012 and Windows Server 2012 R2. Dynamic Access Control works by assigning file classifications to target resources, configuring user and device claims, and then creating rules that describe conditions for access.
- Dynamic Access Control relies on a modified form of Kerberos in which user tokens are expanded to include extra information called claims about the user and the device from which the user is connecting. To support this functionality, you need to enable Key Distribution Center support for claims-based authentication in Group Policy at the Domain Controllers OU level. You also need to define the claims types that you will include in the Kerberos token for each user.
- To assign file classifications, first enable chosen resource properties in Active Directory and add the properties to a property list. Afterward, run the Update-FSRMClassificationPropertyDefinition cmdlet. Then, configure classification values of desired file or folder objects on the Classification tab of the Properties dialog box. You can also use File Server Resource Manager to configure file classification rules that classify files automatically, for example, on the basis of an expression found in the contents of the file.

- A central access rule includes one or more conditional expressions that match target resources and one or more conditional expressions that match users or devices and defines permissions to the target resources. One or more central access rules must be added to a central access policy before it can be deployed to file servers.
- You use Group Policy to make central access policies available to file and folder objects. A central policy must be selected and enforced manually on a file or folder.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the “Answers” section at the end of the chapter.

1. You are a network administrator for Adatum.com. The Adatum.com network consists of a single domain that spans branch offices in New York and London. Within the Adatum.com domain, the users and computers within the New York office are contained in an OU named US, and the users and computers within the London office are contained in an OU named UK.

You want to be able to classify data as originating from either the New York or the London office. You create a resource property named Country and configure the suggested values “US” and “UK.” You want administrators in both the New York and London offices to see the Country resource property appear on the Classification tab of files and folder properties.

What should you do next?

- A. Run the Update-FSRMClassificationPropertyDefinition cmdlet.
 - B. Enable the Country resource property.
 - C. Create a classification rule.
 - D. Add the Country property to a resource property list.
2. Your organization’s network consists of a single Active Directory domain. All servers are running Windows Server 2012 and all clients are running Windows 8. You want to enable claims-based access authorization for users in your domain. Which of the following steps should you take to take to achieve this goal?
 - A. Enable the policy setting KDC Support For Claims, Compound Authentication, And Kerberos Armoring in a GPO at the Domain Controllers OU level.
 - B. Enable the policy setting KDC Support For Claims, Compound Authentication, And Kerberos Armoring in a GPO at the domain level.
 - C. Enable the policy setting Kerberos Support For Claims, Compound Authentication, And Kerberos Armoring in a GPO at the Domain Controllers OU level.
 - D. Enable the policy setting Kerberos Support For Claims, Compound Authentication, And Kerberos Armoring in a GPO at the domain level.

- 3.** You are a network administrator for Proseware.com. The Proseware.com network consists of a single Active Directory domain. All servers in the network are running Windows Server 2012, and all clients are running Windows 8.

On a file server named FileSrv1, your manager has created five new file shares named Finance, Marketing, Sales, Operations, and Legal. On each share, your manager has assigned Full Control to Authenticated Users for both the NTFS and share permissions.

Your manager now asks you to configure permissions to the contents of each departmental file share so that Full Control access is restricted to members of the corresponding department, and that no other users are allowed any access. Your manager also wants you to ensure that files within each departmental share can be traced to their origin even when they are moved from their original share location.

Which of the following steps will allow you to meet these stated goals? (Choose two. Each answer represents part of the solution.)

- A.** On each new shared folder, remove all currently configured NTFS permissions and then grant Full Control NTFS permissions to a security group that includes all the members of the corresponding department only.
- B.** On each new shared folder, remove all currently configured share permissions and then grant Full Control share permissions to a security group that includes all the members of the corresponding department only.
- C.** On each department's shared folder, configure a Department classification property value that corresponds to the name of the department.
- D.** On each department's shared folder, apply a central access policy that assigns to members of the appropriate department Full Control permissions on files assigned with a matching Department value classification.



Thought experiment

Implementing Dynamic Access Control at Adventure Works

You are a network administrator for Adventure Works, Inc., a rapidly growing company based in Seattle that has just opened its first branch office in Denver. The network consists of a single Active Directory domain, Adventureworks.com. All servers are running either Windows Server 2008 R2 or Windows Server 2012, and all clients are running either Windows 7 or Windows 8. The two sites are linked by a site-to-site VPN.

The Seattle and Denver offices each include a main file server, named FSSeattle1 and FSDenver1, respectively, that is shared by all users in the local office location. DFS has been configured so that the same five shares are available to authorized users in both offices. Each share is used by one company-wide department, including Finance, Sales and Marketing, Operations, Human Resources, and Research and Development.

Both office locations include employees from each of the five departments.

A goal for the IT department is to address security concerns about confidential information while making all other information available to members of each department.

With the preceding information in mind, answer the following questions. You can find the answers to these questions in the “Answers” section.

1. If you wanted to limit access to some files within each department share to members of each office site, how can you best achieve this goal by using NTFS file permissions?
2. Given the information provided about the network, what changes might you need to make to ensure that Dynamic Access Control can be implemented on the network?
3. You want to make sure that when employees at one office designate a file in their department share as highly confidential, that the file can be viewed only from computers with account properties in Active Directory that indicate the same physical delivery office name as that of the user. How might you achieve this goal by using Dynamic Access Control permissions only? (Describe or list resource properties, claims types, and the central access rules you would need to create. You can assume that all informational fields are filled out in the properties of both user and computer accounts at both locations.)
4. What changes must you make to the network before you can configure detailed assistance to all users who are denied access to a resource?

Answers

This section contains the answers to the Objective Review and the Thought Experiment.

Objective 11.1: Review

1. Correct answer: D

- A. Incorrect:** You should run this cmdlet after you add the new resource property to a resource property list.
- B. Incorrect:** You don't need to enable new resource properties that you create. They are already enabled when you create them.
- C. Incorrect:** Optionally, you can create a classification rule to classify files and folders automatically. However, you can take this step only later, after you have updated file and folder objects.
- D. Correct:** After you create or enable a resource property, you need to add it to a resource property list. Only then can you update file and folder objects so that they include this resource property on the Classification tab.

2. Correct answer: A

- A. Correct:** To enable claims-based authorization in your domain, you should enable this policy setting at the domain controller level.
- B. Incorrect:** You should enable this policy setting at the domain controller level, not at the domain level.
- C. Incorrect:** This policy setting enables computers to request claims. It is used for policy auditing, not for enabling claims-based authorization.
- D. Incorrect:** This policy setting enables computers to request claims. It is used for policy auditing, not for enabling claims-based authorization.

3. Correct answers: C, D

- A. Incorrect:** Changing the NTFS permissions will restrict access to members of the appropriate department, but it will not provide any information about files that will allow them to be traced when they are moved outside of the shared folder.
- B. Incorrect:** Changing the share permissions will restrict access to members of the appropriate department when they connect over the network, but it will not provide any information about files that will allow them to be traced when they are moved outside of the shared folder.
- C. Correct:** Configuring a Department property value will allow you to classify the files in each departmental shared folder as belonging to that department, even when they leave that folder.
- D. Correct:** Applying this type of central access policy to each shared folder will configure the files within these folders with appropriate access permissions.

Thought experiment

1. You should create a security group for members of each site-specific department, such as Seattle-Finance and Denver-Finance. Then you could create a folder in each department share to which only members of each site-specific department had access.
2. You might need to install a Windows Server 2012 domain controller at each site.
3. You can configure the following:
 - Resource property: Confidentiality
 - Claims types: Office name (Physical-Delivery-Office-Name) for both users and computers
 - Access rule, target resource conditional expression: Resource.Confidentiality Equals High
 - Access rule, permissions: Authenticated users = Full Control. Conditional expression: Device.physicalDeliveryOfficeName Equals User.physicalDeliveryOfficeName
4. You must first upgrade all clients to Windows 8.

Implement business continuity and disaster recovery

This exam domain includes two objectives that, on the 70-417 upgrade exam, each emphasizes one feature new to Windows Server 2012 and Windows Server 2012 R2. For the Configure and Manage Backups objective, the new feature is Windows Azure Backup. (This feature is actually an optional online extension to Windows Server Backup in Windows Server 2012 and Windows Server 2012 R2.) For the Configure Site-level Fault Tolerance objective, the new feature is one of the most interesting Windows Server features ever: Hyper-V Replica.



EXAM TIP

Microsoft has announced that Windows Azure is being renamed Microsoft Azure. You can expect the name change to be reflected in the exams when the exams are updated.

Objectives in this chapter:

- Objective 12.1: Configure and manage backups
- Objective 12.2: Configure site-level fault tolerance

Objective 12.1: Configure and manage backups

Windows Azure Backup (formerly Windows Azure Online Backup) is an online server backup service that was first introduced at approximately the same time as Windows Server 2012. Windows Azure Backup is now much more closely integrated with Windows Azure services than it was when Windows Server 2012 was first released. The original procedures for configuring online server backups to Windows Azure have therefore changed since this feature first appeared.

As with all features new to Windows Server 2012 and Windows Server 2012 R2, Windows Azure Backup is a topic that's likely to appear in an exam question. Fortunately, it's also an easy topic to understand, provided you create a trial Windows Azure account and get some hands-on experience with it.

This section covers the following topics:

- Configuring online backups
- Performing Windows Azure Backups in Windows PowerShell

Certificate requirements for Windows Azure Backup

Before you can configure online backups to Windows Azure, you need to obtain or create a public certificate to upload to Windows Azure when you register your server or servers online. You can use any valid Secure Sockets Layer (SSL) certificate issued by a CA that is trusted by Microsoft, or you can create your own self-signed certificate by using the Makecert.exe tool.

Remember the following details about the certificate requirements. They could easily form the basis of a test question:

- It must be an x.509 v3 certificate.
- The key length of the certificate must be at least 2048 bits.
- The certificate must have a valid ClientAuthentication EKU.
- The certificate must have a validity period that is less than three years.
- The certificate should be installed in the Personal certificate store of the computer account of server or servers you want to back up. (You can install the same certificate on multiple servers as long as you configure each server to be backed up to the same backup vault in Windows Azure.)
- The certificate uploaded to Windows Azure must be in a .cer format. (This version contains only the public key, not the private key.)

Creating a self-signed certificate with Makecert.exe

To access the Makecert.exe utility, you need to download and install the latest version of the Windows Software Development Kit (SDK). At the time of this writing, the latest version is the Windows SDK for Windows 8.1, which you can find at the following address: <http://msdn.microsoft.com/en-us/windows/desktop/aa904949.aspx>.

To install Makecert.exe, you need only to install one component of the Windows SDK for Windows 8.1: the feature named "Windows Software Development Kit." After you have installed this component, you are ready to use the Makecert utility. To create a self-signed certificate that meets the requirements defined by Windows Azure, open an elevated command prompt and then (assuming the Windows SDK version is 8.1) navigate to C:\Program Files (x86)\Windows Kits\8.1\bin\x64 by typing the following:

```
cd C:\Program Files (x86)\Windows Kits\8.1\bin\x64
```

Then type the following command, replacing *CertificateName* with the desired name of the certificate and *mm/dd/yyyy* with a desired expiry date that is less than three years from the present date:

```
makecert.exe -r -pe -n CN=CertificateName -ss my -sr LocalMachine -eku  
1.3.6.1.5.5.7.3.2 -len 2048 -e mm/dd/yyyy CertificateName.cer
```

For example, the following command installs a certificate named AzureBackup in the Personal certificate store of the local computer account and creates a .cer file named AzureBackup.cer, each with an expiry date of 01/01/2017:

```
makecert.exe -r -pe -n CN=AzureBackup -ss my -sr localmachine -eku  
1.3.6.1.5.5.7.3.2 -len 2048 -e 01/01/2017 AzureBackup.cer
```

NOTE The .cer file created by this running command is stored in the same directory in which the command is run. You might want to move this file to a more convenient location, such as the desktop.

If you want to configure online backups for other servers, make sure you export the new certificate from the Personal certificate store with the private key. The resulting file will be a .pfx file. You can then copy this .pfx file to other servers and install it as needed.

Creating a backup vault in Windows Azure management portal

After you have created the .cer file, the next step in configuring online backups is to create a backup vault in Windows Azure that will use this public certificate. To do so, first navigate to the Windows Azure management portal at <https://manage.windowsazure.com/>. In the navigation menu on the left, select Recovery Services, as shown in Figure 12-1.

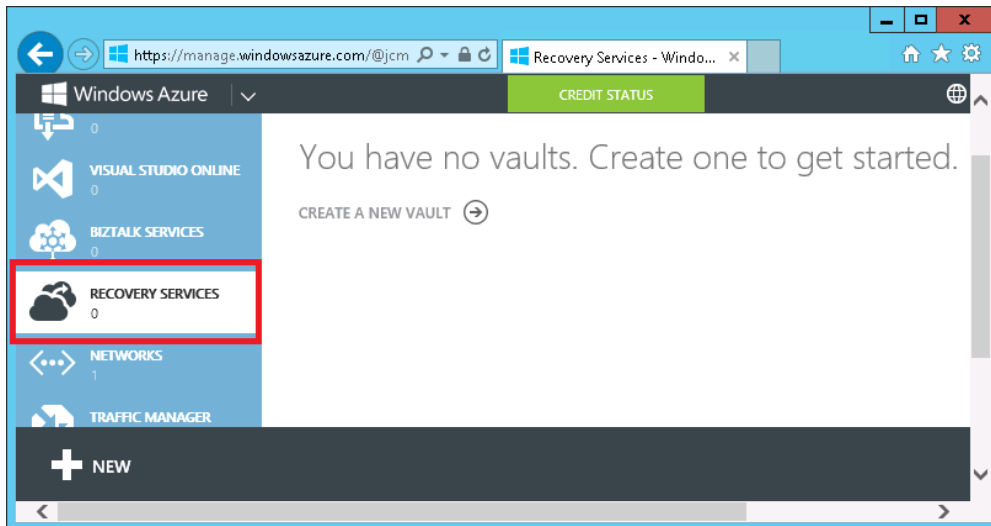


FIGURE 12-1 Creating a new vault in the Windows Azure management portal

Click +NEW and the bottom menu will expand, as shown in Figure 12-2. On the expanded menu, click Backup Vault and then click Quick Create. Enter a name for the vault and then select the most local region. Finally, click Create Vault.

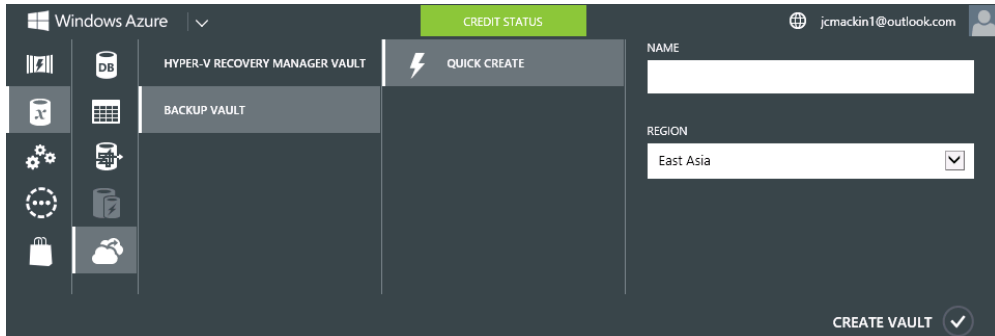


FIGURE 12-2 Creating a new vault in the Windows Azure management portal (continued)

Uploading a certificate to Windows Azure

After the backup vault is created, select the vault in the Windows Azure management portal and then click Manage Certificate. (You can see all of your backup vaults when you click Recovery Services, as shown in Figure 12-1.)

Use the Manage Certificate option to browse to and select the .cer file you either obtained from a CA or created with the Makecert utility.

Downloading and installing the Windows Azure Backup Agent

After you upload the public certificate, you can download the Windows Azure Backup Agent and install it locally. To find the agent in the Windows Azure management portal, double-click on your backup vault and look for a Download Agent link.

A new Backup node then appears in the navigation pane of the Windows Server Backup console, as shown in Figure 12-3. (An earlier version of the Microsoft Azure Backup Agent made this node appear as "Online Backup" instead of just "Backup," so be prepared to see either name on the exam.)

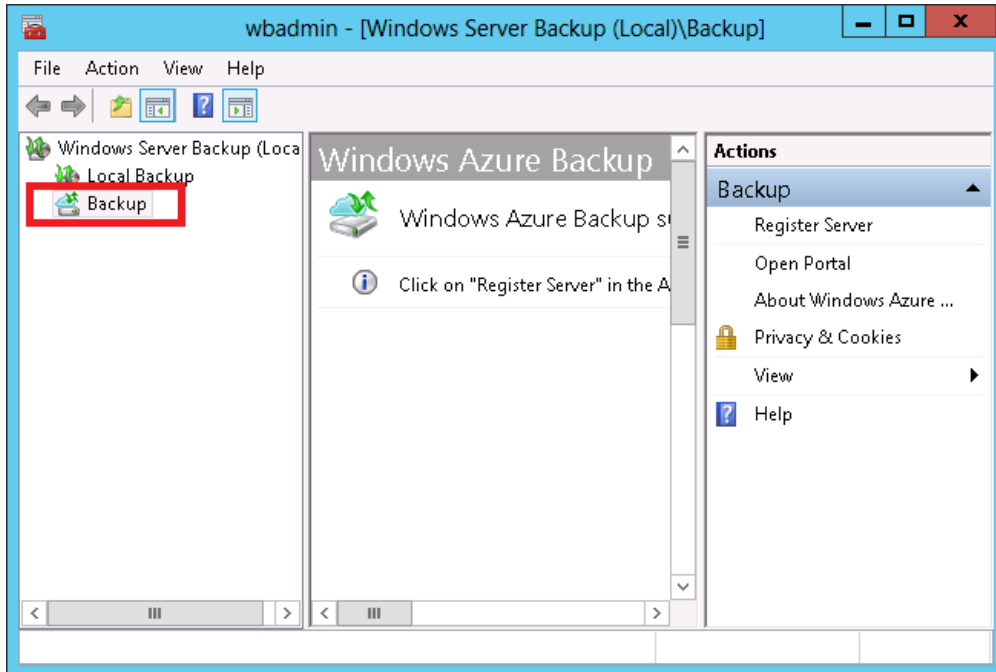


FIGURE 12-3 The Backup node in the Windows Server Backup console

If you prefer, you can also configure online backups from the Windows Azure Backup snap-in, which becomes available through the Microsoft Management Console after you install the agent. The Windows Azure Backup snap-in provides exactly the same set of options as the Backup node in the Windows Server Backup console does.

Registering your server

The next step is to register your server. Registering a server enables you to perform backups from that same server only. (Remember this point for the exam.) To register the server, from the Actions menu, select Register Server.

The Register Server Wizard includes three configuration steps. First, you are given an opportunity to specify a proxy server if desired. Second, you are asked to specify a certificate that you have previously uploaded to a backup vault. (The certificate and backup vault you choose will determine where the backups for the local server will be stored.) Third, you are asked to provide a passphrase that will be used to encrypt your backup data and a location to save this passphrase in a file. You need to provide this passphrase when you perform a restore operation, so it's essential that you don't lose it. (Microsoft doesn't maintain a copy of your passphrase.) A Generate Passphrase option creates the passphrase for you automatically.

After you register a server, new options for Online Backup appear in the Actions pane, including Schedule Backup, Recover Data, Change Properties, and Open Portal.



EXAM TIP

Remember this last sequence of steps: Create an account, create a backup vault, upload a certificate, download and install the agent, and then register the server.

Creating an online backup schedule

Here's an unexpected detail about online backups that could appear on the 70-417 exam: Creating a schedule for your online backup is a requirement. Unlike with local backups in the Windows Server Backup utility, you can't perform a one-time online backup until you have created an automated backup schedule for the same items first.

To start the Schedule Backup Wizard, click Schedule Backup in the Actions pane, as shown in Figure 12-4.

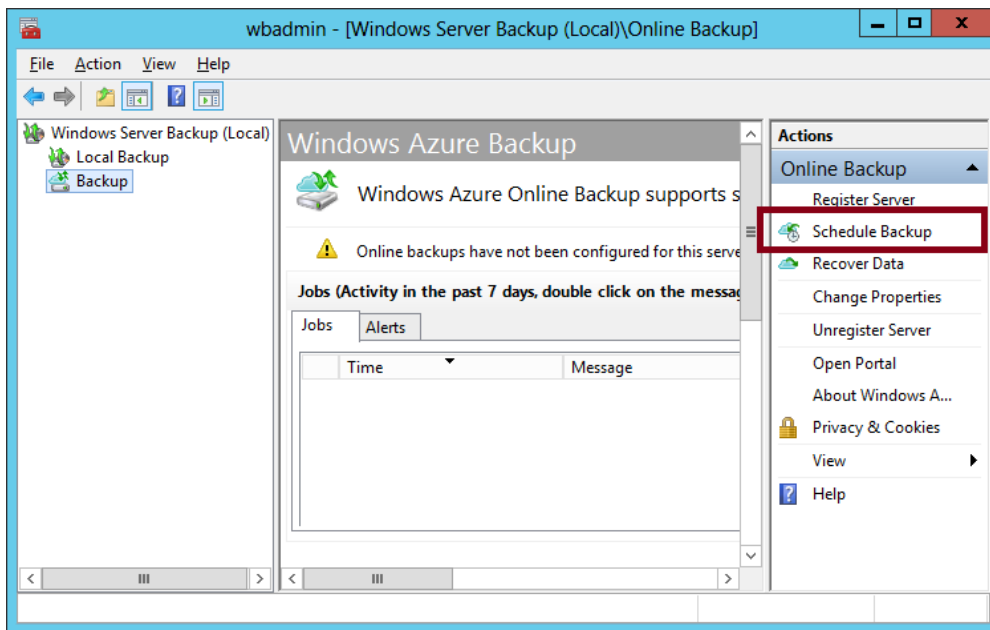


FIGURE 12-4 Scheduling an online backup

SELECTING ITEMS FOR BACKUP

The items you can select to back up in the Schedule Backup Wizard are shown in Figure 12-5. You remember that in Windows Server 2008, you could back up only entire volumes, not folders or files. That's changed now. Beginning in Windows Server 2008 R2 and continuing through to Windows Server 2012 R2, you can now back up selected individual volumes, folders, or files. This improved granularity of backup sets might be difficult for the exam writers to use as the basis for a test question, but you should be aware of it both for the exam and your job in the real world.




FIGURE 12-5 Backup selection for an online backup

It's important to point out a certain limitation related to selecting items in online backup sets, compared to local backup sets. The Select Items dialog box for local backups is shown in Figure 12-6. Compare this to Figure 12-5. Notice the local backup lets you select settings that you can't select for online backups: Bare Metal Recovery, System State, and Hyper-V data (individual VMs or the host component).




FIGURE 12-6 Backup selection for a local (not online) backup



EXAM TIP

Remember that you can't use Windows Azure Backup for Bare Metal Recovery, System State, or Hyper-V data. There is no restriction on individual folders or files.

EXCLUDING ITEMS FROM BACKUP

Beginning in Windows Server 2008 R2 and continuing through to Windows Server 2012 R2, you can exclude files or subfolders from a volume or folder that you have selected for backup. A good example of a folder you might want to exclude from a backup set is a temp folder. When you choose to exclude a folder from your backup set, you are also given an opportunity to exclude its subfolders, as shown in Figure 12-7.

It's possible that you'll see a question on the 70-417 exam that requires some understanding of backup exclusions. Such a question might set up a scenario in which you need to perform a backup more quickly, with less space, or with less network traffic than the current backup set. The "correct answer" might be to exclude a folder with temporary data in the current backup set.

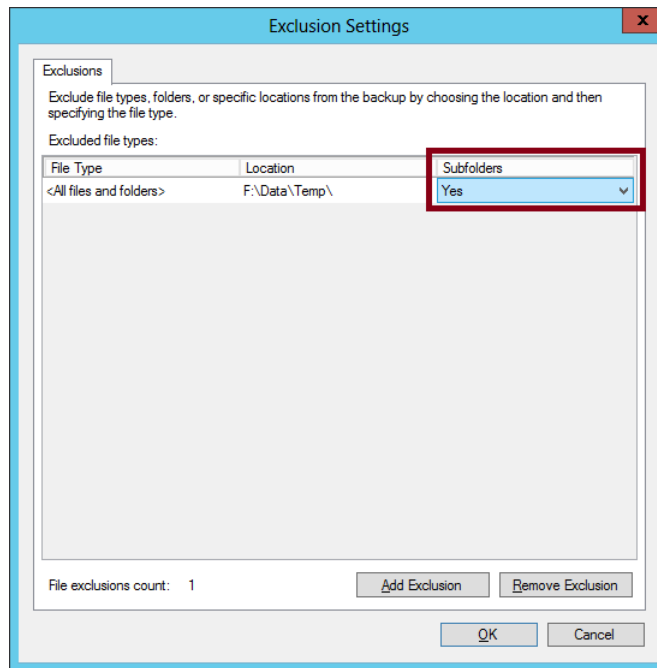


FIGURE 12-7 Excluding a folder and its subfolders from a backup set

SPECIFYING RETENTION SETTINGS

Another feature especially relevant for the exam can be found on the Specify Retention Setting page of the Schedule Backup Wizard, shown in Figure 12-8. The retention setting, also called the retention range, is simply the number of days that the backup cannot be overwritten or deleted to make space for another backup. You can set the retention range for a backup at 7 days (the default), 15 days, or 30 days.

If your Windows Azure Backup account runs out of free space and your retention settings prevent a new backup from overwriting any of the existing backups, the new backup will fail. For example, imagine that the storage quota for your account is 300 GB and you have scheduled a weekly backup job of 200 GB. If you set the retention range of the backup job for 15 days, the backup will fail in the second week. At the default retention setting of 7 days, however, the backup will succeed every week.

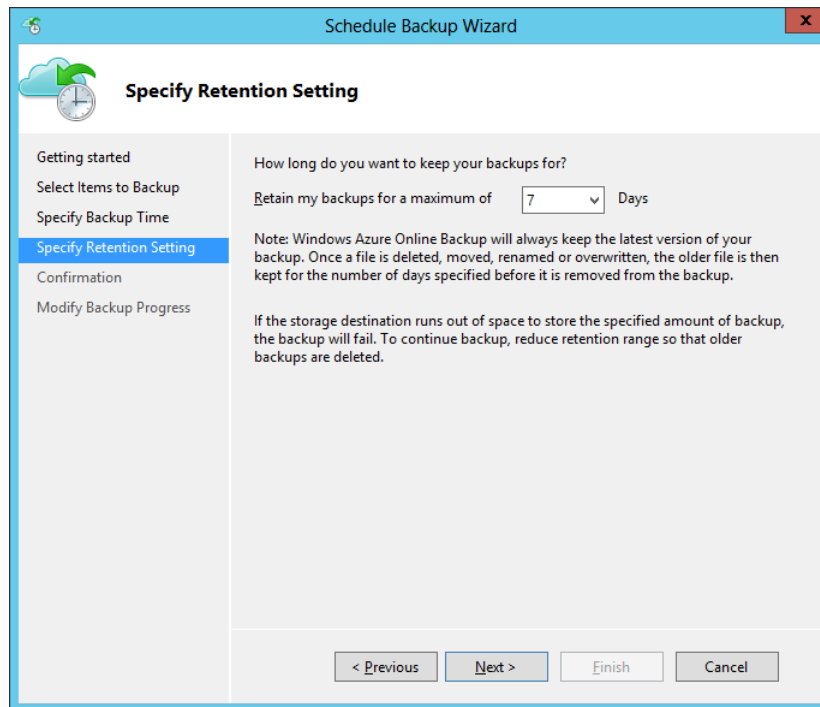


FIGURE 12-8 Backup retention settings

Configuring the Back Up Now option

The Back Up Now option appears in the Actions pane for online backups, as shown in Figure 12-9, but it does so only after you first complete the Schedule Backup Wizard. As stated earlier, Back Up Now for online backups allows you to perform additional online backups only of online backup sets that have been previously defined and scheduled. You *cannot* use this

option to select a new set of volumes, folders, or files and then perform an online backup of that new set.

Aside from this critical difference, the Back Up Now option for online backups resembles the Back Up Once option for local backups.

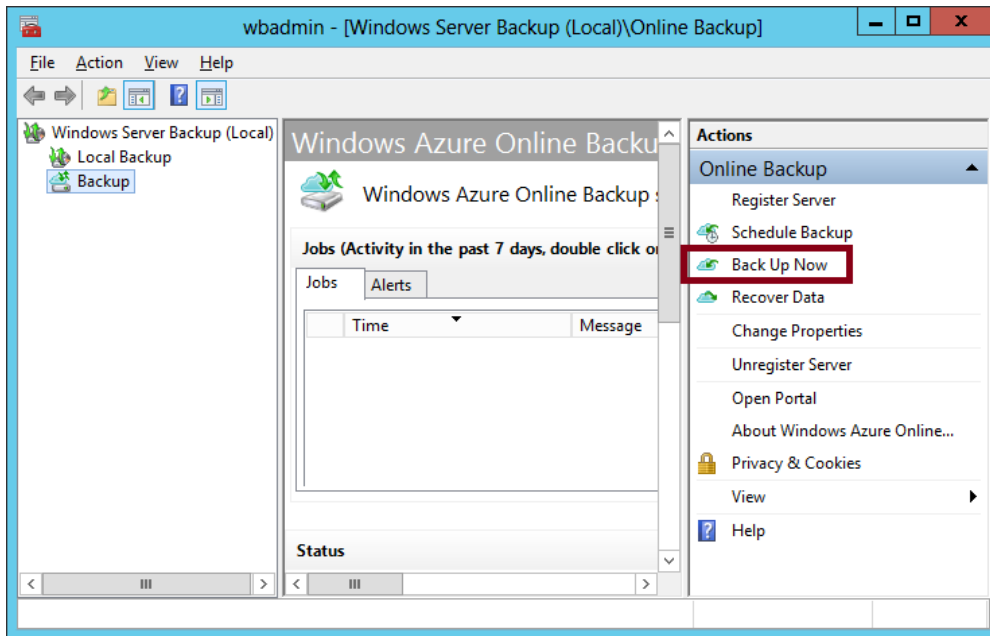


FIGURE 12-9 Performing an additional backup of a scheduled backup

Recovering data

To restore data that has been backed up, choose the Recover Data option in the Actions pane. There isn't anything new or unusual about this option that would likely confuse you in the real world or the exam world. However, it's worth remembering that you can restore online backups to an alternate location such as another volume, a file share, or another server.

Enabling bandwidth throttling

You can restrict the amount of bandwidth used during your online backup operations in a way that depends on when the backup occurs. To enable bandwidth throttling, click Change Properties in the Actions pane, click the Throttling tab, and then select the Enable Internet Bandwidth Usage Throttling For Backup Operations check box, as shown in Figure 12-10.

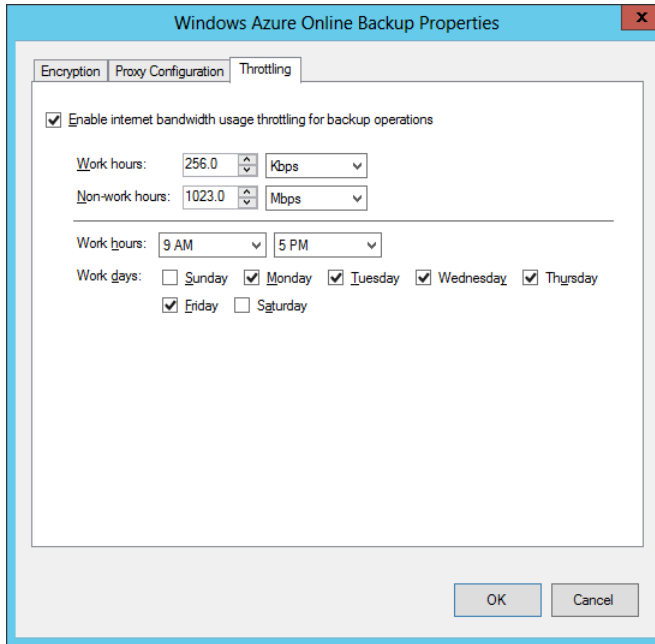


FIGURE 12-10 Configuring bandwidth throttling for online backups

Bandwidth throttling works by letting you set different bandwidth speeds for work and nonwork hours. First you define the hours that should be considered work hours and for which days of the week. You then specify how much Internet bandwidth you want to use for online backup operations during both these work hours and during the remaining nonwork hours.

Bandwidth throttling might be the most likely feature about online backups to appear on the exam. For example, you could see a question that displays the Throttling tab and an accompanying scenario in which you need to adjust the settings in a way that reduces the impact of online backups on your users. In such a case, you might need to redefine the work hours (perhaps by lengthening the work day). Alternatively, you might need to decrease the bandwidth currently assigned to work hours if you want to prevent workday disruption, or increase the bandwidth currently assigned to nonwork hours if you want the online backups to be performed as quickly as possible.



EXAM TIP

If your knowledge is rusty, be sure to review topics related to backing up and restoring that have remained the same since Windows Server 2008. For example, remember that when you enable and configure Shadow Copies settings on a file server, users can use the Previous Versions tab to restore older versions of files, and that you can use the VSSAdmin tool to manage this feature. Remember also the function of the Backup Operators group: It grants users not only the right to perform backups, but also the right to restore and shut down the system.

Performing Windows Azure Backups in Windows PowerShell

After you install the Windows Azure Backup Agent, an additional Windows PowerShell module becomes available: MSOnlineBackup. You can use the cmdlets in this module to configure and perform backups to Windows Azure. Some of these cmdlets are likely to appear on the 70-417 exam, so you need to review them.

To see all 35 available cmdlets in the MSOnlineBackup module, type the following Windows PowerShell command:

```
Get-Command -Module MSOnlineBackup
```

You'll notice that the cmdlets in the MSOnlineBackup module all include the letters "OB," as in Get-OBPolicy, Set-OBMachineSetting, and Start-OBBackup. In case it isn't obvious, "OB" stands for "online backup." That information by itself can already help you answer a question correctly. For example, if you see a question asking you for the cmdlet needed to register a server with Windows Azure Backup, you know you can immediately eliminate any cmdlet that doesn't include "OB."

Although it's a good idea to review all 35 cmdlets, it's essential to emphasize the ones that are most likely to appear on the exam. The most important cmdlets for Windows Azure Backup relate to registering a server, creating an online backup policy, and performing an online backup.

Registering a server with Windows Azure Backup

If you use Windows PowerShell for online backups, you'll typically use PowerShell scripts.

You don't actually need to know PowerShell scripting for the 70-417 exam. You just need to know the key cmdlets that are used to perform each function. For example, the following script might be used to register a server with Windows Azure Backup. The key cmdlet you need to know is in the last line: Start-OBRegistration.

```
$pwd = ConvertTo-SecureString -String <password> -AsPlainText -Force
$cred = New-Object -TypeName System.Management.Automation.PsCredential -ArgumentList
    <username>, $pwd
Start-OBRegistration -Credential $cred
```

After you register the server, you need to use the Set-OBMachineSetting cmdlet to set the encryption passphrase, as in this example:

```
$pass = ConvertTo-SecureString -String <password> -AsPlainText -Force
Set-OBMachineSetting -EncryptionPassphrase $pass
```



EXAM TIP

Remember the Start-OBRegistration and Set-OBMachineSetting cmdlets for the exam.

Creating an online backup policy

You would normally use Windows PowerShell scripting to configure an online backup policy, which is stored in an OBPoicy object. Fortunately, you don't need to know how to configure an online backup policy through Windows PowerShell scripting for the 70-417 exam. However, it helps when you understand that to define an online backup policy, you need to define certain parameters (such as the schedule, file storage locations, and a data retention policy). For example, you could use the following commands to set a new online backup policy for a server:

```
$fspec = New-OBFileSpec -FileSpec C:\test\texttext1.txt
$rpolicy = New-OBRetentionPolicy
$sch = New-OBSchedule
New-OBPolicy | Add-OBFileSpec -FileSpec $fspec | Set-OBRetentionPolicy -RetentionPolicy
    $rpolicy | Set-OBSchedule -Schedule $sch | Set-OBPolicy
```

The most important cmdlets here are New-OBSchedule and New-OBPolicy.

Starting an online backup

After you create an OBPoicy object, you could use the following command to start an online backup immediately by using the values stored in that OBPoicy object:

```
Get-OBPolicy|Start-OBBackup
```



EXAM TIP

Remember the Get-OBPolicy|Start-OBBackup command for the exam.

Table 12-1 includes nine important cmdlets for Windows Azure Backup, along with their official descriptions.

TABLE 12-1 Names and descriptions of important cmdlets for Backup

| cmdlet | Description |
|----------------------|---|
| Get-OBPolicy | Gets the current backup policy set for the server |
| Get-OBSchedule | Gets the OBSchedule object (which includes the days of the week and times of day to create daily backups) for the specified OBPolicy object |
| New-OBPolicy | Generates an empty OBPolicy object |
| New-OBSchedule | Creates a new OBSchedule object based on the days of the week and times of day to create daily backups |
| Set-OBMachineSetting | Sets a OBMachineSetting object for the server with respect to proxy details for accessing the internet, throttling settings, and the encryption passphrase that will be required to decrypt the files during recovery |
| Set-OBPolicy | Sets the OBPolicy object as the backup policy that will be used for scheduled backups |
| Set-OBSchedule | Sets the OBSchedule object (which includes the days of the week and times of day to create daily backups) for the backup policy (OBPolicy object) |
| Start-OBBackup | Starts a one-time backup operation based on the specified OBPolicy |
| Start-OBRegistration | Registers the current computer to the Online Backup Service using the credentials (username and password) created during enrollment |

MORE INFO For more information about the cmdlets in the MSONlineBackup module, see “Windows Azure Online Backup Cmdlets in Windows PowerShell” at <http://technet.microsoft.com/en-us/library/hh770400.aspx>. For in-depth information about how to use Windows PowerShell to configure and implement an online backup policy, see the post named “Microsoft Online Backup Service” by Jeffrey Snover on the Windows Server Blog at <http://blogs.technet.com/b/windowsserver/archive/2012/03/28/microsoft-online-backup-service.aspx>.

Objective summary

- Windows Server 2012 and Windows Server 2012 R2 let you back up selected volumes, folders, and files of the local server over the Internet to cloud storage on Microsoft-owned premises. This functionality is provided by an optional add-on service called Windows Azure Backup, formerly called Windows Azure Online Backup.
- To use Windows Azure Backup, you first need to create an account on the Windows Azure Backup website. Then, create a backup vault, upload a public certificate to the vault, and download and install the Windows Azure Backup Agent to desired local servers.

- After you install the Windows Azure Backup Agent, you can administer online backups in either the Windows Server Backup console or the Microsoft Management Console Windows Azure Backup snap-in. The first step to configuring online backups for a particular server is to register that server online.
- With online backups, you need to create a backup schedule for any backup sets you define. When you run the Schedule Backup Wizard, you select the volumes, folders, and files in the backup, specify any exclusions, set retention settings, and determine the times during the week you want the backup to run.
- Bandwidth throttling is a feature that lets you limit to the amount of Internet bandwidth you want to consume for your online backups. With bandwidth throttling, you define the hours in the week to be considered work hours and then specify the bandwidth in Kbps or Mbps you want online backups to use during these work hours as well as during the remaining nonwork hours.
- You can use Windows PowerShell to configure and perform online backups to Windows Azure, typically through scripting.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the “Answers” section at the end of the chapter.

1. You configure a Hyper-V host running Windows Server 2012 named VHost01 to perform a Windows Azure Backup at 11:00 P.M. every Wednesday. The organization’s Internet connection isn’t used for any other operations until 8:00 A.M. the following day. After running the online backup for the first time, you discover that the backup operation completes at 10:00 A.M. Thursday, after the start of the workday. You open the bandwidth throttling settings for the server and see the configuration shown in Figure 12-11.

You want the online backup of VHost01 to complete before 8:00 A.M. on Thursday. Which of the following solutions is most likely to help you accomplish your goal with the minimum disruption for workers?

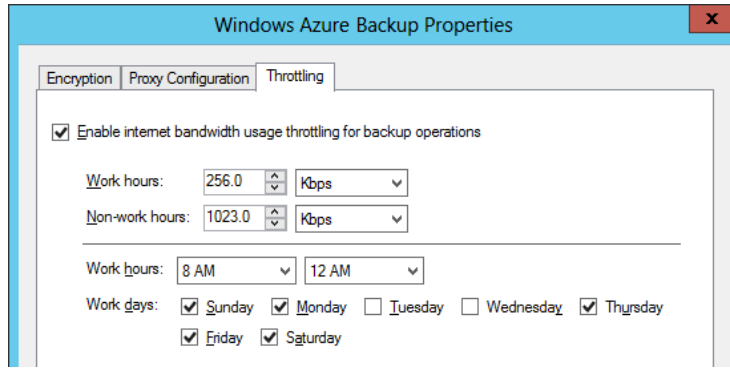


FIGURE 12-11 Bandwidth throttling settings on VHost01

- A. Change the bandwidth setting assigned to work hours.
 - B. Change the bandwidth setting assigned to nonwork hours.
 - C. Change the hours defined as work hours.
 - D. Change the days defined as work days.
2. You have a Windows Azure Backup account with a storage quota of 300 GB. You use this account to configure a single weekly backup of a file server named FileSrv01 that is running Windows Server 2012. The total amount of data on FileSrv01 does not significantly change from week to week. No other backups are configured with your account. The online backup of FileSrv01 completes successfully the first week, but the second week, the backup fails. You receive an error indicating that the usage associated with your Windows Azure Backup account has exceeded its quota. The Windows Azure Backup console displays the information shown in Figure 12-12 about the backup:




FIGURE 12-12 Backup settings and destination usage

You want to be able to perform the weekly backup of FileSrv01 without failure. Which of the following actions is most likely to allow you to accomplish your goal?

- A. Configure an exclusion for C:\Windows\Temp, and choose to exclude its subfolders.
- B. Configure an exclusion for C:\Windows\Temp, and choose not to exclude its subfolders.
- C. Change the retention range to 7 days.
- D. Change the retention range to 30 days.

3. You are a network administrator for a company based in Mumbai, India. You want to configure a local file server named FS02 that is running Windows Server 2012 to perform a daily Windows Azure Backup at 3:00 A.M. You also want to ensure that if the online backup operation extends into the beginning of the next work day at 9:00 A.M., that it will have a minimal impact on network performance for users. The work week in your organization runs from Monday through Friday.

You enable Internet bandwidth usage throttling for backup operations and find the default settings shown in Figure 12-13. What should you do next?

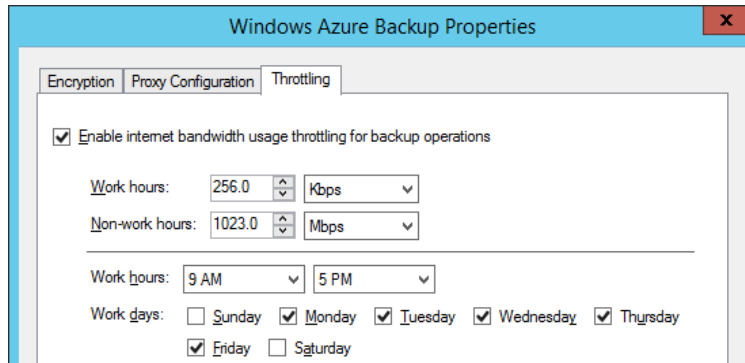


FIGURE 12-13 Bandwidth throttling settings on FS02

- A. Leave the default settings.
- B. Increase the bandwidth setting assigned to work hours.
- C. Increase the bandwidth setting assigned to nonwork hours.
- D. Change the selected work days.

Objective 12.2: Configure site-level fault tolerance

Hyper-V Replica is a new feature in Windows Server 2012 and Windows Server 2012 R2 that provides for a virtual machine (VM) a warm standby copy (or *replica virtual machine*) that can exist anywhere in the world. If the primary VM fails, you can manually fail over to the replica VM. In Windows Server 2012 R2, you can also extend replication to a third VM if desired. Hyper-V Replica in both Windows Server 2012 and Windows Server 2012 R2 can thus provide fault tolerance for a VM even if an entire host site should go offline.

Unlike a failover cluster, Hyper-V Replica doesn't rely on shared storage between the VMs. The replica VM instead begins with its own copy of the primary VM's virtual hard disk. The primary VM then sends updates of its changes (called *replication data*) every few minutes, and this data is repeatedly saved by the replica VM. The replica thus remains up-to-date.

Hyper-V Replica is one of the most important features first introduced in Windows Server 2012, and there's no doubt that it will appear on the 70-417 exam. In fact, you'll probably

see more than one question about it. Fortunately, it's not an especially difficult feature to understand or implement, so your study efforts in this area will likely reap large dividends on the test.

This section covers the following topic:

- Configure Hyper-V replication

Configuring Hyper-V physical host servers

It's important to understand the sequence of steps in configuring Hyper-V Replica. The first step is to configure the server-level replication settings for *both* physical Hyper-V hosts, called the primary server and the replica server. You can access these settings in Hyper-V Manager by right-clicking a host server in the navigation pane, selecting Hyper-V Settings, and then selecting Replication Configuration in the left column of the Hyper-V Settings dialog box, as shown in Figure 12-14. By default, replication is not enabled, and no options are selected or configured.

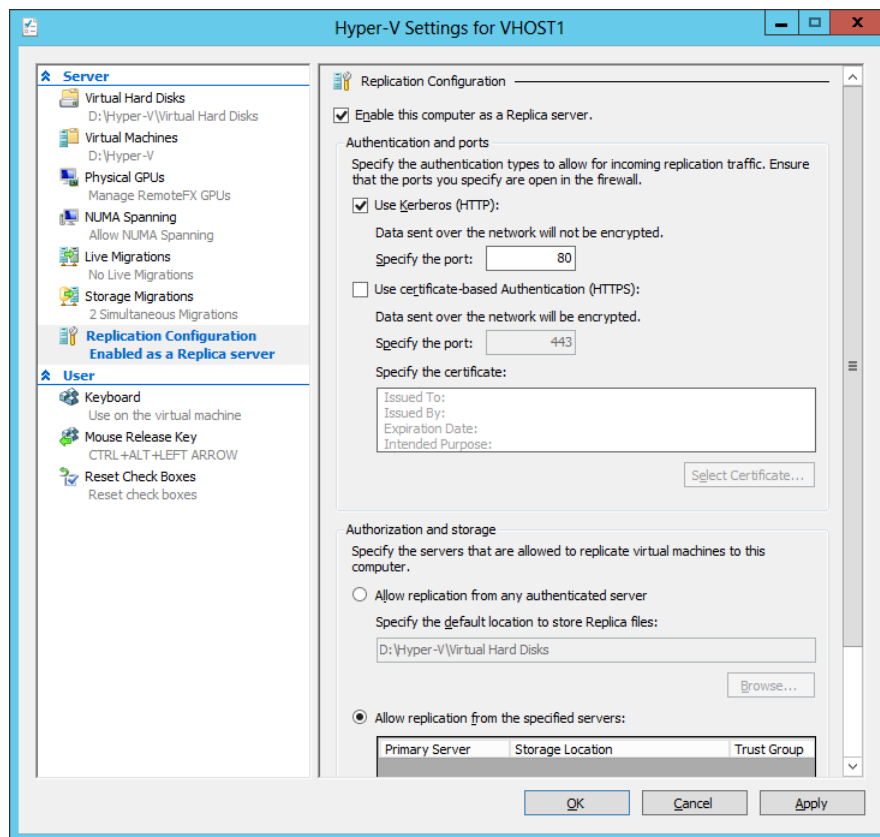


FIGURE 12-14 Host server settings for Hyper-V Replica

To enable a physical host for Hyper-V Replica, first select the Enable This Computer As A Replica Server check box. Then configure the settings in the Authentication And Ports area and the Authorization And Storage area, as shown in Figure 12-14. You need to repeat these configuration steps on both primary and replica servers before configuring a VM for replication.

- **Authentication And Ports** In this area you choose which authentication methods you want to be available later as options when you configure a locally hosted VM for replication. You can enable Kerberos (HTTP), Certificate-Based Authentication (HTTPS), or both.
 - You can enable Kerberos (HTTP) only if the local server is domain-joined. The advantage of choosing Kerberos is that it requires no further configuration. The two disadvantages are first that it doesn't encrypt data sent over the network, and second that it can be used for authentication only when the remote host server is located in a trusted domain. Note also that when you choose this authentication protocol, you need to enable the firewall rule named Hyper-V Replica HTTP Listener (TCP-In).
 - You can enable Certificate-Based Authentication (HTTPS) regardless of whether the local server is domain-joined. In fact, when the local server is a standalone server, it is the only authentication protocol option. The two advantages of enabling Certificate-Based Authentication (HTTPS) are first that it encrypts replication data, and second that it allows you to replicate with a remote host when there is no trust relationship with that host through Active Directory. The disadvantage of this authentication method is that it is more difficult to configure: It requires you to provide an X.509v3 certificate for which Enhanced Key Usage (EKU) must support both Client Authentication and Server Authentication (through the Computer certificate template, for example) and that specifies (typically) the fully qualified domain name (FQDN) of the local server in the subject name field. The certificate can be self-signed or issued through a public key infrastructure (PKI). When you choose this authentication protocol, you need to enable the firewall rule named Hyper-V Replica HTTPS Listener (TCP-In).

It's important to remember that Windows Server 2012 and Windows Server 2012 R2 don't automatically enable the firewall rules you need for the authentication protocols you choose. Depending on which protocol(s) you have enabled, you also need to enable the firewall rule "Hyper-V Replica HTTP Listener (TCP-In)", "Hyper-V Replica HTTPS Listener (TCP-In)", or both. You can enable a rule either in Windows Firewall with Advanced Security or by using the `Enable-NetFirewallRule -DisplayName` command in Windows PowerShell followed by the name of the rule (including quotation marks).



EXAM TIP

Remember that encrypted replication of a VM requires the host servers to have installed a certificate including both Client Authentication and Server Authentication extensions for EKU.

MORE INFO For more information about configuring certificate-based authentication with Hyper-V Replica, search for Hyper-V Replica - Prerequisites for certificate-based deployments or visit <http://blogs.technet.com/b/virtualization/archive/2012/03/13/hyper-v-replica-certificate-requirements.aspx>.

- **Authorization And Storage** This area allows you to configure security settings on the local server that are used when the local server acts as a replica server. More specifically, your choice here determines the remote primary servers from which the local server will accept replication data. Even if you are configuring your local server as the primary server, the settings here are required so that—if you ever need to fail over to a remote replica—you can later fail back to the local server.

You need to choose one of two security options, both of which also provide a default path you can modify to store replication data:

- **Allow Replication From Any Authenticated Server** This option is somewhat less secure. When you choose this option, the local server can receive replication data from any authenticated server.
- **Allow Replication From The Specified Servers** This option requires you to specify the primary server(s) authorized for the local replica server. You can add multiple entries to authorize different primary servers by DNS name. To add an entry authorizing a primary server address, click Add as shown in Figure 12-15. This step opens the Add Authorization Entry dialog box shown in Figure 12-16.

For each entry, a default storage path (the middle field) is already provided, but the other two fields must be filled in manually. In the Specify The Primary Server field, you enter an FQDN that can include a wildcard character (for example, “*.adatum.com”). You also have to provide a tag called a trust group. If you want to allow replication traffic from a set of primary servers, you should assign those primary servers the same trust group name.




FIGURE 12-15 Authorizing primary servers for the local replica server




FIGURE 12-16 Adding an authorized primary server address

How might these settings in the Authorization And Storage area appear on the 70-417 exam? One could imagine a question based on an unsuccessful failover. In such a question, authorization settings might not be configured at all on the replica server. Or the FQDN provided in the Specify The Primary Server field shown in Figure 12-16 might be configured incorrectly; you might be asked to identify the answer that fixes that problem. Another possible question could involve a new organizational requirement that security be tightened on a replica server. Incorrect answer choices might refer to IPSec or other security-tightening methods, but the correct answer will refer to adding an authorization entry on the replica server.

Configuring VMs

After you configure both physical host servers, the next step in configuring Hyper-V Replica is to configure the chosen VM for replication on the primary server. Begin by right-clicking the VM and selecting Enable Replication, as shown in Figure 12-17.

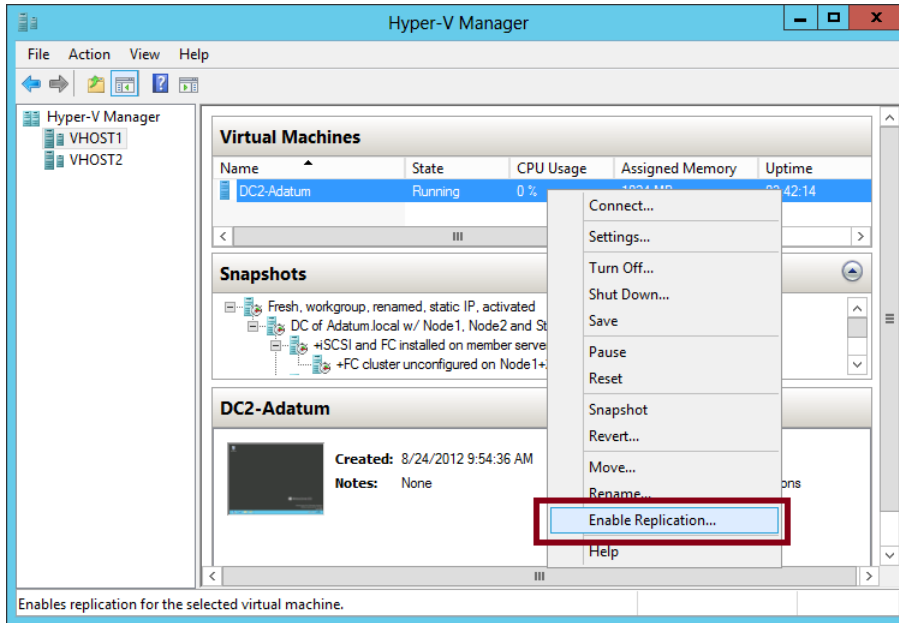


FIGURE 12-17 Creating a replica of a virtual machine

This step opens the Enable Replication Wizard. The wizard includes the following configuration pages:

1. **Specify Replica Server page** Use this page to specify the remote replica server by name.
2. **Specify Connection Parameters page** This page, shown in Figure 12-18, asks you to specify which of the authentication types enabled at the server level in Hyper-V Settings you want to use to support this replicated VM. If you have enabled only one of these two authentication methods at the server level, that same method is the only option here. Naturally, the replica server must support the same authentication method.

This page also provides an option that lends itself fairly well to an exam question: the Compress The Data That Is Transmitted Over The Network check box. This compression option reduces bandwidth requirements for replication at the expense of increased processor usage. If this option does appear on the exam, this trade-off is likely to be the key to getting the right answer.



EXAM TIP

If both authentication types are available for the VM and you want to change the authentication type later, you have to remove replication and complete the Enable Replication Wizard again. Before you do, though, make sure that certificate-based authentication is also enabled in the Hyper-V Settings on the remote host server.

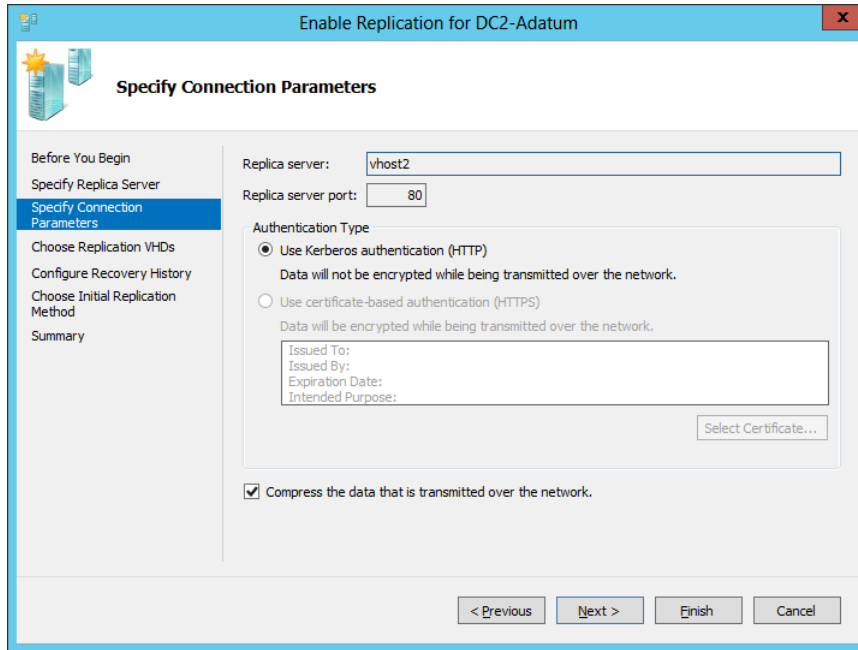


FIGURE 12-18 Selecting authentication and compression settings for a replicated VM

3. **Choose Replication VHDs page** By default, all virtual hard disks (VHDs) attached to the VM are enabled for replication. You can use this page to deselect any VMs that you don't want to be replicated.
4. **Configure Replication Frequency page** (Windows Server 2012 R2 only) In the first release of Windows Server 2012, changes were normally sent from the primary server to the replica server every 5 minutes, an interval that was not configurable. In Windows Server 2012 R2, as shown in Figure 12-19, you can now choose a replication frequency of 30 seconds, 5 minutes, or 15 minutes. With more frequent replication, you can reduce the amount of data that can be lost in case the replicated VM goes down, but at the expense of increased strain on resources (especially bandwidth). With less frequent replication, resources are not used as much, but more data can be lost. (Because this feature is new in Windows Server 2012 R2, you should expect to see an exam question about it. Make sure that you understand the tradeoffs involved when you choose more or less frequent replication.)

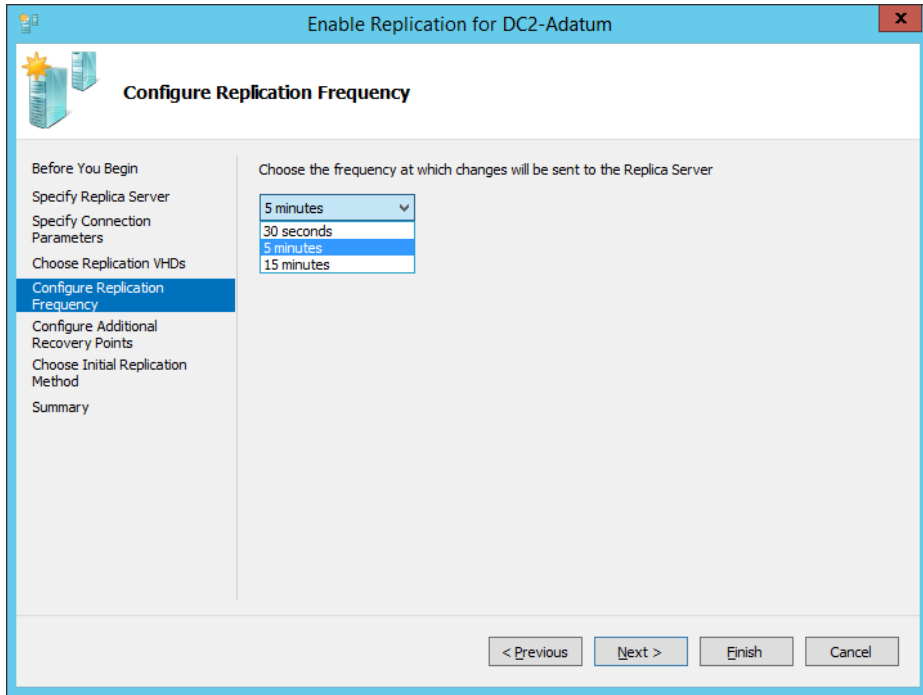


FIGURE 12-19 Configuring replication frequency

- 5. Configure Additional Recovery Points page** This page, shown in Figure 12-20, includes the settings to configure recovery points. By default, the Maintain Only The Latest Recovery Point option is selected, and no other options are enabled or configured.

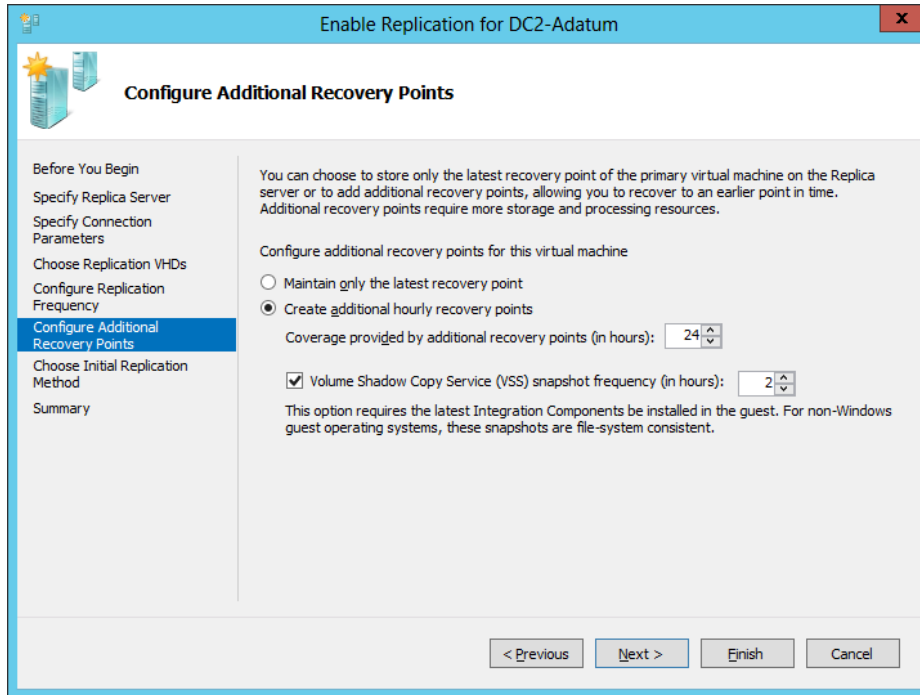


FIGURE 12-20 Configuring additional recovery points.

Recovery points are VM snapshots saved on a replica server. New snapshots are sent from the primary to the replica server according to the replication frequency you've configured on the previous page, but only the latest snapshot is saved on the replica by default. Selecting the Create Additional Hourly Recovery Points option configures the replica server to keep one extra snapshot per hour, up to the amount of coverage specified in the next configuration setting, "Coverage Provided By Additional Recovery Points (In Hours)." If you later perform a failover operation at the replica server, you then have the option of recovering either the most recent version of the VM, which is always available, or one of these earlier, hourly snapshots.

A menu of available recovery points on a replica server is shown in Figure 12-21. If the Configure Additional Recovery Points page were left at the default setting (Maintain Only The Latest Recovery Point), only the first option named Latest Recovery Point would appear in this menu.

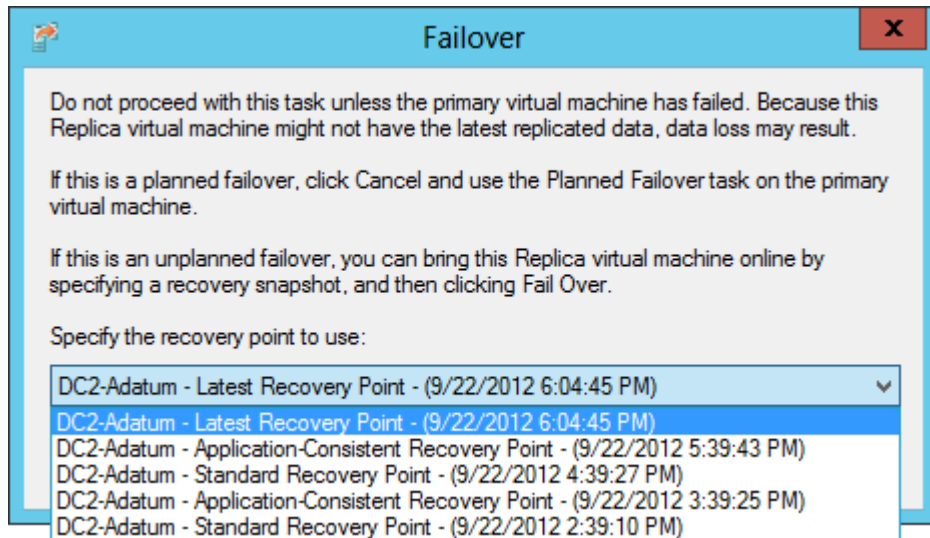


FIGURE 12-21 The latest recovery point and previous hourly snapshots of a VM that can be restored in a failover on the replica server

When you enable the Create Additional Hourly Recovery Points option on the Configure Additional Recovery Points page, the replica server by default will keep an hourly snapshot for each of the past 24 hours in addition to the latest recovery point. However, you can change this setting if you want to store fewer of these recovery points on the replica server. The main drawback to keeping many recovery points is the use of storage resources required to do so.

The last configuration settings on the Configure Additional Recovery Points page relate to incremental Volume Shadow Copy Service (VSS) copies, also known as *application-consistent recovery points*. These are high-quality snapshots taken during moments in which the VM momentarily “quiesces” (gracefully pauses) activity in VSS-aware applications such as Microsoft Exchange and SQL Server. The advantage of these snapshot types is that they help ensure that the failover will be free of errors in these applications. The disadvantage is that they are more processor-intensive and cause important applications to pause briefly. (However, it should be noted that the pause is normally too brief for users to detect.)

You enable incremental VSS copies by selecting the Volume Shadow Copy Service (VSS) Snapshot Frequency (In Hours) check box, and then selecting the frequency of the application-consistent recovery point. (You can see these options in Figure 12-20.) If you leave the default frequency of 4 hours, then every fourth recovery point will be an application-consistent recovery point. If you select a frequency of 2 hours, then the standard recovery point will be replaced by an application-consistent recovery point every 2 hours, and so on. Figure 12-22 shows the snapshots stored on a replica server for which incremental VSS copies are scheduled every two hours.

Two final points to note about VSS snapshots: They require Integration Components to be installed in the guest VM, and they are not possible with non-Windows operating systems (such as Linux).

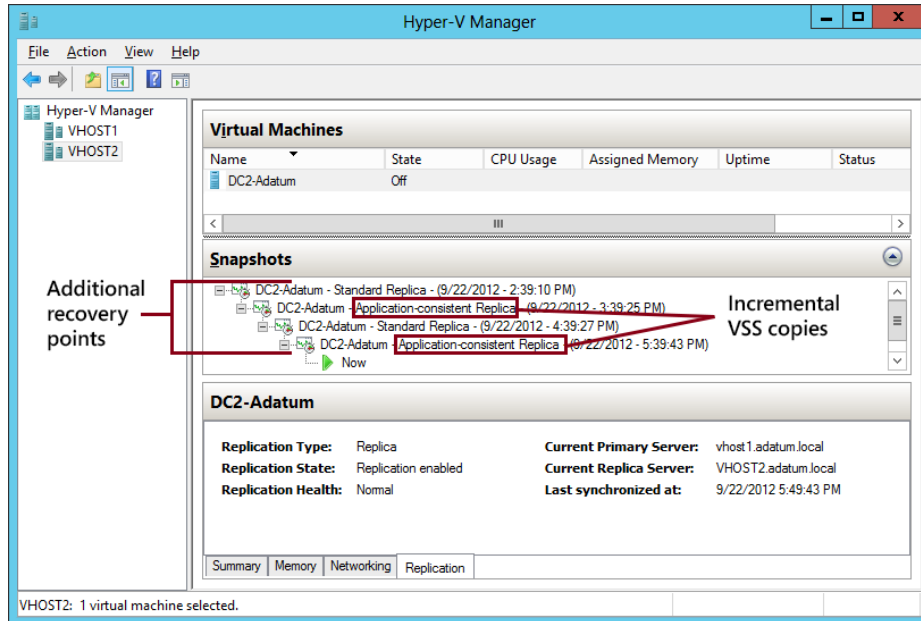


FIGURE 12-22 Incremental VSS copies and standard recovery points



EXAM TIP

Expect to see a question about VSS snapshots on the 70-417 exam.

- 6. Choose Initial Replication Method page** This page, shown in Figure 12-23, allows you to specify how the initial copy of the VHDs attached to the primary VM will be sent to the replica server. By default, the VHDs are sent over the network. Sending very large files over a network such as the Internet isn't always a realistic option, however. As an alternative, you can choose the second option, to export the VHDs to external media (and then physically transport them to the replica server). The final option is to use an existing VM on the replica server as the initial copy. You can choose this option if you have restored an exact copy of the VM and its VHDs on the replica server. This page also allows you to configure the initial network transfer to take place at a specified future time. You can use this option to minimize user disruption.

NOTE Typically, the initial transfer of the VHD is far more bandwidth-intensive than the updates sent through replication are. After the initial copies of the VHDs are sent, only the changes (deltas) to these VHDs are sent during replication according to the frequency you have specified.

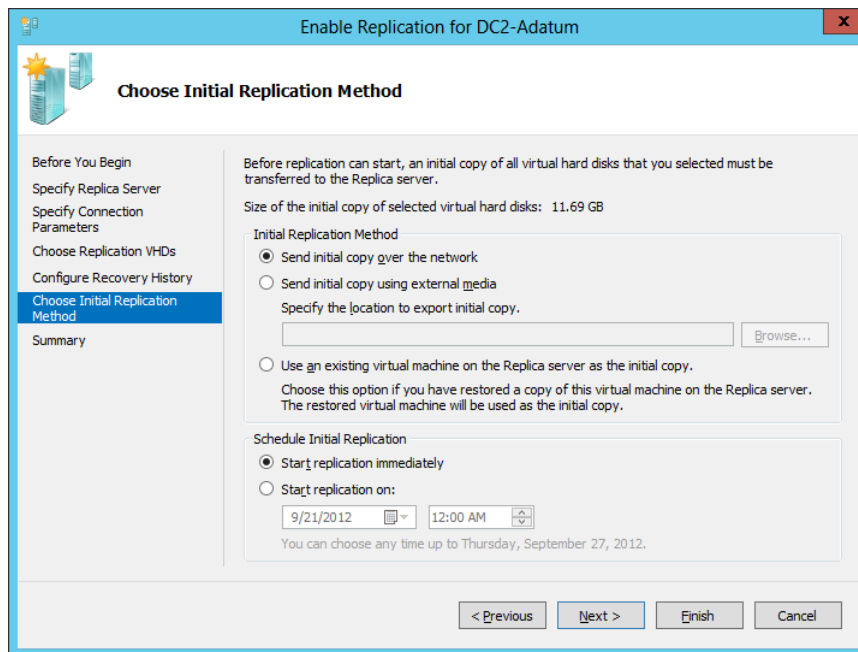


FIGURE 12-23 Determining how to send the base copy of the VHDs attached to a primary VM

Configuring failover TCP/IP settings

After you enable replication on a VM, you might need to specify the TCP/IP settings that will apply to the replica VM after failover. By default, the replica VM will inherit the same IPv4 and IPv6 configuration as the primary VM. In many cases, however, the replica VM will need a different IP configuration to communicate in its environment.

To assign a different IP configuration to the replica VM, in Hyper-V Manager on the replica server, right-click the replica VM and select Settings from the shortcut menu. In the Settings dialog box, expand Network Adapter in the left column and then select Failover TCP/IP, as shown in Figure 12-24. In the right pane, assign the new IP configuration as appropriate.

Then, on the primary server, assign the original IP configuration in the same settings area. Otherwise, the replica settings will persist if you fail back to the original location. (Remember this last point for the exam.)

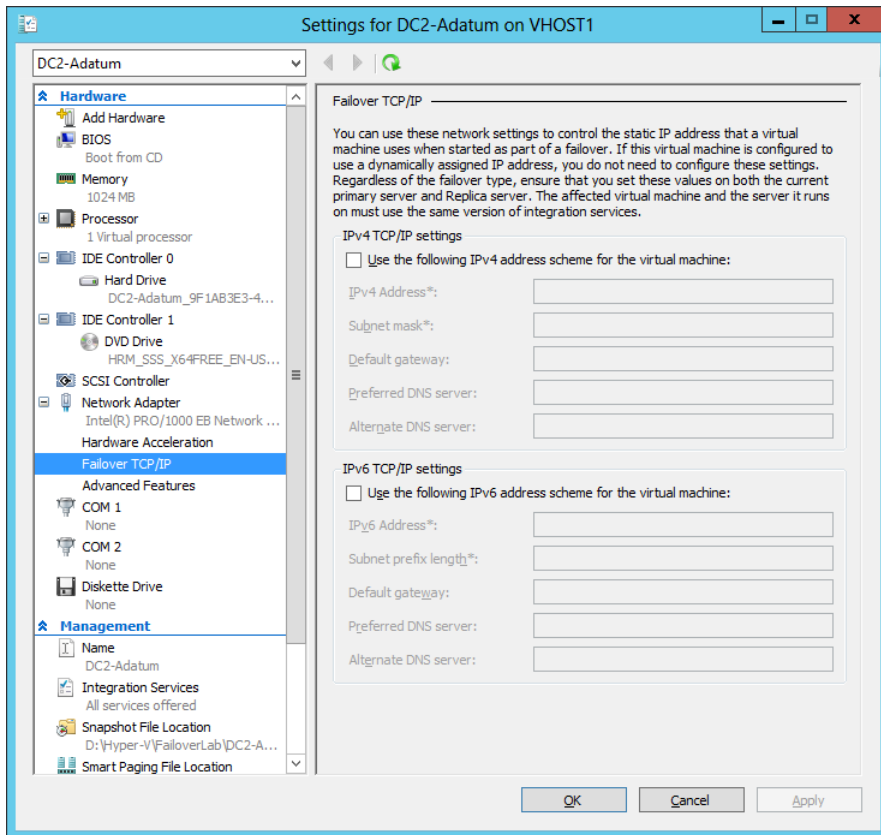


FIGURE 12-24 Assigning a different IP configuration to a replica VM

Resynchronizing the primary and replica VMs

After you complete the Enable Replication Wizard, you can modify the replication settings for a VM in the Settings dialog box for that VM. Replication settings appear in the Management category in the menu on the left, as shown in Figure 12-25.

One configuration setting appears here that does not appear in the Enable Replication Wizard: Resynchronization. Resynchronization is a highly resource-intensive operation that is performed occasionally between a primary and replica VM. By default, resynchronization can occur at any time. You have the option, however, to restrict resynchronizations to selected off-peak hours. Alternatively, you can opt to perform resynchronization manually.

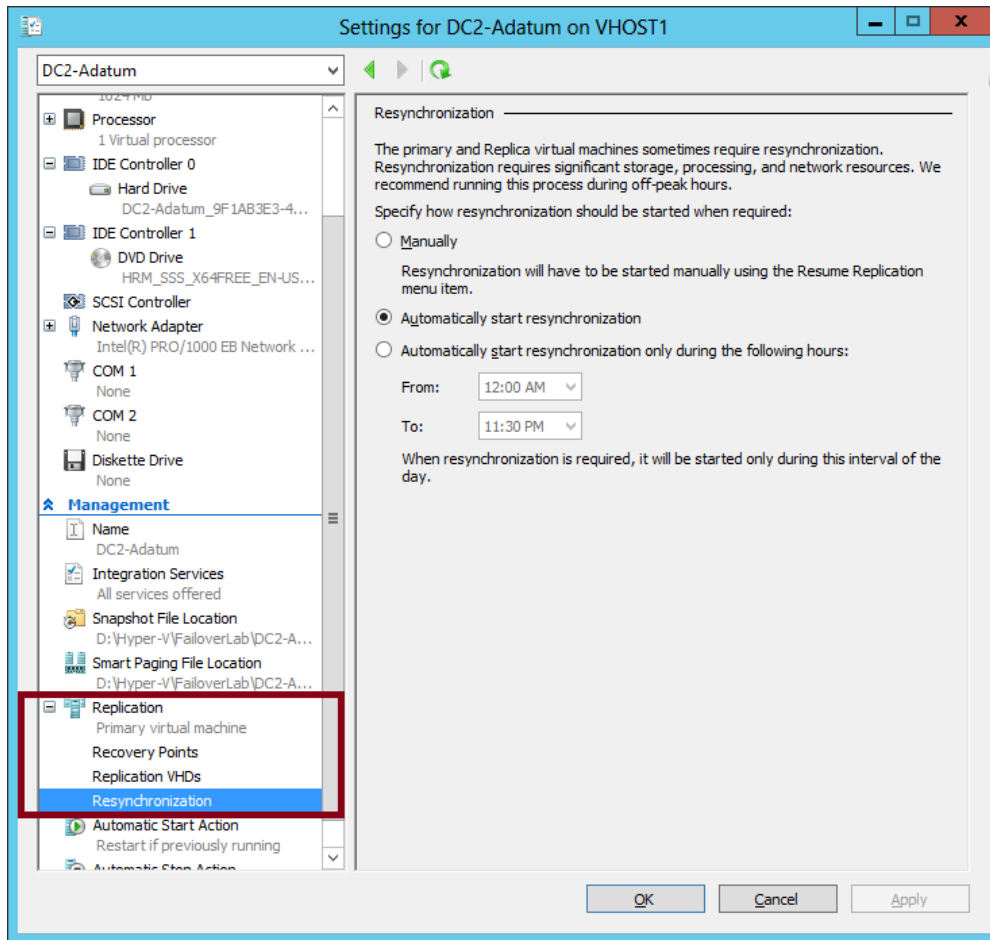


FIGURE 12-25 Replication settings for a VM

Performing Hyper-V Replica failover

You can perform three types of failovers with Hyper-V Replica after it is configured: planned failovers, unplanned failovers, and test failovers. It's likely you'll see an exam question in which you need to understand the difference among them and when they are used.

- **Planned failover** A planned failover is the only failover you initiate from the primary server. You use this method whenever you can manually shut down the primary VM, and the primary and replica servers can still communicate.

A planned failover is the preferred failover type because no data is lost. In fact, you cannot even use this option to fail over to the latest recovery point or to any earlier recovery point. With a planned failover, only an exact copy of the current primary VM and its VHDs can be failed over to the replica server.

A planned failover is a good option in the following situations:

- You want to perform host maintenance on the primary server and temporarily want to run the VM from the replica.
- Your primary site is anticipating a possible power outage and you want to move the VM to the replica site.
- You are expecting a weather emergency such as a flood and you want to ensure business continuity.
- Your compliance requirements mandate that you regularly run your workloads for certain periods of time from the replica site.

To perform a planned failover, you begin by *shutting down the primary VM*. You then right-click the VM in Hyper-V Manager, click Replication, and then click Planned Failover, as shown in Figure 12-26. The latest updates are then sent to the replica server, the VM is failed over, and the replica VM is automatically started on the remote server. At the end of this operation, the replication relationship is reversed, so what was the replica server becomes the primary server, and vice versa.

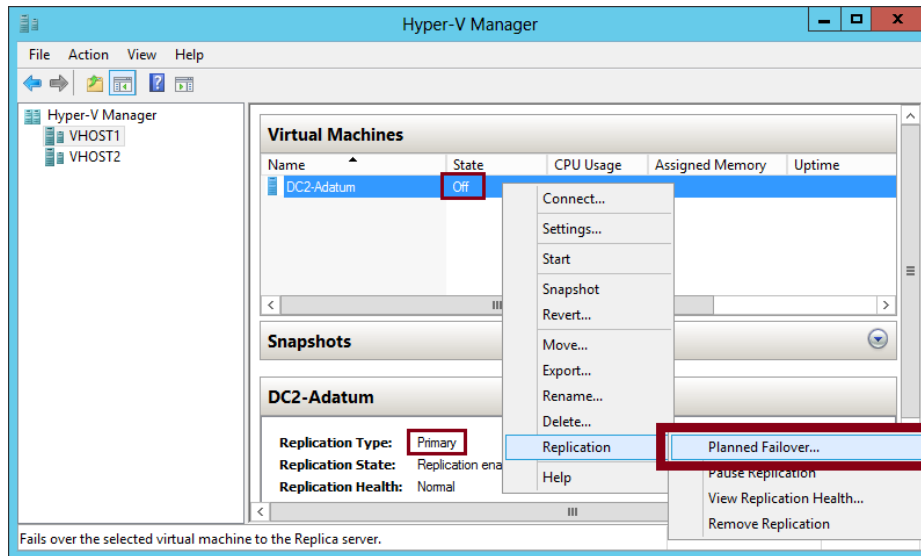


FIGURE 12-26 Performing a planned failover from the primary server

- **(Unplanned) failover** This type of failover is called an unplanned failover in documentation about the feature, but in the actual interface, it's called just "failover." On the 70-417 exam, you might see it referred to either way.

An unplanned failover is performed at the replica server. You perform this failover type when the primary VM fails suddenly and cannot be brought back online. An unplanned failover is a good option in the following situations:

- Your primary site experiences an unexpected power outage or a natural disaster.

- Your primary site or VM has had a virus attack and you want to restore your business quickly with minimal data loss by restoring your replica VM to the most recent recovery point before the attack.

To perform an unplanned failover, in Hyper-V Manager on the replica server, right-click the replica VM, click Replication, and then click Failover, as shown in Figure 12-27.

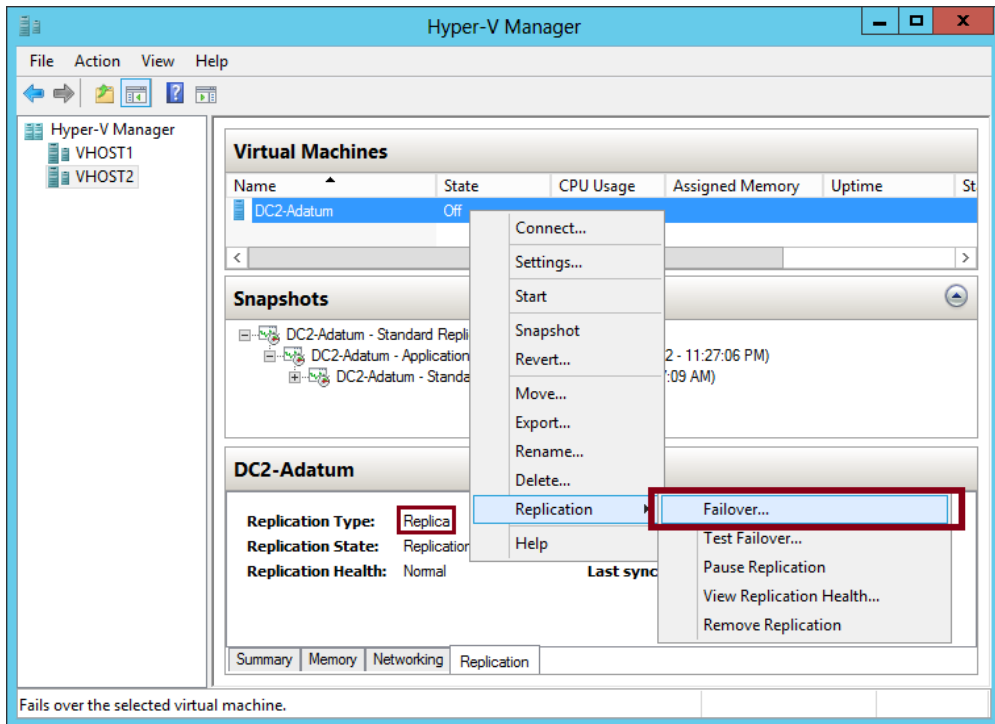


FIGURE 12-27 Performing an unplanned failover on the replica server

When you perform an unplanned failover, you have to choose a recovery point, as shown earlier in Figure 12-21. The VM is then started on the replica server.

After the replica VM is started, the replica relationship with the primary VM is broken and replication stops. If at some later point you can bring the original primary VM online, you can resume replication by reversing the replication relationship. After you perform this operation, the local replica server becomes the new primary and the remote primary becomes the new replica. To reverse replication in this way, right-click the VM on the replica server, click Replication, and then click Reverse Replication, as shown in Figure 12-28. This step starts the Reverse Replication Wizard, which allows you to reenter the settings for the replica.

Another option you can see on the Replication submenu in Figure 12-28 is Cancel Failover. You can safely choose this option after you perform an unplanned failover as long as no changes have been made to the replica. After you cancel a failover, you

have to manually resume replication on the primary VM by right-clicking it and selecting Resume Replication. Cancelling a failover is a good idea if you quickly discover after performing an unplanned failover that the primary VM can be brought online.



EXAM TIP

Remember the Reverse Replication option and the Cancel Replication option for the exam.

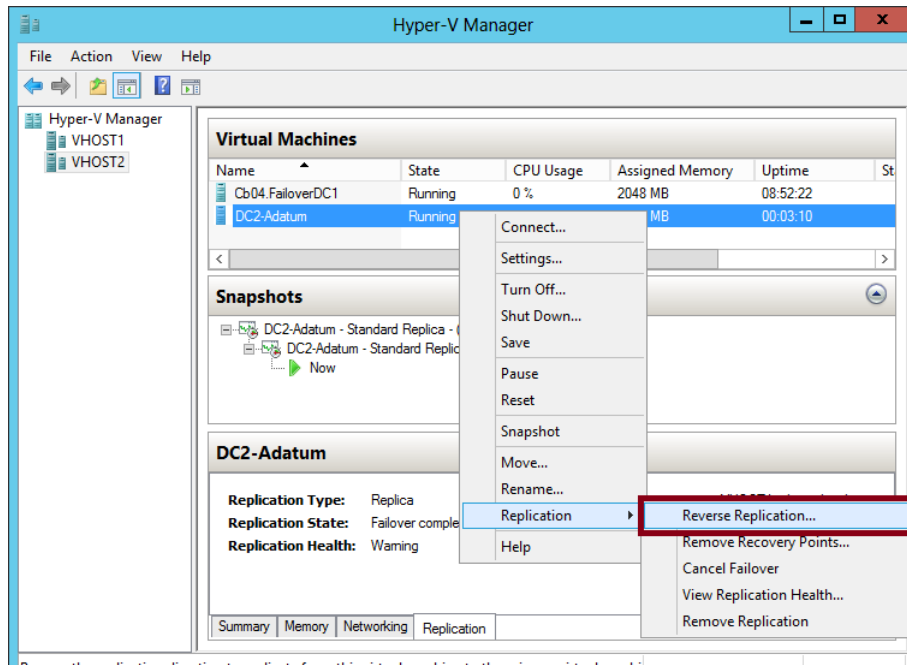


FIGURE 12-28 Reversing replication

- **Test failover** A test failover is the only failover operation you can perform while the primary VM is still running. The purpose of this failover type is to simulate an unplanned failover so that you can ensure that it will function as planned in case of an emergency.

To perform a test failover, in Hyper-V Manager on the replica server, right-click the replica VM, click Replication, and then click Test Failover. You then have to select a recovery point, just as you do with an unplanned failover. Next, a local, disposable copy of the replica VM is created on the replica server. The new copy of the VM appears in Hyper-V Manager in a stopped state with the tag “- Test.” For example, a test failover of a VM named “MyVM1” would result in a new VM called “MyVM1 - Test.” You can then start the new VM manually to see if it works as expected.

By default, the virtual network adapters of the test VM are disconnected from all virtual switches. If desired, you can preattach the adapter(s) of the test VM to a virtual

switch of your choice. To do so, open the settings of the base replica VM, expand Network Adapter, and then click Test Failover, as shown in Figure 12-29. Make sure you choose a virtual switch that will not create any conflicts in a production network. After you examine the functioning of the test VM, you can safely delete it in Hyper-V Manager.

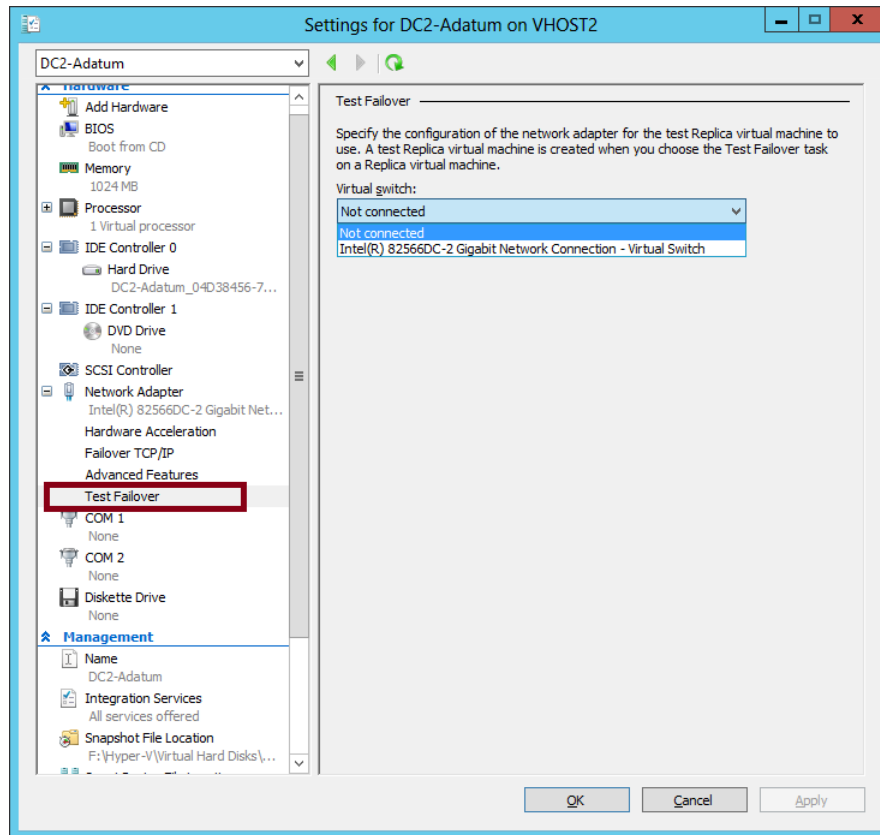


FIGURE 12-29 Preattaching the network adapter of a failover test VM to a virtual switch



EXAM TIP

Remember for the exam when you should perform a planned failover, a failover, and a test failover.

Extending replication to a third site in Windows Server 2012 R2

Windows Server 2012 R2 introduces the ability to perform replication of a VM to a third site. The most important thing to understand is that this option extends replication from the replica server only. In other words, you cannot configure the primary VM to be replicated directly to two different sites. Instead, with extended replication, a first site replicates to a second site and the second site replicates to a third site. This concept is illustrated in Figure 12-30.




FIGURE 12-30 Windows Server 2012 R2 allows you to extend replication from the replica server to a third server.

Configuring extended replication is easy. You simply right-click the replica VM, select Replication, and then select Extend Replication, as shown in Figure 12-31. This step opens the Extend Replication Wizard, which provides the same set of configuration options as does the Enable Replication Wizard described earlier in this chapter.

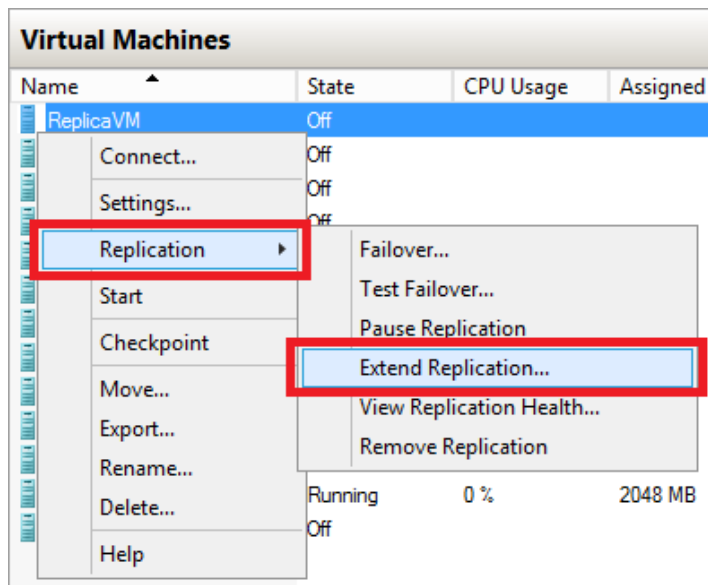


FIGURE 12-31 Extending replication from a replica VM

Using Hyper-V Replica in a failover cluster

The configuration steps previously described apply to VMs that are not hosted in a failover cluster. However, you might want to provide an offsite replica VM for a clustered VM. In this scenario, you would provide two levels of fault tolerance. The failover cluster is used to provide local fault tolerance, for example, if a physical node fails within a functioning data center. The offsite replica VM, on the other hand, could be used to recover only from site-level failures, for example, in case of a power outage, weather emergency, or natural disaster.

The steps to configure a replica VM for a clustered VM differ slightly from the normal configuration, but they aren't complicated. The first difference is that you begin by opening Failover Cluster Manager, not Hyper-V Manager. In Failover Cluster Manager, you then have to add a failover cluster role named *Hyper-V Replica Broker* to the cluster. (Remember, the word "role" is now used to describe a hosted service in a failover cluster.)

To add the Hyper-V Replica Broker role, right-click the Roles node in Failover Cluster Manager and select *Configure Role*. This step opens the High Availability Wizard. In the High Availability Wizard, select *Hyper-V Replica Broker*, as shown in Figure 12-32.

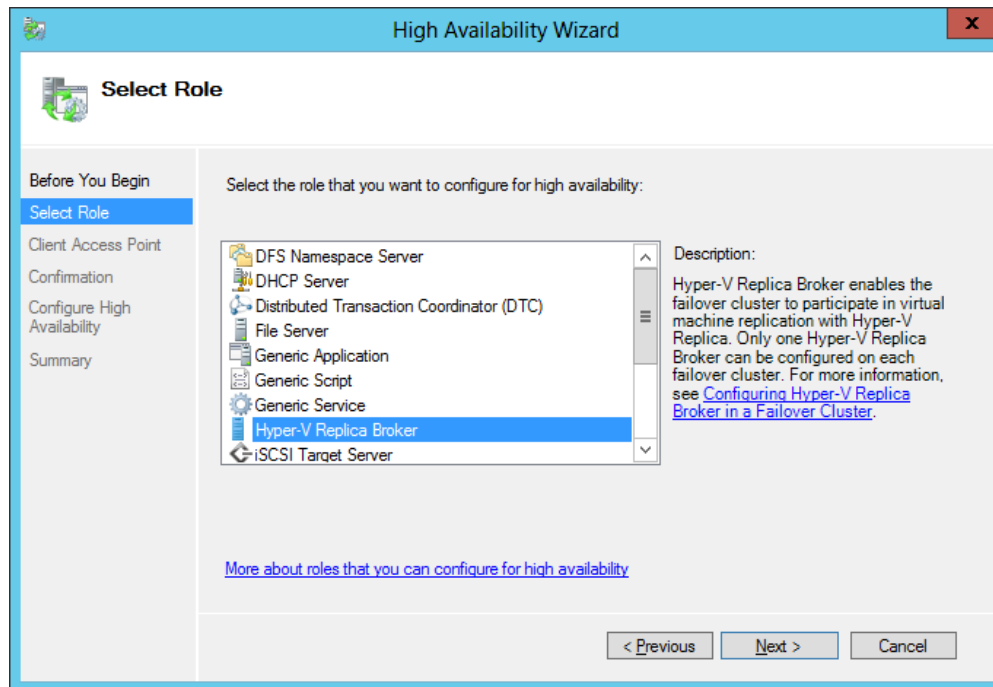


FIGURE 12-32 Adding the Hyper-V Replica Broker role to a failover cluster

When you choose this role, the High Availability Wizard will ask you to provide a NetBIOS name and IP address to be used as the connection point to the cluster (called a *client access point*, or *CAP*). This step is shown in Figure 12-33.

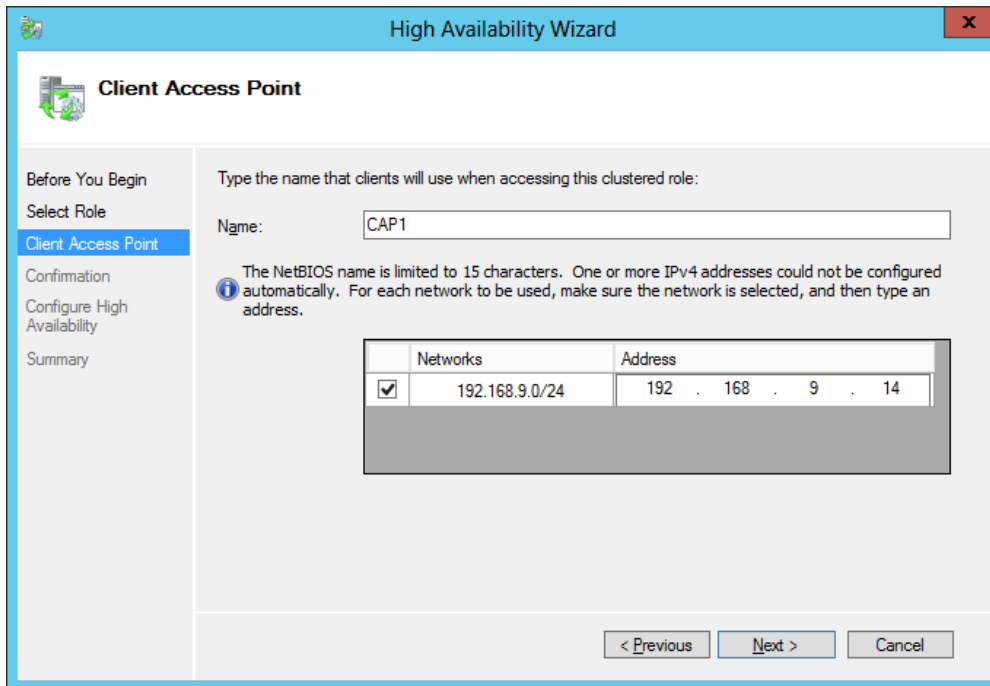


FIGURE 12-33 Providing a name and address for the client access point

Next, you configure the equivalent of the server replication settings shown earlier in Figure 12-14. To do so, right-click the Hyper-V Replica Broker node in Failover Cluster Manager and select Replication Settings from the shortcut menu, as shown in Figure 12-34. The difference between the settings here and the settings in Figure 12-14 is that in this case, the settings apply to the entire cluster as a whole.

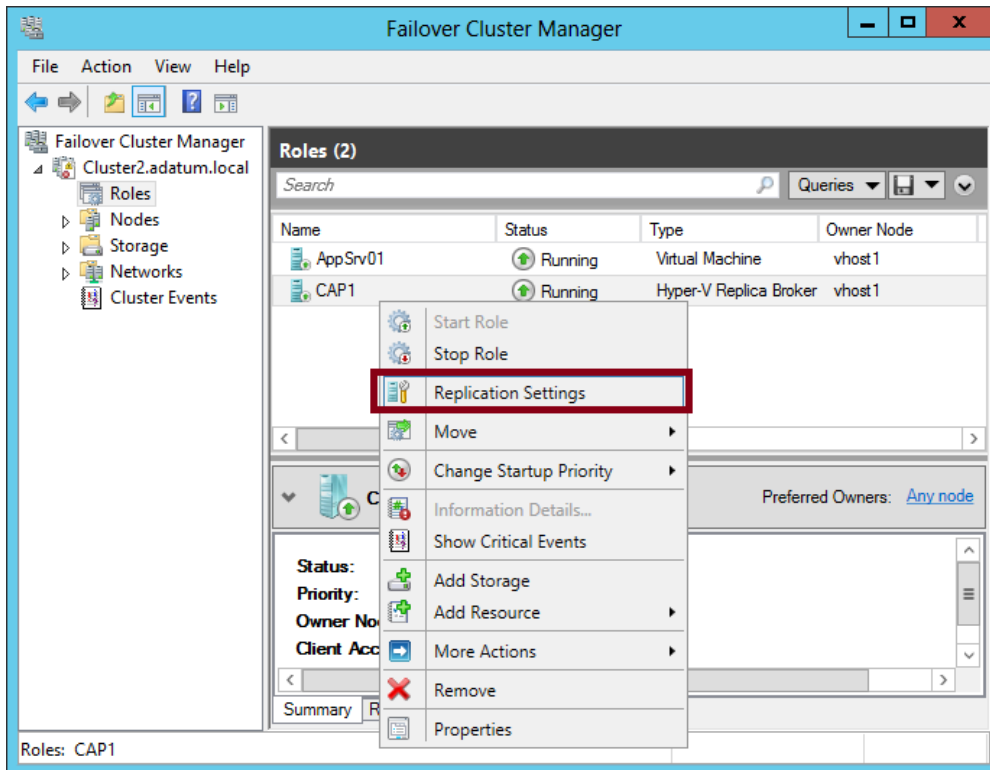


FIGURE 12-34 Configuring replication settings for the cluster

On the remote Replica server, you configure replication as you normally would by configuring Hyper-V Settings in Hyper-V Manager as described in the earlier section named “Configuring Hyper-V physical host servers.” However, if you want the remote Replica also to be a multi-node failover cluster, then you would need to configure that remote failover cluster through Failover Cluster Manager (by adding and configuring the Hyper-V Replica Broker role).

After you configure the host server settings, you can configure replication on the VM in Failover Cluster Manager just as you would in Hyper-V Manager. Right-click the clustered VM, click Replication, and then click Enable Replication, as shown in Figure 12-35.

This step opens the same Enable Replication Wizard that you see when you configure replication on a nonclustered VM. The remaining configuration steps are therefore identical. To perform failover to the replica server of a clustered VM, use the same options on the shortcut menu of the VM as you would on a nonclustered VM. The only difference, again, is that you perform the operation in Failover Cluster Manager as opposed to Hyper-V Manager.

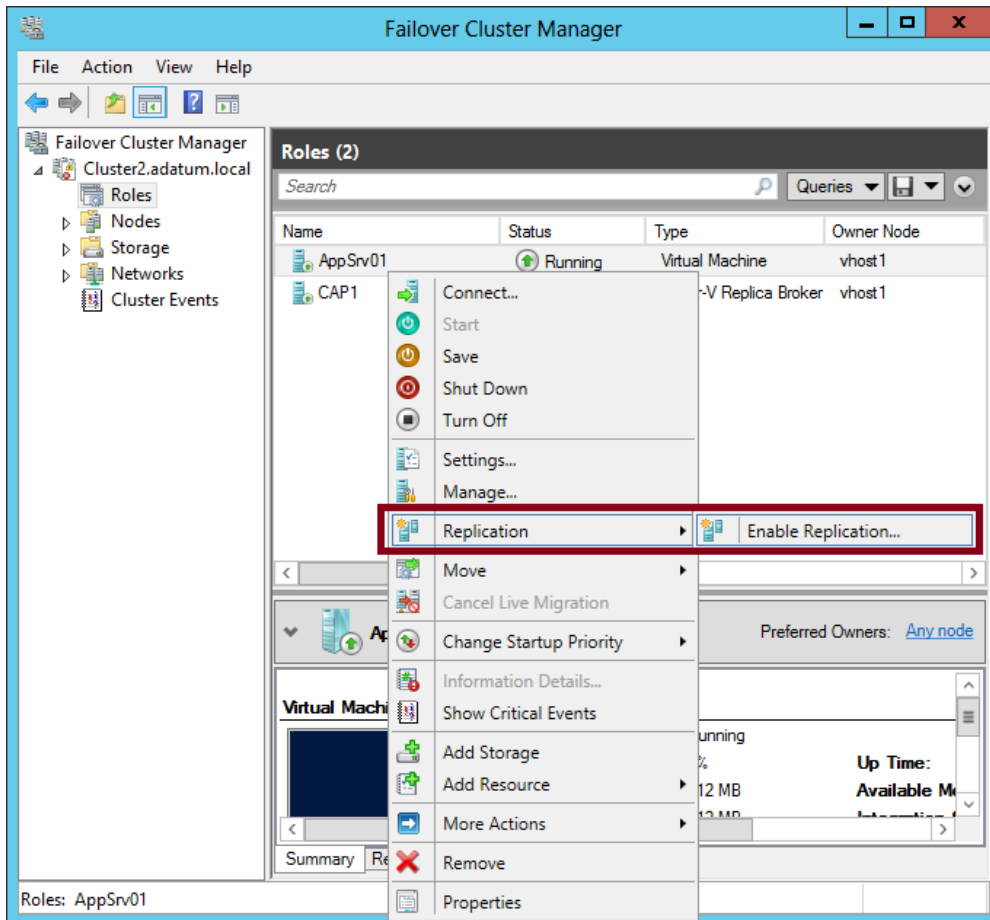


FIGURE 12-35 Enabling replication on a clustered VM

For the 70-417 exam, there's a good chance you'll be asked about basic concepts related to configuring replication on clustered VMs. Remember, first of all, that you use Failover Cluster Manager to configure replication for a clustered VM at the primary site but still use Hyper-V Manager at the Replica site. Remember that in Failover Cluster Manager at the primary site, you need to add the Hyper-V Replica Broker role to the failover cluster and that this role is used to configure Hyper-V Replica "server" settings for the cluster. Finally, you also need to remember that when you configure Hyper-V Replica in a failover cluster, the CAP name and address are used as the server name and address.

Objective summary

- Hyper-V Replica is a new feature in Windows Server 2012 and Windows Server 2012 R2 that creates an offline copy (replica) of a running VM and its storage. This replica can exist anywhere in the world. The online original (called the primary VM) periodically

sends the replica updates of any changes. In case the primary VM fails, you can fail over to the replica and bring it online.

- To configure Hyper-V Replica, you first configure authentication and authorization settings for both physical host servers, called the primary server and replica server. Then, in Hyper-V Manager on the primary server, run the Enable Replication Wizard for the desired VM.
- In Windows Server 2012 R2, you can choose a replication frequency between the primary and replica VMs from among the options of 30 seconds, 5 minutes, or 15 minutes. (In the first release of Windows Server 2012, replication frequency is 5 minutes and cannot be changed.)
- By default, you can fail over only to the most recent recovery point, which is the point when the most recent updates were received. However, you can choose to store additional, older recovery points that allow you to return to point-in-time snapshots of the primary VM.
- A planned failover is performed on the primary server after you shut down the primary VM. A planned failover brings the replica VM online with no loss of data. You can perform an unplanned failover on the replica server if the primary server fails without warning. With an unplanned failover, the replica VM recovers a copy of the primary VM that is normally no more than 5 to 15 minutes old. Finally, you can also perform a test failover while the primary VM is still running. A test failover brings a copy of the replica VM online in a state that is disconnected from the network.
- Windows Server 2012 R2 also enhances Hyper-V Replica by allowing you to extend replication from the replica server to a third server.
- If you want to configure Hyper-V Replica for a VM that is hosted in a failover cluster, you need to add the Hyper-V Replica Broker role to the cluster. You also need to provide a CAP name and address for the cluster that will act as the server name.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the “Answers” section at the end of the chapter.

1. You are configuring Hyper-V Replica on a VM that is hosting Microsoft Exchange. You want to help ensure that if you fail over to the replica VM, the application data will remain in a consistent state.

What should you do? (Choose all that apply.)

- A.** Configure the replica server to save additional recovery points.
- B.** Configure the primary server to replicate incremental VSS copies.
- C.** Configure a resynchronization schedule for the primary and replica VM.
- D.** Configure Hyper-V Replica Broker.

2. You have configured Hyper-V Replica for a VM named AppSrv1, which is hosted on a primary server named VMhost1 in Cleveland. The replica server is named RepHost1 and is located in Denver.

An unexpected power outage suddenly brings the entire Cleveland site offline. You perform a failover at the Denver site and start the replica VM on RepHost1. Power is returned to the Cleveland site after several hours, but only after changes have been made to AppSrv1.

You are able to bring VMhost1 back online and now want to return AppSrv1 to its original host. Which step should you take next?

- A. Perform an unplanned failover.
 - B. Choose the option to cancel the failover.
 - C. Perform a planned failover.
 - D. Choose the option to reverse replication.
3. Within your organization, a clustered VM named SQL1 is hosting SQL Server. The failover cluster hosting SQL1 is named Cluster1 and includes three nodes, named Node1, Node2, and Node3. Node1 is the preferred owner of the SQL1 VM. All three nodes are located in the same data center.
- You want to configure an offsite replica of SQL1 to protect the VM in case the entire failover cluster is brought down because of a power outage or other emergency.
- You deploy a physical server named RepSrv2 at a remote site. You want to configure RepSrv2 as the replica server. You install Windows Server 2012 and then the Hyper-V role on RepSrv2. You then connect the server to the Internet and establish a VPN connection between the two sites.
- Which of the following steps should you take? (Choose two.)
- A. At the primary site, configure Hyper-V Replica Broker and provide a CAP name.
 - B. At the replica site, configure Hyper-V Replica Broker and provide a CAP name.
 - C. In the replication settings on Cluster1, restrict authorization to the CAP.
 - D. In the replication settings on RepSrv2, restrict authorization to the CAP.



Thought experiment

Implementing business continuity and disaster recovery at Adatum

You are a network administrator for Adatum.com, an organization with headquarters in San Francisco and a branch office in Montreal. You are designing fault tolerance and business continuity for a new application server and VM that will be named AppVM1. AppVM1 will be hosted in the San Francisco office.

You want to meet the following goals:

- You want to prevent any disruption of service and data loss in case an individual server fails unexpectedly.
- You want to be able to resume service with minimal data loss in case a catastrophe such as an earthquake brings the main office offline for an extended period.
- You always want to retain daily backups from the previous two weeks.

With these goals in mind, answer the following questions. You can find the answers to these questions in the “Answers” section.

1. Which feature(s) in Windows Server 2012 can enable you to meet the first goal?
2. How might you design fault tolerance so that you can meet the first goal even after a catastrophe brings the main office offline for an extended period?
3. Describe two ways you might design fault tolerance for AppVM1 so that you can continue to meet the third goal even through a catastrophe that brings the main office offline for an extended period.

Answers

This section contains the answers to the Objective Reviews and the Thought Experiment.

Objective 12.1: Review

1. Correct answer: B

- A. Incorrect:** Changing the bandwidth assigned to the work hours will not help you achieve your goal of having the backup operation complete before the work day begins at 8:00 A.M.
- B. Correct:** The bandwidth setting assigned to nonwork hours is restricted to 1023.0 Kbps, which is much lower than the default setting of 1023 Mbps. This low setting could be unnecessarily limiting the bandwidth allowed at night. If you raise this value, the backup operation could proceed much more quickly (assuming more bandwidth is available.)
- C. Incorrect:** Adjusting the work hours could potentially cause disruption for workers, and it will not help you meet your goal of completing the backup operation before 9:00 A.M.
- D. Incorrect:** The work days are not currently affecting the backup because the backup is being performed outside of work hours. If you include Wednesday as a work day, you would actually apply bandwidth throttling to the first hour of the backup operation, and slow the procedure down for that hour.

2. Correct answer: C

- A. Incorrect:** This step would exclude the C:\Windows\Temp folder and its subfolders from the backup set, but it would not meet your goal of allowing the backup to be performed weekly. This folder is too small to reduce the size of the backup in any significant way.
- B. Incorrect:** This step would exclude the C:\Windows\Temp folder but not its subfolders from the backup set, but it would not meet your goal of allowing the backup to be performed weekly. Too little data is stored in this folder to reduce the size of the backup in any significant way.
- C. Correct:** This setting would allow the previous week's backup to be deleted to make space for the current week's backup. The size of the backup from the previous week is approximately 220 GB, and your storage quota is 300 GB. Consequently, you need to be able to remove the previous week's backup to make room for the current week's backup.
- D. Incorrect:** This setting would not fix your problem. It would require all backups to be kept at least 30 days on Microsoft servers. If there is insufficient space to allow a new backup, as is the case in this scenario, the new backup will fail.

3. Correct answer: A

- A. Correct:** You don't need to modify the default settings. The bandwidth of the backup operation will be throttled to 256 Kbps beginning at 9:00 A.M. every weekday.
- B. Incorrect:** You don't want to increase the bandwidth settings assigned to work hours because this would increase the impact on network performance for users during work hours.
- C. Incorrect:** Increasing the default setting would not have any effect. No Internet download speeds currently available are higher than the default setting of 1023.0 Mbps.
- D. Incorrect:** You don't need to adjust work days because the current selection reflects the Monday through Friday schedule of the organization.

Objective 12.2: Review

1. Correct answers: A, B

- A. Correct:** You need to enable the option to save additional recovery points. This step allows you to configure some of these additional recovery points as incremental VSS copies, which are application-consistent.
- B. Correct:** Incremental VSS copies are snapshots that are application-consistent for VSS-aware applications like Microsoft Exchange.
- C. Incorrect:** Resynchronization does not affect the consistency of applications within recovery point snapshots.
- D. Incorrect:** Hyper-V Replica Broker is used for failover clustering, not for application consistency.

2. Correct answer: D

- A. Incorrect:** You have already performed an unplanned failover. You cannot perform failover to the other site until replication is reestablished between the two servers.
- B. Incorrect:** It's too late to cancel the failover because changes have already been made to AppSrv1.
- C. Incorrect:** You cannot perform a planned or unplanned failover to the other site until replication is reestablished.
- D. Correct:** Choosing the option to reverse replication starts the Reverse Replication Wizard. This wizard lets you reestablish replication between the two servers, with the local server in Denver acting as the new primary. After you complete this wizard, you can perform a planned failover to return the VM to the site in Cleveland.

3. Correct answers: A, D

- A. Correct:** You need to configure the Hyper-V Replica Broker role for the failover cluster if you want to add an offsite replica to a clustered VM.
- B. Incorrect:** To configure the Hyper-V Replica Broker at the replica site, you would need to create a failover cluster at the replica site. This step is unnecessary because you want to configure RepSrv2 as the replica server. Your goal is not to create a replica cluster.
- C. Incorrect:** In the replication settings for Cluster1, you want to restrict authorization to RepSrv2. However, this step is not immediately necessary. It would be required only if the VM were failed over to the replica site, and you later wanted to fail back to the original site.
- D. Correct:** The server-level replication settings allow you to limit which remote servers can act as a primary server to the local replica server. In this case, you need to configure the CAP as the name of the primary server.

Thought experiment

- 1.** Only failover clustering can prevent any disruption of service and data loss in case of an individual server failure.
- 2.** You can configure Hyper-V Replica on failover clusters in both the San Francisco and Montreal offices. The failover cluster in the San Francisco office can act as the primary server, and the failover cluster in the Montreal office can act as the replica server.
- 3.** One option is to use a cloud backup service such as Windows Azure Backup to back up AppVM1 daily and specify a retention range of 15 days. Another option is to perform daily backups of AppVM1 to local file storage on a file server that is itself a VM. You can then configure this file server as a primary VM with a replica VM in the replica site (Montreal). In case of site-level failure at the primary site, the replica VMs of AppVM1 and the file server at the replica site will continue to operate as before with no loss of backup data.

This page intentionally left blank

Configure network services

The Configure Network Services domain includes a single objective: Deploy and manage IP Address Management (IPAM). IPAM is a new feature used for managing your organization's entire IP address space, public and private.

IPAM is a large topic, but as a new feature, the questions you'll see on the 70-417 exam will most likely not require unusually deep knowledge. That said, be sure to supplement the information in this chapter with some hands-on practice so that you can develop a feel for how IPAM works.

Objectives in this chapter:

- Objective 13.1: Deploy and manage IPAM

Objective 13.1: Deploy and manage IPAM

On the surface, IPAM seems easy. Here's the quick explanation of how it works: You have an IPAM server that automatically collects information from your infrastructure servers about the IP address ranges used on your network. You then use the IPAM management interface as a reference about these same address ranges. What could be complicated about that?

Unfortunately, IPAM is not quite as easy to master as it might appear at first. The difficulty is not with the feature on a conceptual level, but rather in the sheer number of component features and functionalities that IPAM includes. The first official Microsoft whitepaper about IPAM was almost 100 pages long, which gives you some idea of its feature depth. There's a lot there—too much, in fact, to be covered in one objective.

The best way to handle the large subject of IPAM, as a result, is to focus on three basic topics: What can you use the feature to accomplish? How do you configure it? And finally, how do you use it?

This section covers the following topics:

- Configuring IPAM
- Creating and managing IP blocks and ranges
- Delegating IPAM administration
- Role-based access control

What is IPAM?

IPAM is a useful new feature in Windows Server 2012 and Windows Server 2012 R2 that lets you centrally view, manage, and configure the IP address space in your organization. With IPAM, you can look at all your address blocks and ranges, find free IP addresses, manage DHCP scopes across multiple servers, create DHCP reservations and DNS host records, and even search for address assignments by device name, location, or other descriptive tag.

IPAM works by first discovering your infrastructure servers and importing from them all available IP address data. You then manually add whatever additional data you need to complete the picture of your organization's IP address assignments. Once you have this information in place, you can track updates to your IP address space.

Problems solved by IPAM

IPAM has many component features that help you manage IP addressing in different ways. To better understand the purpose and functionality of IPAM, it's helpful to view IPAM as a means to solve the following kinds of administrative problems:

- How can I track my organization's address space and know the addresses that are either in use or available across different locations?
- How can I find a free static IP address for a new device and register it in DNS?
- How can I find out which DHCP scopes in my organization are full or close to being full?
- How can I efficiently change a DHCP option across dozens of scopes residing on multiple servers?
- How can I find an unused address range within our organization's address space to dedicate to a new subnet?
- How can I determine which public and private address ranges are used by my organization?
- How can I determine which portion of the address space used by the organization is dynamically assigned, and which part is statically assigned?
- How can I search for and locate an IP address or set of addresses by name, device, location, or another descriptive tag?

Limitations of IPAM

IPAM is a new feature, and as such, it's important to recognize some of the limitations in this first release:

1. IPAM can import data only from Windows servers running Windows Server 2008 and later that are members of the same Active Directory forest.

2. IPAM does not support management and configuration of non-Microsoft network elements.
3. IPAM does not check for IP address consistency with routers and switches.
4. Address utilization trends and reclaiming support are provided only for IPv4.
5. IPAM does not support auditing of IPv6 stateless address auto configuration on an unmanaged machine to track the user.

Installing and configuring IPAM

To install the IPAM feature in Windows Server 2012 and Windows Server 2012 R2, you can use the Add Roles And Features Wizard or the following Windows PowerShell command:

```
Install-WindowsFeature IPAM -IncludeManagementTools
```

There are some important limitations you need to know about where to install IPAM. First, you can't install IPAM on a domain controller. Second, you should not install an IPAM server on a network infrastructure server, such as one running DNS or DHCP. If you install IPAM on a DHCP server, discovery of DHCP servers will be disabled.



EXAM TIP

Remember these deployment limitations for the exam because they could easily form the basis of a question. For example, if your IPAM server cannot discover DHCP servers, make sure it is not installed on a DHCP server.

IPAM must be installed on a domain member computer running Windows Server 2012 or later. The IPAM server is intended as a single purpose server. Once IPAM is installed, you configure and manage the feature through the IPAM client in Server Manager, as shown in Figure 13-1, or by using Windows PowerShell cmdlets from the IpamServer module. (There is no other graphical IPAM console.)



EXAM TIP

You can install just the IPAM client tool without installing the server component. To accomplish this task by using the Add Roles And Features Wizard, select IPAM in the wizard, choose to install the prerequisite features of IPAM, clear the selection of IPAM you have just selected, and then complete the wizard. *The IPAM client doesn't appear by default in Server Manager, however.* To make the IPAM client appear, you need to add the remote IPAM server to Server Manager by using the "Add Other Servers To Manage" option (visible in Figure 13-1).

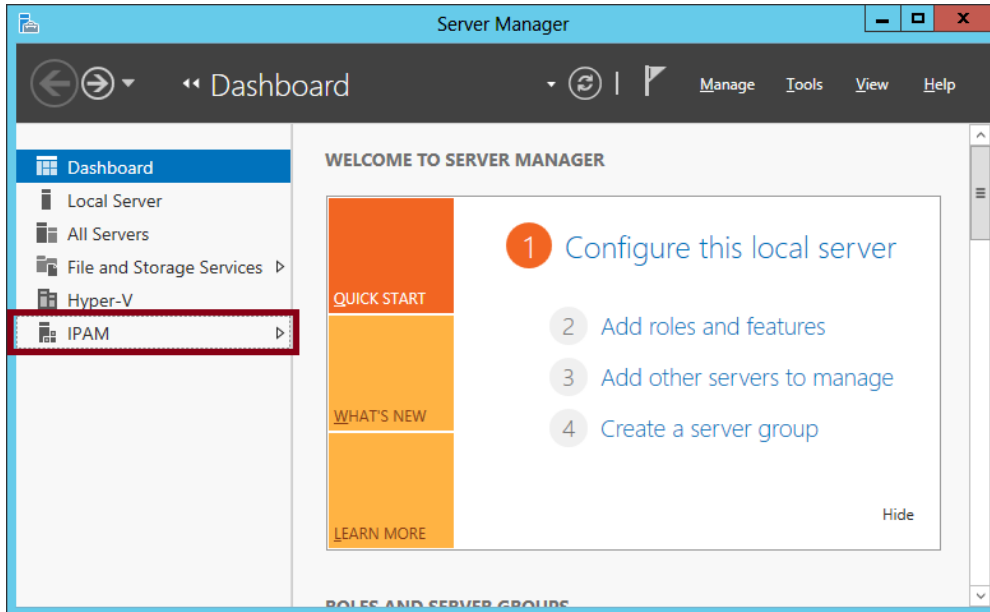


FIGURE 13-1 IPAM in Server Manager

When you click IPAM in the navigation pane of Server Manager, the navigation pane narrows, and the details pane reveals the IPAM Overview page, shown in Figure 13-2.

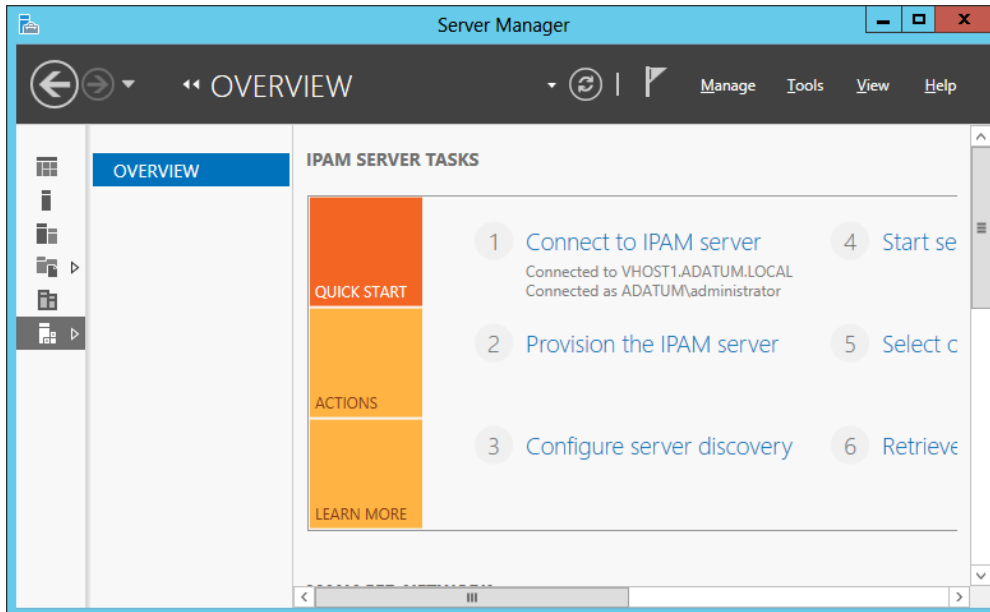


FIGURE 13-2 The IPAM Overview page preconfiguration

The Overview page presents the following six links that help guide you through configuration:

1. Connect To IPAM Server
2. Provision The IPAM Server
3. Configure Server Discovery
4. Start Server Discovery
5. Select Or Add Servers To Manage And Verify IPAM Access
6. Retrieve Data From Managed Servers

We'll use these same steps to cover the configuration process in the next sections.

1. Connect To IPAM Server

You use this step only if you need to connect to a remote IPAM server. By default, Server Manager is connected to the local IPAM server.

2. Provision The IPAM Server

Clicking this step on the Overview page starts the Provision IPAM Wizard. *Provisioning* the IPAM server is the term used to prepare the IPAM server by performing steps such as creating the IPAM database, creating IPAM security groups, and configuring access to IPAM tasks and folders.

Windows Server 2012 R2 introduces a new page in this wizard, which is the Configure Database page that is shown in Figure 13-3. In the first release of Windows Server 2012, the IPAM database was always a Windows Internal Database (WID). In Windows Server 2012 R2, you can now choose between a WID database and a Microsoft SQL Server database. The Microsoft SQL Server database you connect to, if you choose that option, can be installed on the local machine or on a remote computer.

The main advantages of choosing a Microsoft SQL Server database are increased scalability, improved disaster recovery capabilities, and enhanced reporting. Remember these advantages for the 70-417 exam.



EXAM TIP

The Deploy and Manage IPAM objective for 70-417 was updated for Windows Server 2012 R2 in January 2014 to include just one additional task: Configure IPAM database storage. This task requires you to know the setting shown in Figure 13-3. You should also know the `Get-IpamDatabase` cmdlet, which provides information about how the IPAM database is currently stored, and the `Move-IpamDatabase` cmdlet, which lets you migrate the IPAM database to a SQL Server database, either from a WID database or from another SQL Server database.

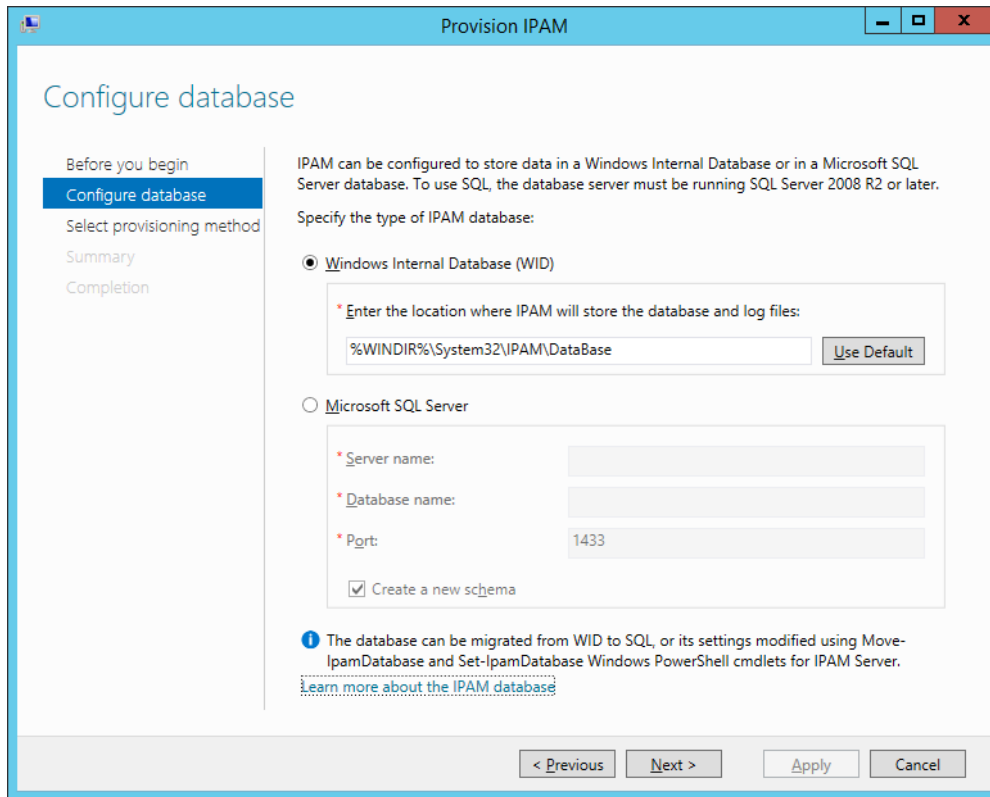


FIGURE 13-3 Choosing the IPAM database type

You also use the Provision IPAM Wizard to determine how you want to configure the infrastructure servers that IPAM will manage. The two choices are to configure the infrastructure servers manually or to do so through Group Policy, as shown in Figure 13-4. If you choose to use Group Policy, you specify a prefix for the three Group Policy Objects (GPOs) that will later be created automatically when you use the `Invoke-IpamGpoProvisioning` cmdlet.

You wouldn't select Manual here unless there was some sort of unusual factor that made the Group Policy Based option impossible or ineffective. Despite this limited real-world applicability of the Manual option, configuring IPAM manually is one of the tasks officially mentioned in the Deploy and Manage IPAM objective. (The process of manual configuration is discussed in the section "5. Select Or Add Servers To Manage" later in this chapter.)

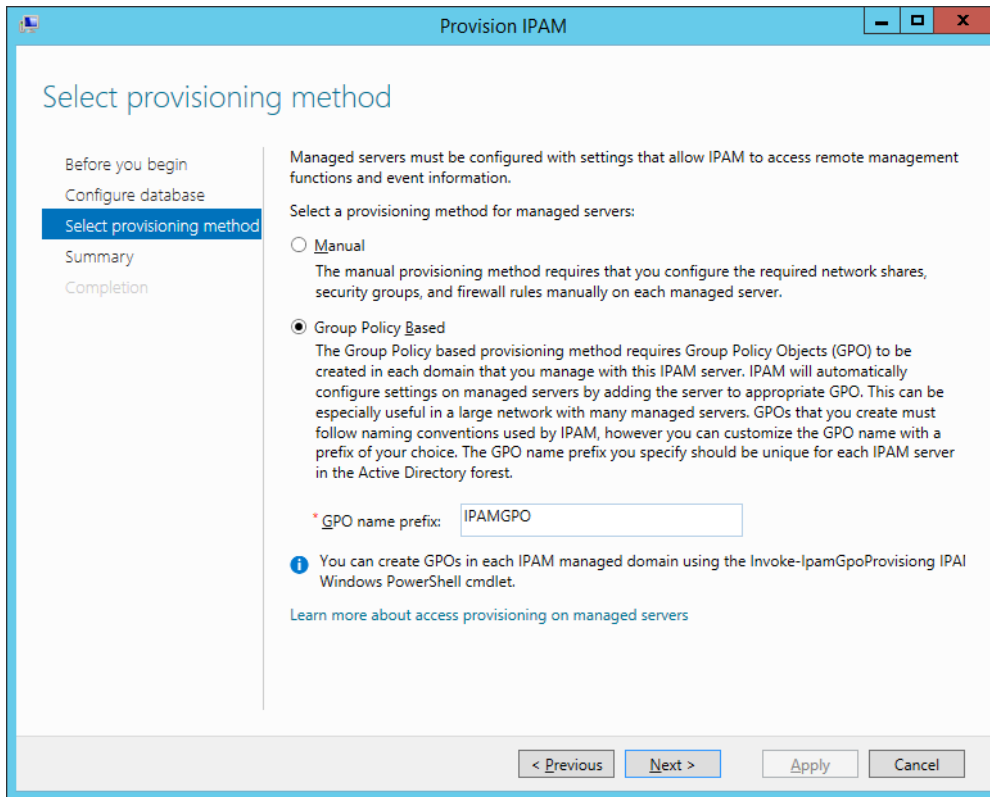


FIGURE 13-4 Choosing Group Policy configuration with a GPO name prefix



EXAM TIP

In the first release of Windows Server 2012, you couldn't change the provisioning method after you completed the Provision IPAM Wizard. If you needed to change the provisioning method, your only option was to uninstall and reinstall IPAM.

In Windows Server 2012 R2, the situation has improved. You can use the following Windows PowerShell command after completing the Provision IPAM Wizard to change the provisioning method *in one direction only*—from manual to GPO-based (automatic).

```
Set-IpamConfiguration -ProvisioningMethod Automatic
```

Even in Windows Server 2012 R2, however, you can't switch from GPO-based (automatic) to manual after you run the Provision IPAM Wizard. If you do need to change the provisioning method from GPO-based to manual, you still have to uninstall IPAM, reinstall IPAM, and finally run the Provision IPAM Wizard again with the proper selection.

Remember these points along with the complete Windows PowerShell command shown above.

3. Configure Server Discovery

Clicking this link on the Overview page opens the Configure Discovery Settings dialog box, shown in Figure 13-5. You use this step to specify which types of infrastructure servers you want to discover. By default, all three possible infrastructure types are selected: Domain Controller, DHCP Server, and DNS Server.

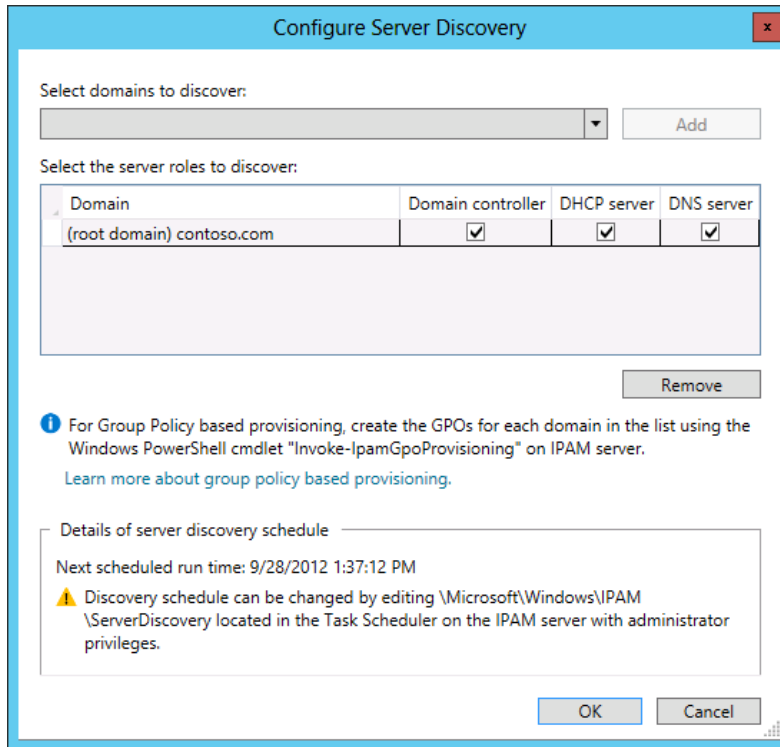


FIGURE 13-5 Selecting infrastructure server types to discover

4. Start Server Discovery

This link begins the process of discovering infrastructure servers in your environment. To determine when the process is complete, click the notification flag in Server Manager, and then click Task Details. The process is complete when the IPAM ServerDiscovery task displays a status of Complete, as shown in Figure 13-6.

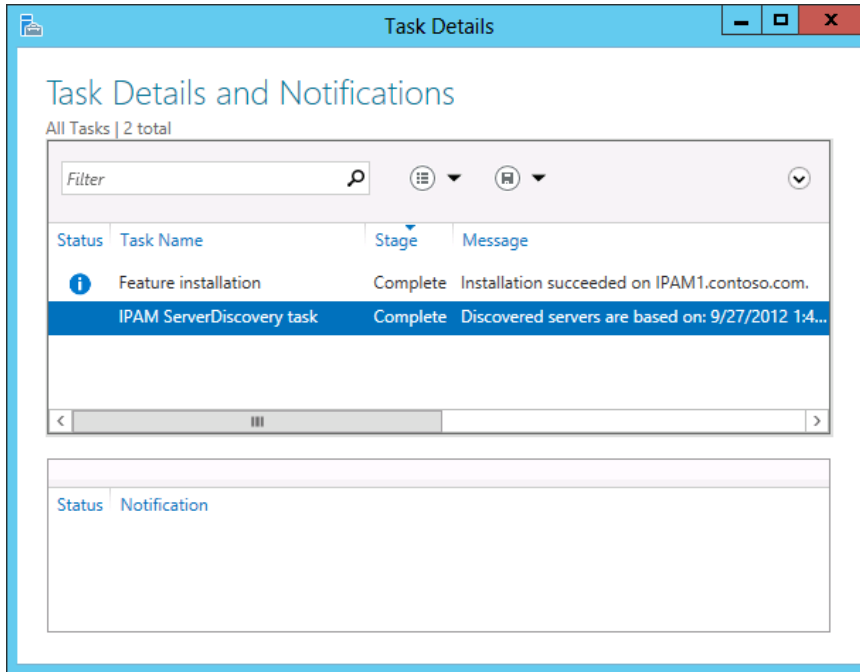


FIGURE 13-6 Server discovery complete

5. Select Or Add Servers To Manage

Clicking this link on the Overview page displays the SERVER INVENTORY page in the IPAM context of Server Manager. This page shows the servers that have been discovered by the server discovery task in the previous step. At first, the discovered servers display a Manageability Status of Unspecified and an IPAM Access Status of Blocked, as shown in Figure 13-7. This status simply means you still need to configure the servers for IPAM management. To perform this step, you need to run a Windows PowerShell command and then designate the desired servers as managed. (You need to perform this step if you have chosen the Group Policy Based option on the Select Provisioning Method page shown in Figure 13-4. If you have chosen the Manual option, the entire IPAM configuration process is different. For instructions on manual configuration, see the sidebar “Manual configuration of managed servers” later in this chapter.)

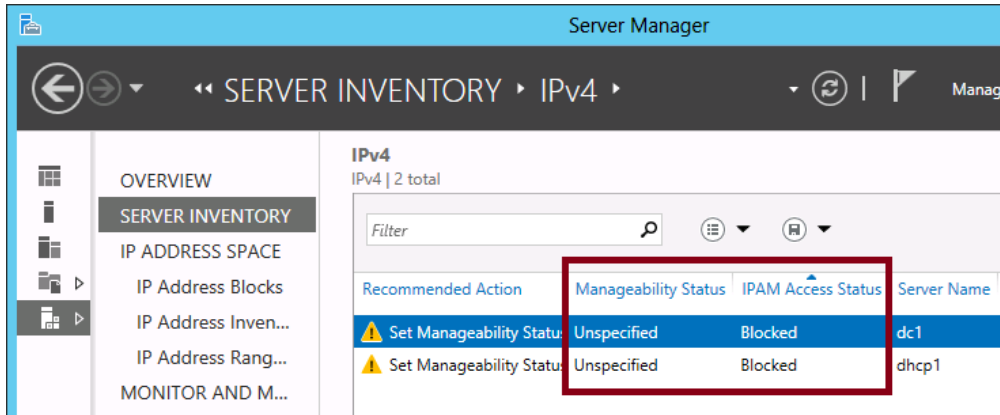


FIGURE 13-7 Discovered servers that need to be configured for IPAM management

To configure the servers through the Group Policy Based provisioning method, you need to create IPAM GPOs. You can do this by running following Windows PowerShell command:

```
Invoke-IPamGpoProvisioning [-Domain] <String> [-GpoPrefixName] <String> [-IpamServerFqdn <String> ]
```

The GPO prefix name should be the same one that you specified in the Provision IPAM Wizard. For example, if you specified a prefix of IPAMGPO in the Provision IPAM Wizard, you could enter the following command at an elevated Windows PowerShell prompt:

```
Invoke-IPamGpoProvisioning -Domain contoso.com -GpoPrefixName IPAMGPO -IpamServerFqdn ipam1.contoso.com
```

This command creates the three GPOs shown in Figure 13-8.

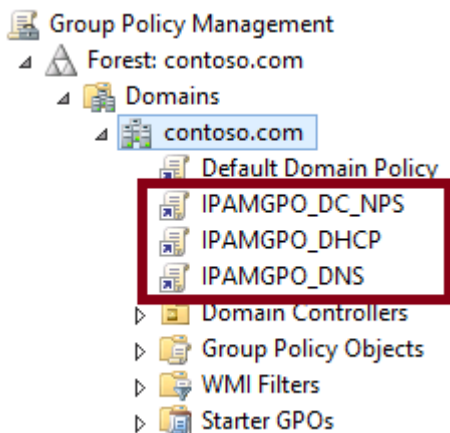


FIGURE 13-8 GPOs created for IPAM



EXAM TIP

If you forget the GPO prefix that you specified when you ran the Provision IPAM Wizard, use the `Get-IpamConfiguration` cmdlet. This cmdlet will display the GPO prefix along with other basic information about the IPAM configuration, such as the version number, the port used, and the provisioning method.

These three new GPOs apply only to servers that you designate as managed, but no servers are designated as managed by default. (Remember this last point for the exam because it could easily serve as the basis for a test question.) To change the manageability status of servers, right-click each server you want to manage on the SERVER INVENTORY page in the IPAM context in Server Manager, and then click Edit Server. In the Add Or Edit Server dialog box that opens, in the Manageability Status drop-down list, select Managed (as shown in Figure 13-9), and then click OK.

NOTE The IPAM provisioning process creates a domain security group named IPAMUG. This group is used to grant permissions to managed servers.

The screenshot shows the 'Add or Edit Server' dialog box with the following details:

| Field | Value |
|----------------------|--|
| * Server name (FQDN) | DC1.corp.contoso.com Verify |
| * IP address | 10.0.0.1 |
| * Server type | <input checked="" type="checkbox"/> DC <input checked="" type="checkbox"/> DNS server <input checked="" type="checkbox"/> DHCP server <input type="checkbox"/> NPS server |
| Manageability status | Unspecified |
| Owner | Unspecified |
| Description | Unmanaged |

The 'Manageability status' dropdown menu is open, showing the following options:

- Managed

FIGURE 13-9 Setting a server's manageability status to Managed

Finally, you need to force an update of Group Policy on all the servers you have designated as managed. You can do this, of course, either by running `Gpupdate /force` on each of these servers, by restarting them, or by invoking `Gpupdate` centrally in the methods described in Chapter 9, “Configure and manage Group Policy.”

Next, click the refresh icon in Server Manager in the menu bar next to the notification flag. (Alternatively, you can right-click your servers on the SERVER INVENTORY page and select the Refresh Server Access Status option. You can see this option on the shortcut menu in Figure 13-11.) After you refresh the server status, the Manageability Status of the servers will appear as Managed, and the IPAM Access Status will appear as Unblocked on the SERVER INVENTORY page, as shown in Figure 13-10. Note that the manageability status can require a relatively long time to be updated in the interface. If it isn't updated within a few minutes, you might need to wait an hour or more before refreshing reveals the status change .



FIGURE 13-10 Servers that are configured to be managed by IPAM

MANUAL CONFIGURATION OF MANAGED SERVERS

Configuring IPAM manually from start to finish without the use of Group Policy involves a far more elaborate and cumbersome process than is normally required of you to learn for Microsoft certification exams.

However, there are aspects of manual configuration that are easily summarized and could plausibly appear on the exam, shown in Table 13-1. The most likely elements to appear in an exam question are the firewall ports created on each server and the security groups the IPAM server needs to join.

If you want to learn the full step-by-step procedure for configuring IPAM manually, consult the document entitled “Understand and Troubleshoot IP Address Management (IPAM) in Windows Server “8” Beta,” available at <http://www.microsoft.com/en-us/download/details.aspx?id=29012>. (The steps for manual configuration appear in the first appendix of the document.)

TABLE 13-1 Manual configuration steps for managed infrastructure servers in IPAM

| On this Managed Server... | Perform this configuration step | enable these Firewall Rules | Associated IPAM functionality |
|----------------------------------|---|---|---|
| DHCP | Add the IPAM server to the local DHCP Users security group | DHCP Server (RPC-In) DHCP Server (RPCSS-In) | DHCP address space, settings, and utilization data collection |
| | Assign to the IPAM server Read access in the DHCP Server service access control list (ACL) | Remote Service Management (RPC) Remote Service Management (RPC-EPMAP) | DHCP Service monitoring |
| | Add the IPAM server to the local Event Log Readers security group | Remote Event Log Management (RPC) Remote Event Log Management (RPC-EPMAP) | DHCP configuration event monitoring |
| | Create a network share named DHCPAudit for %windir%\system32\dhcp and assign read access to the IPAM server on this share | File and Printer Sharing (NB-Session-In) File and Printer Sharing (SMB-In) | DHCP lease event collection for IP address tracking |
| DNS | For DNS servers that are also domain controllers, assign to the IPAM server Read access in the domain-wide DNS ACL* OR For DNS servers that are not domain controllers, add the IPAM server to the local Administrators group | DNS Service RPC DNS Service RPC Endpoint Mapper | DNS zone configuration collection |
| | Add the IPAM server to the local Event Log Readers security group Assign to the IPAM server Read access in the ACL stored in the DNS CustomSD registry key | Remote Event Log Management (RPC) Remote Event Log Management (RPC-EPMAP) | DNS zone event collection for DNS zone monitoring |
| | Assign to the IPAM server Read access in the DNS Server service ACL | Remote Service Management (RPC) Remote Service Management (RPC-EPMAP) | DNS service monitoring |
| DC/NPS | Add the IPAM server to the local Event Log Readers security group | Remote Event Log Management (RPC) Remote Event Log Management (RPC-EPMAP) | Logon event collection for IP address tracking |
| IPAM (local server) | Add the local Network Service to the local Event Log Readers security group | N/A | IPAM configuration event monitoring |



EXAM TIP

The IPAM server needs to be able to read the event logs on the DHCP, DNS, DC and NPS servers. For this reason, it needs to be added to the local Event Log Readers security group on all of these servers.

6. Retrieve Data From Managed Servers

The final step in configuring IPAM is to load data from your managed servers into the IPAM database. To do so, on the Overview page, click Retrieve Data From Managed Servers. Then click the notification flag and wait for all tasks to complete.

Alternatively, you can select and right-click the managed servers on the SERVER INVENTORY page and then select Retrieve All Server Data from the shortcut menu, as shown in Figure 13-11.

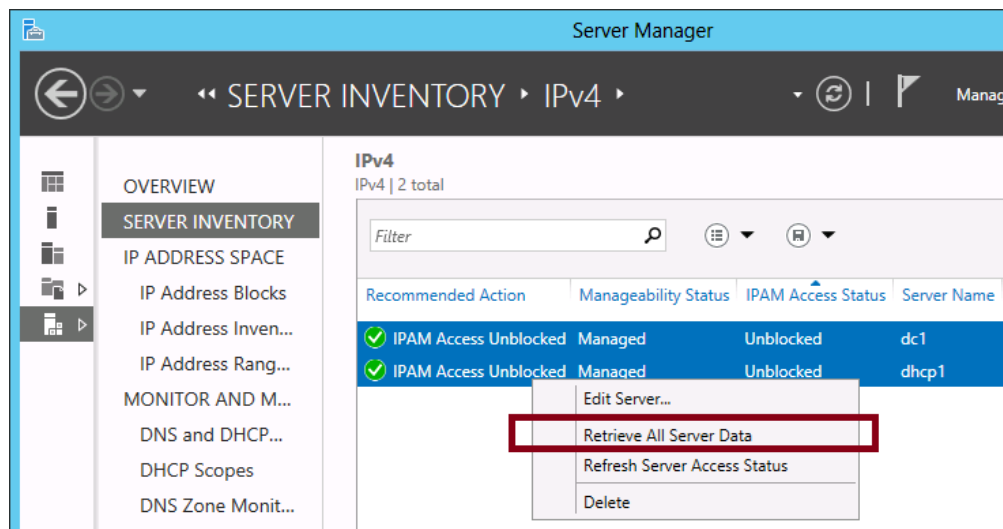


FIGURE 13-11 Retrieving data from managed servers

Managing address space

The most basic function of IPAM is to let you view, monitor, and manage the IP address space in your organization. With IPAM, you can search and sort IP blocks, ranges, and individual addresses based on built-in fields or user-defined custom fields. You can also track IP address utilization within scopes or display utilization trends.

Adding your IP address space to the IPAM database

You can browse and search your organization's address space in the IPAM client in Server Manager, but only after you add this data to the IPAM database. In IPAM, the IP address space is broken down into blocks, ranges, and addresses. Blocks represent the largest sections

of IP address space used by a company, such as 10.0.0.0/8. Ranges are portions of blocks that typically correspond to DHCP scopes. IP addresses exist within ranges.

How are these elements added to IPAM? IPAM discovers DHCP scopes automatically and it automatically imports the corresponding address ranges into the IPAM database. However, IPAM doesn't import blocks or addresses into its database automatically. You have to add them manually or import them from a comma-delimited file. (You can also export addresses to a file in comma-delimited format.) In addition, you might want to create additional ranges that have not been discovered by IPAM. For example, these address ranges might correspond to space reserved for statically assigned addresses or the ranges might be used by DHCP servers that are not members of the local Active Directory forest.

To add an IP address block, subnet, range, or address to the IPAM database, click IP Address Blocks in the IPAM navigation menu, and then, from the Tasks menu, select Add IP Address Block, Add IP Address Subnet, Add IP Address Range, or Add IP Address, as shown in Figure 13-12. (Note that the option to add a new subnet is new to Windows Server 2012 R2.)

NOTE Although IPAM doesn't automatically import addresses associated with DHCP leases into the IPAM database, you can use the EVENT CATALOG page and the IP Address Tracking tool to search the DHCP server log directly for DHCP leases by client name or address.

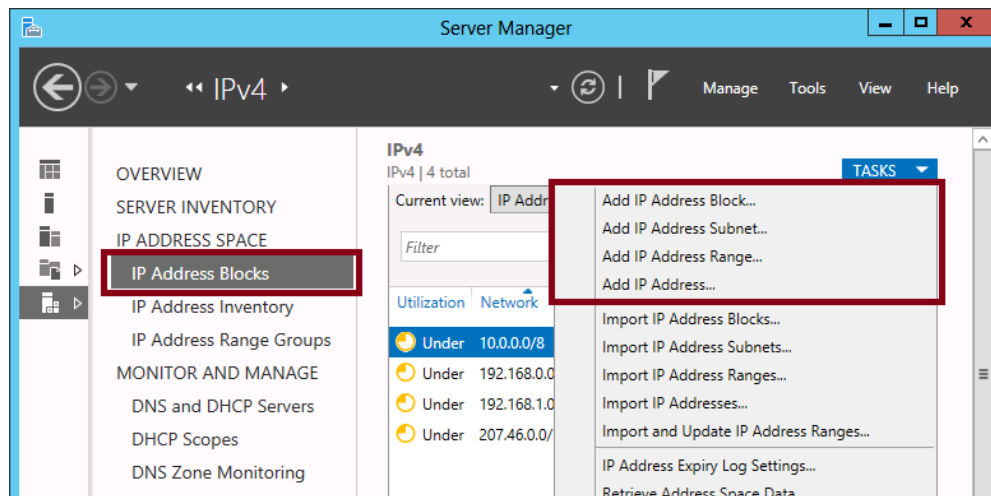


FIGURE 13-12 Adding a block, subnet, range, or address to the IPAM database

After this step, you can view the elements you have added by selecting the appropriate category (IP addresses, address blocks, address ranges, or subnets) in the Current View drop-down list, as shown in Figure 13-13.

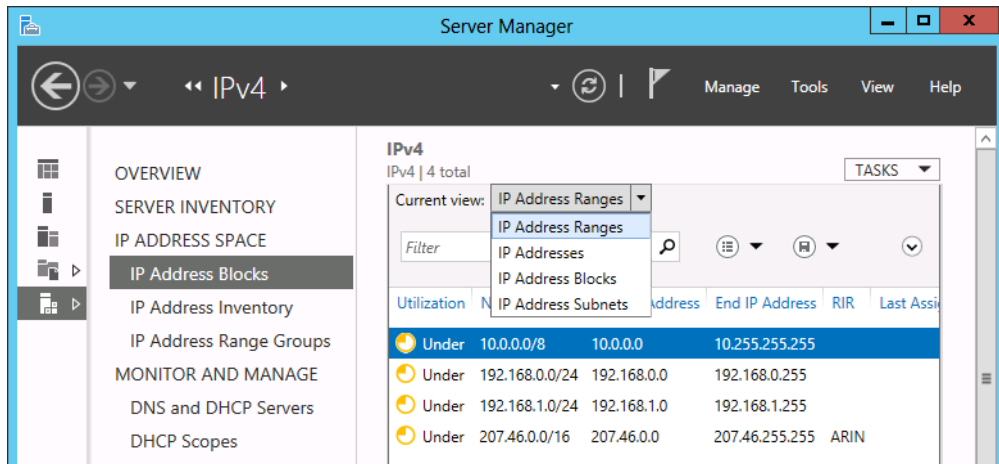


FIGURE 13-13 Viewing added blocks, addresses, or ranges



EXAM TIP

Remember that the only address ranges automatically imported into the IPAM database are ones that correspond to DHCP scopes on managed DHCP servers.

NOTE IPAM in Windows Server 2012 R2 introduces the ability to view and manage the address space used for virtual networks separately from the address space used for physical networks. To view virtual address space in IPAM in Windows Server 2012 R2, click the VIRTUALIZED ADDRESS SPACE node in the upper navigation pane of the IPAM console.

Creating custom fields

You can create custom data fields that you can later apply to your blocks, ranges, and individual addresses. You can then use these fields to sort or locate IP address information in a way that is useful to you, such as by office location, building, floor, or department.

To create a custom field for IPAM, first select IPAM Settings from the Manage menu in Server Manager, as shown in Figure 13-14.

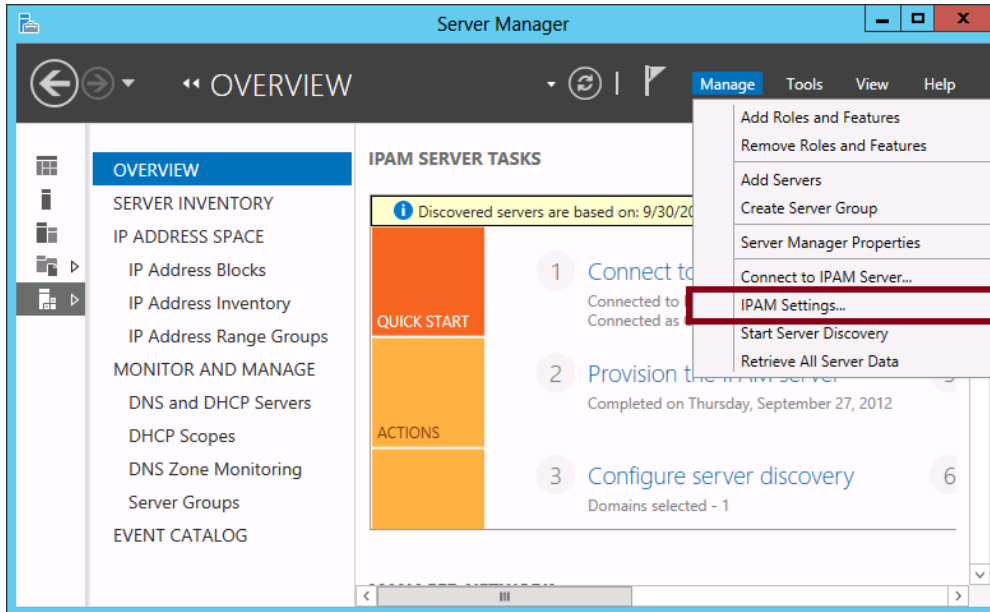


FIGURE 13-14 Opening IPAM settings

The IPAM Settings dialog box is shown in Figure 13-15. In this dialog box, select Configure Custom Fields.

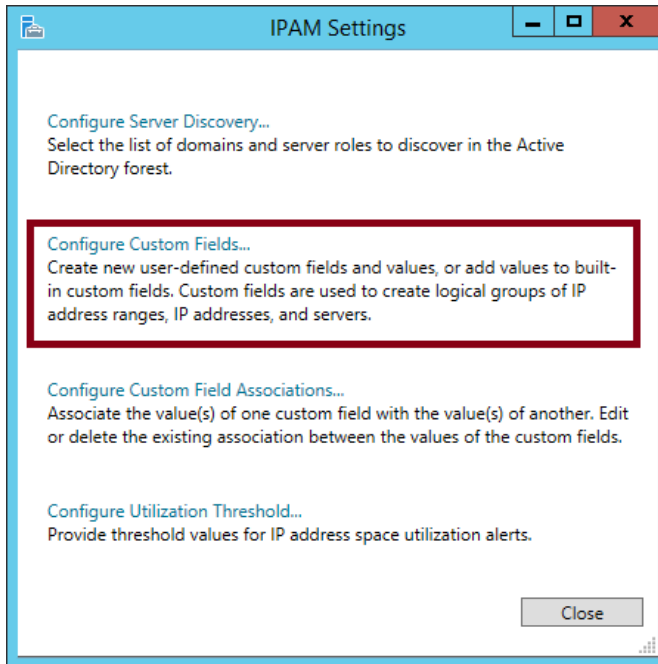


FIGURE 13-15 Configuring custom fields

This step opens the Configure Custom Fields dialog box, shown in Figure 13-16. Use this dialog box to create custom field names such as Building and then possible values for that value, such as Headquarters, Sales, Operations, and Data Center.

Configure Custom Fields

You can extend built-in custom fields below by adding additional values, or create new user-defined custom fields. Custom fields can be associated with IP address ranges, IP addresses, and servers to create logical groups.

Step 1:
Add custom fields below:

| Custom Field Name | Multi-Value | Category |
|---------------------|-------------|------------|
| VMM Logical Network | No | Built-in |
| Building | Yes | User defin |

Delete custom field

Step 2:
Select a multi-value custom field above and provide unique values for the field below:

| Custom Field Value |
|--------------------|
| Headquarters |
| Operations |
| Sales |
| Data Center |

Delete value

Note: Changes to custom field names or values will affect all associated entities and logical groups.

OK Cancel

FIGURE 13-16 Configuring custom fields



EXAM TIP

Remember that you can use custom fields to categorize the IP addresses and ranges in your IPAM database.

Applying a custom field to addresses and ranges

To add a custom field to an IP address range or IP address, right-click the element and select the option to edit it. (You can edit multiple ranges or addresses simultaneously.) Then click Custom Configuration in the associated dialog box and provide the desired field and value.

Creating IP address range groups

An IP address range group is a view of IP addresses or ranges sorted by stacked categories, as shown in Figure 13-17. In this figure, an IP address range group named Building/Floor has been created. When you select this IP address range group, you can view or search ranges and address by building name and then by floor name.

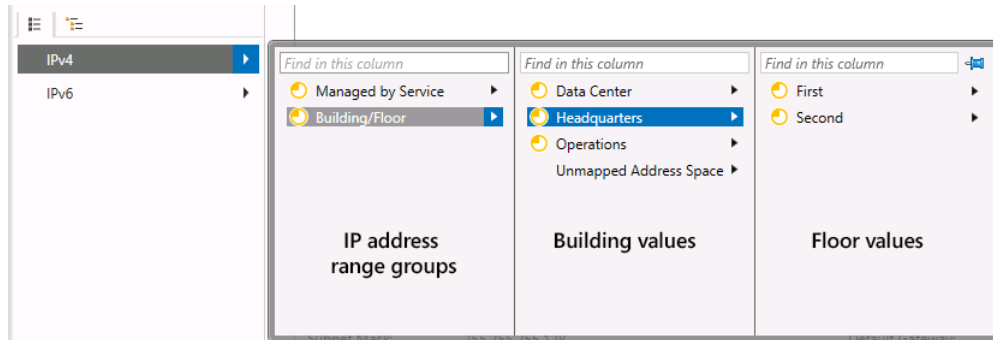


FIGURE 13-17 An IP address range group

To create an IP address range group, select the IP Address Range Groups page in the IPAM navigation pane in the IPAM client in Server Manager. Then, from the Tasks menu, select Add IP Address Range Group, as shown in Figure 13-18. A simple dialog box will then open that allows you to specify parent and child values for your new address range group.

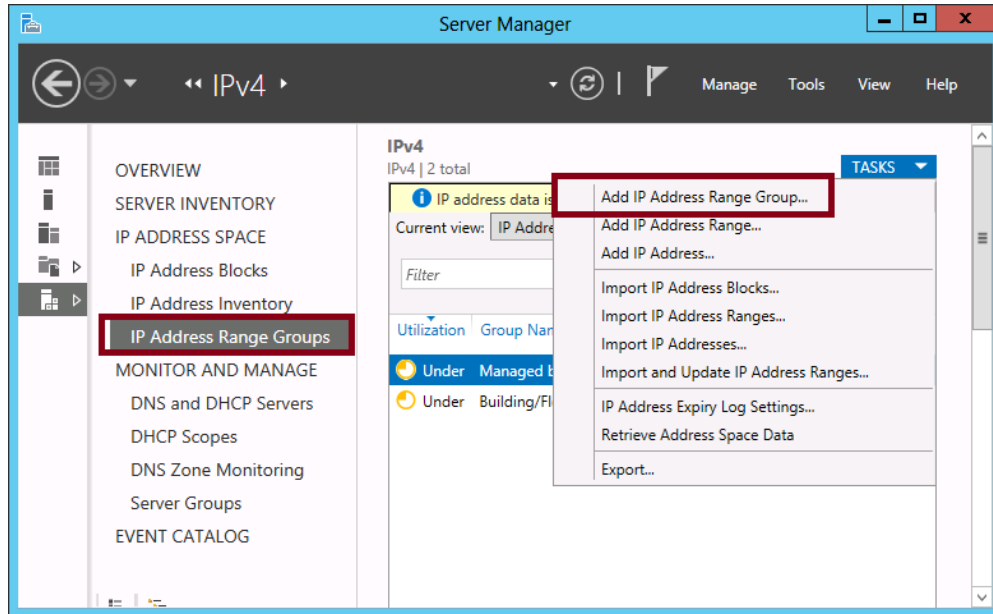


FIGURE 13-18 Adding an IP address range group

Finding and allocating an address from a range

Finding an unused address to allocate to a device is a common task for a network administrator. IPAM lets you perform this task within a chosen IP address range and it verifies for you that the IP address is unused.

To perform this function, select the IP Address Blocks page in the IPAM client in Server Manager and verify that IP Address Ranges is selected in Current View. Then right-click the desired IP Address range and select Find And Allocate Available IP Addresses, as shown in Figure 13-19.

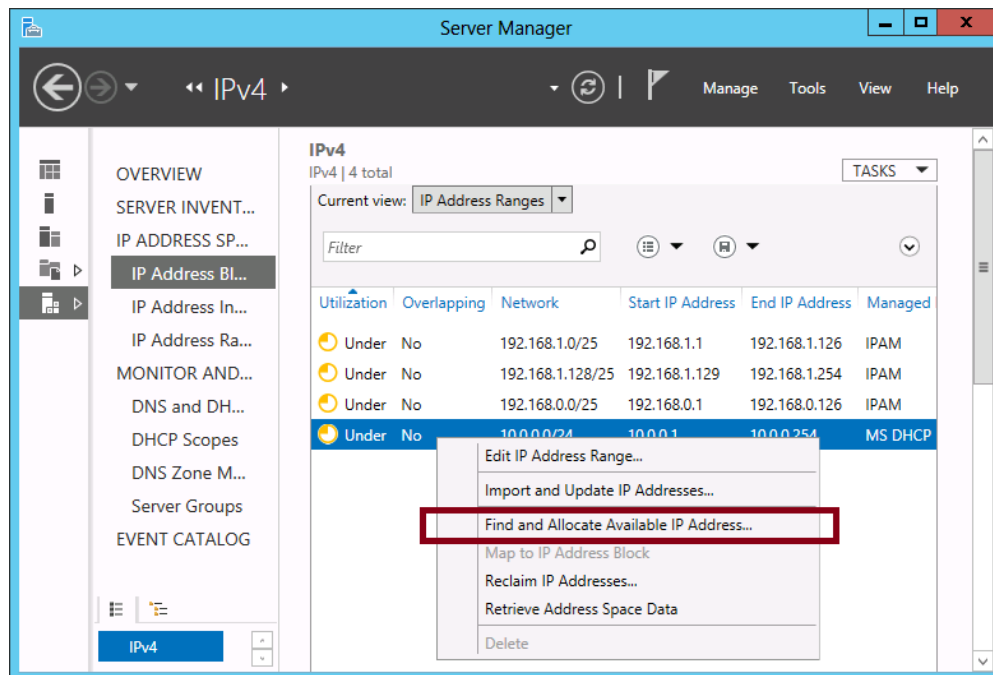


FIGURE 13-19 Finding an available IP address

Viewing and configuring IP utilization thresholds

When you select the IP Address Blocks page, the first column in the list of blocks or ranges displayed is Utilization, as shown in Figure 13-20. This value lets you know how much of a displayed address block or range is already assigned to devices. By default, if fewer than 20 percent of the addresses defined in the range are in use, the status reads Under. If more than 80 percent of the addresses are in use, the status reads Over. Optimal utilization is shown when the IP address usage is between 20 percent and 80 percent.

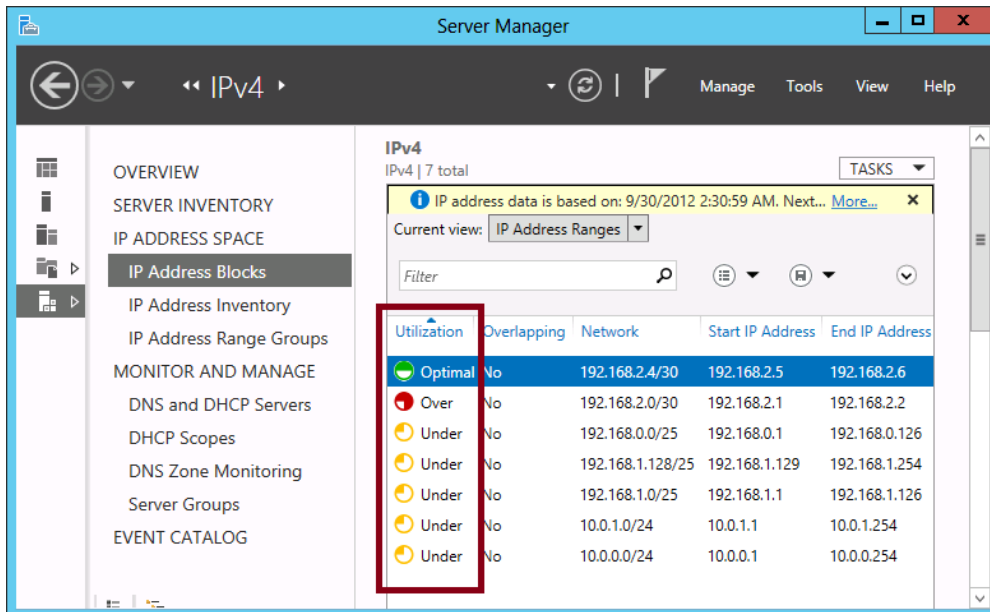


FIGURE 13-20 Utilization rates of IP address ranges

You have the option of altering these 20 percent and 80 percent parameters that separate Under, Optimal, and Over status messages. To do so, in Server Manager, select IP Settings from the Manage menu and then select the Configure Utilization Threshold option in the IPAM Settings dialog box. (You can see these options in Figures 13-14 and 13-15 earlier in this chapter.) In the Configure IP Address Utilization Threshold dialog box that opens, shown in Figure 13-21, adjust the percentage for Under Utilized or Over Utilized, as desired.

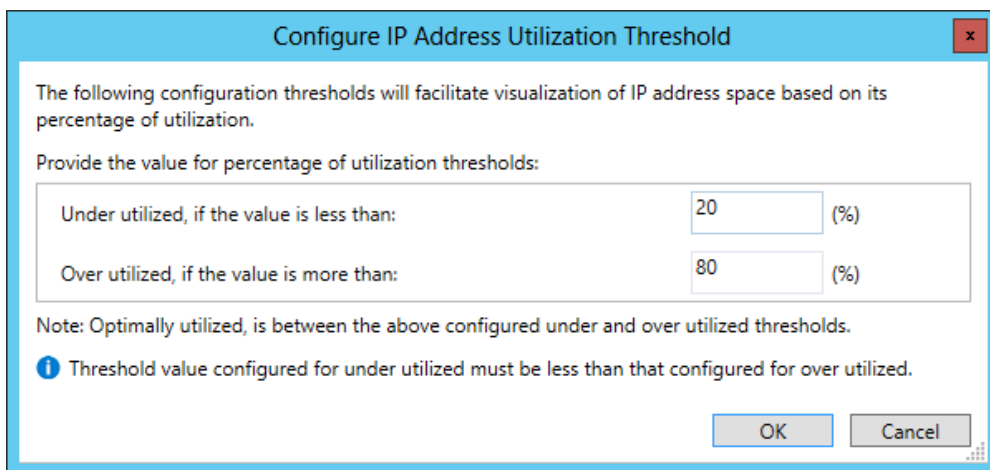


FIGURE 13-21 Changing the overutilization and underutilization thresholds

Delegating IPAM administration

IPAM setup creates on the IPAM server the five local security groups shown in Table 13-2. You can use these groups to delegate particular aspects of IPAM administration to different users.

Make sure you learn these five groups for the 70-417 exam. You might see a question in which each answer choice corresponds to an IPAM security group.



EXAM TIP

If you are a domain user connecting to the IPAM server by using Server Manager, you must be a member of the 'WinRMRemoteWMIUsers__' group on the IPAM server, in addition to being a member of the appropriate IPAM security group listed in Table 13-2.

TABLE 13-2 Local security groups created on the IPAM server

| Group Name | Description |
|------------------------------|---|
| IPAM Users | Members of this group can view all information in server inventory, IP address space, and server management consoles of IPAM. They can view IPAM and DHCP server operational events, but cannot view IP address tracking information. |
| IPAM MSM Administrators | Members of this group have all the privileges of the IPAM Users group and can manage DHCP and DNS server instance-specific information. Such users are Multi Server Management (MSM) Administrators. |
| IPAM ASM Administrators | Members of this group have all the privileges of the IPAM Users group and can perform add and modify address space management operations. Such users are Address Space Management (ASM) Administrators. |
| IPAM IP Audit Administrators | Members of this group have all the privileges of the IPAM Users group and can view IP address tracking information. |
| IPAM Administrators | Members of this group have privileges to view all IPAM information and perform all IPAM tasks. |

MORE INFO For more in-depth knowledge about IPAM, refer to the IPAM Step-by-Step Guide at <http://technet.microsoft.com/en-us/library/hh831622.aspx>, the IPAM Test Lab Guide at <http://www.microsoft.com/en-us/download/details.aspx?id=29020>, or the IPAM virtual lab at <http://go.microsoft.com/?linkid=9838457>. (Microsoft's virtual labs require Internet Explorer.)

Role-based access control for IPAM in Windows Server 2012 R2

The traditional administrator groups that are used for individual computers, domains, and OUs are often not optimal for IP address management functions. At the IP address level, networks are divided up into subnets and DHCP scopes, not into domains and OUs. In addition, IP address management rights within existing administrator groups are highly generalized to functions such as “DHCP Administrators.” What’s needed is a way to assign users very specific IP address management rights to very specific portions of the IP address space. This is the purpose of role-based access control within IPAM in Windows Server 2012 R2.

You can find role-based access control configuration options in Server Manager by selecting Access Control in the IPAM context menu on the left, as shown in Figure 13-22.

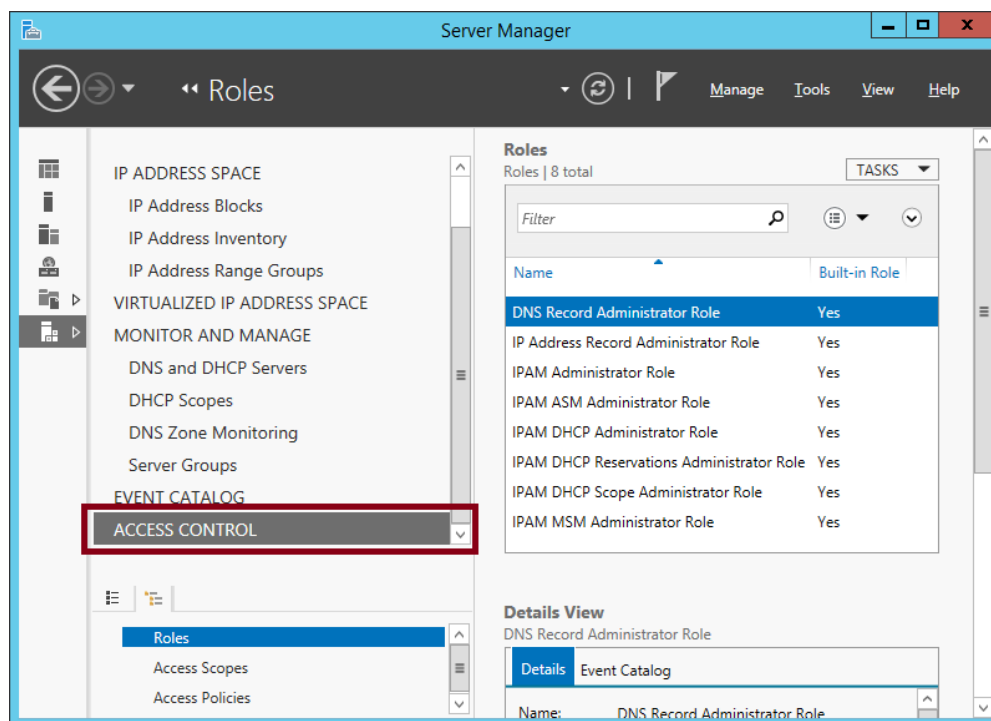


FIGURE 13-22 Role-based access control for IPAM in Windows Server 2012 R2

Role-based access control is built on roles, access policies, and access scopes.

ROLES

Roles are roughly equivalent to administrator groups. The following list shows the eight built-in roles in Windows Server 2012 R2 IPAM and their official descriptions:

- **DNS Record Administrator Role** This built-in role provides permissions to manage the DNS resource records.

- **IP Address Record Administrator Role** This built-in role provides permissions to manage the IP addresses, including finding unallocated IP addresses and creating and deleting IP address instances.
- **IPAM Administrator Role** This built-in role provides all permissions that are provided by the IPAM ASM Administrator Role and IPAM MSM Administrator Role in addition to permissions to manage access scopes, access policies, and logical groups.
- **IPAM ASM Administrator Role** This built-in role provides permissions to completely manage IP address spaces, IP address blocks, IP address subnets, IP address ranges, and IP addresses.
- **IPAM DHCP Administrator Role** This built-in role provides permissions required to completely manage a DHCP server, its associated DHCP scopes, MAC address filters, DHCP policies, and DHCP reservations.
- **IPAM DHCP Reservations Administrator Role** This built-in role provides permissions required to manage the DHCP reservations.
- **IPAM DHCP Scope Administrator Role** This built-in role provides permissions required to manage the DHCP scopes.
- **IPAM MSM Administrator Role** This built-in role provides permissions to completely manage DHCP servers, DHCP superscopes, DHCP scopes, MAC address filters, DHCP policies, DNS zones, and DNS resource records.



EXAM TIP

Make sure you are familiar with these built-in roles for the exam.

ACCESS POLICIES

Unlike with security groups, you don't simply add a user to a role in IPAM. Instead, you use access policies to assign a user or group to one or more roles. Within an access policy, you also have to specify an access scope.

ACCESS SCOPES

Access scopes specify the portion of the address space for which a user or group is assigned to a role, as specified in an access policy. By default, only the Global access scope is created, but you can create other access scopes and assign them to various DHCP scopes in IPAM.

MORE INFO For more information about role-based access control and other new features in Windows Server 2012 R2 IPAM, perform the Windows Server 2012 R2 IPAM walkthrough at <http://technet.microsoft.com/en-us/library/dn268503.aspx>. Another good option is to perform the Windows Server 2012 R2 IPAM virtual lab at <http://go.microsoft.com/?linkid=9839251>. (Microsoft virtual labs require Internet Explorer.)



EXAM TIP

The original Configure Network Services domain on the 70-412 exam includes additional objectives about DHCP and DNS. You might see a question or two about these topics, even though they are not officially listed for the 70-417 exam. For DNS, review the basics about primary and secondary zones, DNS on DCs, zone properties, and basic DNS files such as the zone file and Netlogon.dns. For DHCP, be sure to study a new feature in Windows Server 2012 and Windows Server 2012 R2 called DHCP failover. With DHCP failover, you can configure a second DHCP server in hot standby mode (taking over from the DHCP server in case the main DHCP fails) or load sharing mode (sharing the DHCP requests). For more information about DHCP failover, visit <http://technet.microsoft.com/en-us/library/hh831385.aspx>.

Objective summary

- IPAM is a new feature in Windows Server 2012 and Windows Server 2012 R2 that lets you centrally manage the IP addressing information in your organization. IPAM works only with Microsoft servers in a domain environment.
- To configure IPAM, you run the Provision IPAM Wizard, start a process to discover your infrastructure servers automatically, mark chosen servers as managed, and then run a special cmdlet (Invoke-IpamGpoProvisioning) to create GPOs that automatically configure required settings on those servers.
- DHCP scopes discovered on the network are automatically imported into the IPAM database as IP address ranges. To these IP address ranges you can add larger IP address blocks and individual IP addresses.
- IPAM includes many features for IP address management. These features allow you to describe data in a way that helps you sort and find information about your address space. They also help you keep track of the addresses used in available ranges and update DHCP and DNS servers directly.
- You can delegate aspects of IPAM administration to different users by assigning these users to any of five IPAM security groups.
- Role-based access control is a new IPAM feature in Windows Server 2012 R2. Using this feature, you can assign very specific IP address management tasks to users or groups for a very specific portion of your company's address space.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the "Answers" section at the end of the chapter.

1. You have installed the IPAM feature on a server named IPAM1 that is a member of the Contoso.com domain. You want to configure the IPAM server to retrieve data from all DHCP servers, DNS servers, and domain controllers in the domain.

You choose the option to use the Group Policy–based provisioning method to discover the servers, and you select all three server roles to discover. You start server discovery, and all of the DHCP servers, DNS servers, and domain controllers in the domain are discovered. However, they appear in the Server Manager with an IPAM Access Status of Blocked.

You want the IPAM Access Status of the discovered infrastructure servers to appear as Unblocked. What should you do? (Choose all that apply.)

- A. Run the `Invoke-IpamGpoProvisioning` cmdlet.
 - B. Mark the servers as managed.
 - C. Add the IPAM server to the local Event Log Readers security group.
 - D. Refresh Group Policy on the discovered servers.
2. You work as a network administrator for Fabrikam.com, a company with 2,500 employees and offices in New York, London, Paris, and Munich. Each office site includes its own Active Directory Domain Services domain within the Fabrikam.com forest.

As an administrator, you occasionally need to know the IP address ranges used in various parts of your organization. You want to be able to browse the IP address ranges assigned to each city.

You install IPAM on a server running Windows Server 2012 named IPAM2 in your local office. You then configure IPAM; perform server discovery of the DHCP servers, DNS servers, and domain controllers in your organization; and finally retrieve addressing data from these servers.

Your goal is to be able to use Server Manager to browse the IP address ranges assigned to each city. Which of the following steps do you need to take? (Choose all that apply.)

- A. Configure a custom field.
 - B. Create an IP address range group.
 - C. Assign custom values to IP address ranges.
 - D. Edit the Description field in IP address ranges.
3. You have installed the IPAM feature on a server named IPAM3 that is a member of the Litwareinc.com domain. You want to allow a certain user named Pam to view the IPAM database, server inventory, DHCP events, and IP address tracking.
- You don't want to assign any additional rights to Pam. To which security group on IPAM3 should you add her?
- A. IPAM Users
 - B. IPAM MSM Administrators
 - C. IPAM ASM Administrators
 - D. IPAM IP Audit Administrators



Thought experiment

Deploying and managing IPAM at Adatum

You are a network administrator for Adatum.com. The Adatum.com network is spread out over four buildings on a single campus near Sydney, Australia. The network consists of 1,000 clients and 60 servers, including the following infrastructure servers:

- 12 DHCP servers, including
 - 8 running Windows Server 2012 (four scopes)
 - 2 running Windows Server 2008 R2 (one scope)
 - 1 running Windows Server 2003 R2 (one scope)
 - 1 running CentOS distribution of Linux (one scope)
- 4 domain controllers that are also DNS servers, including
 - 2 running Windows Server 2012
 - 2 running Windows Server 2008 R2
- 2 DNS servers running Debian Linux

All clients and servers on the network that are running Windows are also members of the Adatum.com domain.

As a network administrator, you want to use the IPAM feature in Windows Server 2012 to manage the company's address space. The 12 DHCP servers are used to support seven /24 IPv4 networks within the 10.0.0.0/16 address block. The first 10 addresses in each scope are configured as exclusions and are reserved for static IP assignments. Twelve public servers are hosted on a /28 IPv4 network obtained from the APNIC regional registry. All 12 of these public addresses are statically assigned.

With this information in mind, answer the following questions. You can find the answers to these questions in the "Answers" section.

1. How many of the infrastructure servers are not compatible with IPAM?
2. Assuming you want to add all of the organization's public and private addresses to the IPAM database, how many IP address blocks should you add? How many IP address ranges must be added manually?
3. Server Manager displays the public IP address range utilization as "Over." Your organization isn't intending to assign any more public IP addresses to servers. What step can you take to allow your IP ranges to use 90 percent of the available addresses before displaying the "Over" status?
4. You want to assign a static IP address to a new device in the logical subnet that contains the DHCP server running Windows Server 2012. What is the most efficient way to discover an unused address in this range?

Answers

This section contains the answers to the Objective Review and the Thought Experiment.

Objective 13.1: Review

1. **Correct answers:** A, B, D
 - A. **Correct:** You use this cmdlet to create the GPOs needed to configure discovered servers for IPAM management.
 - B. **Correct:** The GPOs created by running the Invoke-IpamGpoProvisioning cmdlet apply only to servers you mark as Managed in IPAM.
 - C. **Incorrect:** This setting needs to be configured on the domain controller you want to manage in IPAM, but it is automatically configured by the GPO created by the Invoke-IpamGpoProvisioning cmdlet. You therefore don't need to take this step.
 - D. **Correct:** After running the Invoke-IpamGpoProvisioning cmdlet, you need to refresh Group Policy on the discovered infrastructure servers so that the settings in the newly created GPOs are applied to these servers.

2. **Correct answers:** A, B, C
 - A. **Correct:** You need to configure a custom field that includes the values New York, London, Paris, or Munich. You can create a new custom field for these values and name it City, or you can use an existing custom field such as Location and add the four cities as new custom values.
 - B. **Correct:** An IP address range group allows you to sort and browse your IP address ranges by any field you choose.
 - C. **Correct:** You need to tag each IP address range with the City (or Location) value of New York, London, Paris, or Munich.
 - D. **Incorrect:** You cannot browse IP addresses based on the values included in the Description tag. To accomplish that, you need an IP address range group.

3. **Correct answer:** D
 - A. **Incorrect:** The IPAM Users group will assign all the rights Pam needs except for the ability to use IP address tracking.
 - B. **Incorrect:** The IPAM MSM Administrators group doesn't provide the ability to use IP address tracking. It also provides unnecessary rights for managing DHCP and DNS servers.
 - C. **Incorrect:** The IPAM ASM Administrators group doesn't provide the ability to use IP address tracking. In addition, it provides the unnecessary right to perform add and modify address space management operations
 - D. **Correct:** The IPAM IP Audit Administrators group provides exactly the user rights (and no others) that are required by the scenario.

Thought experiment

1. Four. One running Windows Server 2003 R2 and three running Linux.
2. You should add two IP address blocks, one for the 10.0.0.0/16 private address block and one for the public /28 address block. You need to add three IP address ranges manually, one for the public address range, one for the scope hosted on the Windows Server 2003 R2 server, and one for the scope hosted on the CentOS Linux server. (The others are imported automatically.)
3. Change the Over Utilized utilization threshold in IPAM Settings to 90 percent.
4. Use the Find And Allocate Available IP Addresses function in IPAM.

This page intentionally left blank

Configure identity and access solutions

In its original form on the 70-412 exam, the Configure Identity and Access Solutions content area covers three general topics: Active Directory Federation Services (AD FS), Active Directory Certificate Services (AD CS), and Active Directory Rights Management Services (AD RMS). However, of these three, just AD FS has been tagged as a 70-417 exam topic.

The purpose of AD FS is to use claims (described in Chapter 11) to securely extend the reach of Active Directory authentication across multiple networks and platforms. One major benefit of extending authentication in this way is to provide users with single-sign-on (SSO) access to applications such as cloud-hosted applications that are hosted outside of company premises.

AD FS works by taking claims about users from a provider such as Active Directory Domain Services (AD DS), packaging these claims into security tokens, and then issuing the tokens to another security provider or application (called a “relying party”). The tokens can conform to various industry-standard protocols, including WS-Trust, WS-Federation, and Security Assertion Markup Language (SAML) 2.0. AD FS can thus create and issue tokens containing claims information for almost any application that requests them, and as a result, the server role can be used in a wide range of contexts.

AD FS has undergone a lot of changes in recent releases, and this trend has continued in Windows Server 2012 R2. Additionally, some of the most important new features in Windows Server 2012 R2, such as Workplace Join and multi-factor authentication, rely on AD FS. You should expect, therefore, to see plenty of questions about AD FS on the 70-417 exam.

Objectives in this chapter:

- Objective 14.1: Implement Active Directory Federation Services (AD FS)

Objective 14.1: Implement Active Directory Federation Services (AD FS)

AD FS is used to authenticate Active Directory users to resources that are located outside of the users’ Active Directory forest. AD FS helps perform this procedure by issuing user-authenticating tokens for resources configured to request them. The tokens that AD FS issues can be created in

a number of requested formats, but they all contain validation of a user's identity along with other user information such as group memberships and other attributes stored in claims. Active Directory normally serves as the original source of users and claims about them that AD FS packages into the various token types. Besides assisting initial authentication, AD FS tokens also enable web single-sign-on (SSO) to applications across different organizations and platforms.

For the 70-417 exam, you won't need to understand AD FS in great depth, but you should expect to see a number of questions about various AD FS-related topics on the exam. Therefore, make sure you study this material well and perform the home and virtual labs suggested at the end of the chapter.

This section covers the following topics:

- AD FS scenarios
- How AD FS works
- Initial configuration of AD FS
- Configuring Relying Party Trusts
- Configuring Workplace Join
- Configuring multi-factor authentication

AD FS scenarios

AD FS is most often used to authenticate Active Directory users and provide these users with web-based SSO to resources outside of their Active Directory environment. For example, AD FS supports SSO to enterprise applications that are hosted in the cloud. A similar example appears with AD FS when it is used to support authentication and web-based SSO to resources in a partner organization.

However, AD FS can also be used to authenticate Active Directory users within their own Active Directory forest. This scenario could be used to support employees connecting from outside the company network and provide them with web-based SSO to resources located on company premises. But even more broadly, AD FS can be used to support authentication to virtually any claims-based application, regardless of whether that application is located on company premises or off premises.

IMPORTANT It is considered best practice not to expose an AD FS server directly to the Internet. When AD FS in Windows Server 2012 R2 needs to support users through a standard Internet connection, it can do so through a new Remote Access role service called *Web Application Proxy*. The Web Application Proxy server is accessible from the Internet and can be configured to communicate with an internal AD FS server on behalf of an external user. In the first release of Windows Server 2012, this role is performed by an AD FS role component called AD FS Proxy. AD FS Proxy is not available in Windows Server 2012 R2.

MORE INFO For more information about how to use Web Application Proxy to connect users to applications, visit "Walkthrough Guide: Connect to Applications and Services from Anywhere with Web Application Proxy" at <http://technet.microsoft.com/en-us/library/dn280943.aspx>.

How AD FS works

The process of authenticating a user through AD FS is illustrated in Figure 14-1 and Figure 14-2. For simplicity, the scenario assumes no Web Application Proxy is being used. To conform to best practices without using a Web Application Proxy, the connection between the user and the AD FS server should be established over a VPN.

The process begins in Figure 14-1 and Step 1, when a user for the first time attempts to connect to a particular claims-aware application that has been configured to trust AD FS. In Step 2, the application requests that the user obtain a digitally signed user token from AD FS. In Step 3, the user is redirected to the AD FS server in order to request the token. In Step 4, AD FS contacts AD DS to obtain information about the user.




FIGURE 14-1 Authentication through AD FS, part 1

The final steps are illustrated in Figure 14-2. In Step 5, AD DS answers the request for information about the user—information that AD FS then packages into a token and digitally signs. In Step 6, AD FS provides the user with the requested token. In Step 7, the user provides the application with the same requested token. If the user is allowed access to the application, the user can then successfully sign on.

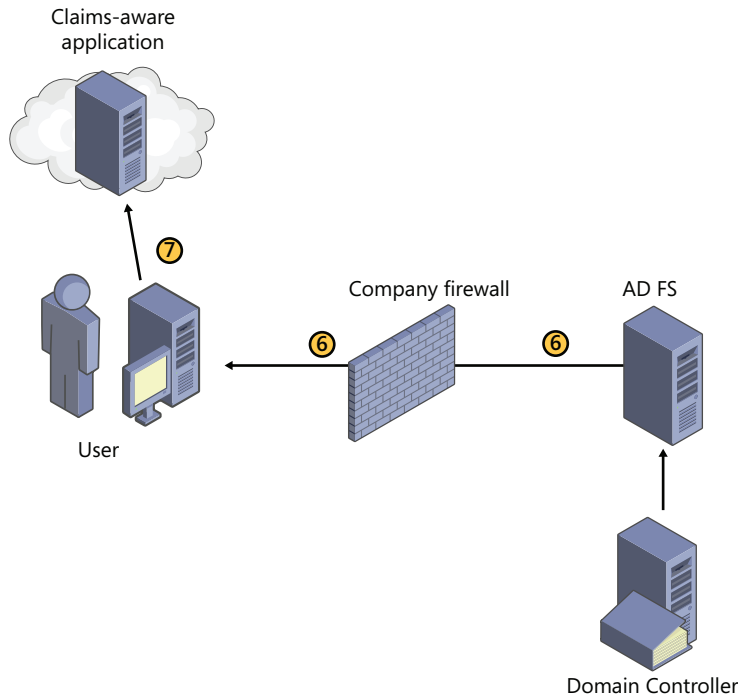


FIGURE 14-2 Authentication through AD FS, part 2

Active Directory Federation Server Configuration Wizard

The AD FS management console installs automatically when you install the AD FS server role in Windows Server 2012 and Windows Server 2012 R2, but the console tree remains unpopulated until you run the Active Directory Federation Server Configuration Wizard. You can access this wizard by clicking the notification flag in Server Manager immediately after you install the AD FS server role and then selecting the option to configure AD FS.

Prerequisite #1: Install an SSL certificate on the AD FS server

Before you run the Active Directory Federation Server Configuration Wizard, you need to install a Secure Sockets Layer (SSL) certificate in the Personal certificate store for the computer account of the local computer. This certificate is used for AD FS service communications. In the first release of Windows Server 2012, the same certificate also needs to be installed in the Default Web Site of the local instance of Internet Information Services (IIS). However, in

Windows Server 2012 R2, IIS is no longer needed for AD FS and is not installed automatically when you install the AD FS server role.

The SSL certificate must meet the following criteria:

- The subject name of the certificate should be a Common Name type and set to the name of your ADFS server, such as adfs1.contoso.com. (You must also specify this same name on the certificate for a Web Application Proxy server, if you later deploy one.)
- You must add a DNS-type alternative name that specifies the same ADFS server name, such as adfs1.contoso.com.
- If you use a separate DNS namespace internally within your organization (such as contoso.local or corp.contoso.com), you should add this internal DNS name as a DNS-type alternative, such as adfs1.corp.contoso.com.
- If you might ever want to use the Workplace Join feature, which is new to Windows Server 2012 R2, you need to configure an additional DNS-type alternative name named "enterpriseregistration" followed by the UPN suffix (domain name) of the domain, as in, for example, enterpriseregistration.contoso.com.

Figure 14-3 shows the configuration of the certificate request to AD CS for an AD FS server.

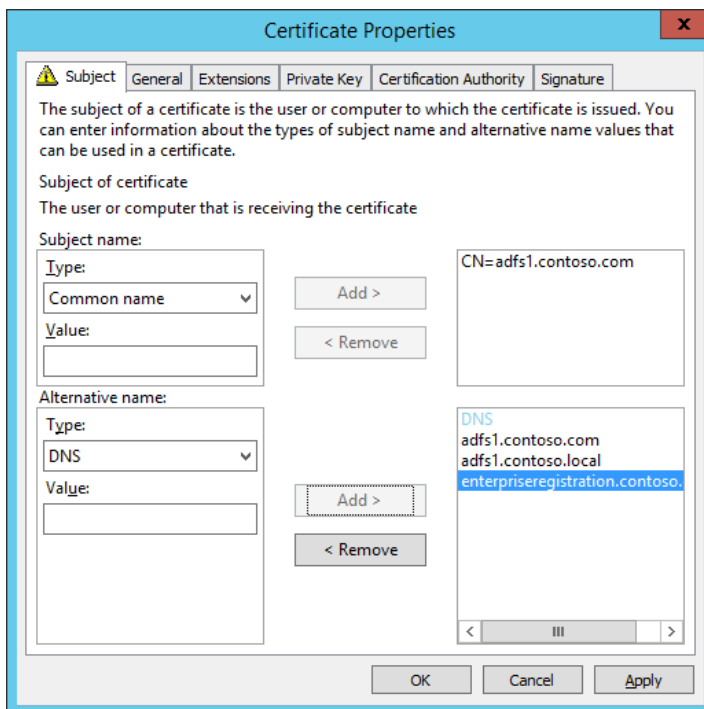


FIGURE 14-3 Configuring the common name and alternative name for the AD FS service communications certificate

Note that in many production scenarios, it's preferable for this SSL certificate to be issued by a public CA instead of from AD CS. Even so, for the exam, you need to remember how to configure such an SSL certificate through AD CS by using the Web Server template, and also how to enroll and install the certificate on another server (here, the AD FS server). These methods haven't changed since Windows Server 2008, but if you need a refresher, search the TechNet website for the article named "Configure SSL/TLS on a Web Site in the Domain with an Enterprise CA."



EXAM TIP

Although Certificate Services is not included in the objectives of this domain for the 70-417 exam, it's difficult to imagine you won't see at least one question on the test related to this topic. If your knowledge of Certificate Services is rusty, be sure to review features such as enrollment, auto-enrollment, user and computer templates, and public key infrastructure (PKI) components.

Prerequisite #2: Create the KDS root key for group Managed Service Accounts

The Active Directory Federation Server Configuration Wizard requires you to create a new group Managed Service Account (gMSA) or to specify an existing gMSA to be used as the service account for AD FS. A feature first introduced in Windows Server 2012, gMSAs are service accounts that can be used on multiple computers in a domain. This type of account is useful for services such as AD FS that can be configured on multiple servers in a multi-server farm. All the servers in the farm run AD FS in the context of the same gMSA and use the same (and only) password for that gMSA.

You can create the AD FS gMSA account as part of the Active Directory Federation Server Configuration Wizard, but you need to take a preparatory step beforehand. Before you can create any gMSA in a domain, you first need to create a KDS root key for that domain. This KDS root key is stored on all domain controllers and is used to automatically generate passwords for gMSAs.

To create the KDS root key for the domain and allow gMSAs to be created, run the following command at a Windows PowerShell prompt:

```
Add-KdsRootKey -EffectiveImmediately
```

You then need to wait 10 hours for the new KDS root key to propagate to all domain controllers in the domain before you can create the gMSA. In a test environment with one domain controller, however, you can work around this requirement by running the following version of the command at a Windows PowerShell prompt. This version allows you to create a gMSA immediately:

```
Add-KdsRootKey -EffectiveTime ((get-date).addhours(-10))
```



EXAM TIP

Remember the importance of the `Add-KdsRootKey` cmdlet for the exam.

AD FS management console

Figure 14-4 shows the new AD FS management console in Windows Server 2012 R2 after you run the AD FS Federation Server Configuration Wizard. The AD FS management console tree in Windows Server 2012 R2 includes three main nodes: the Service node, the Trust Relationships node, and the Authentication Policies node.

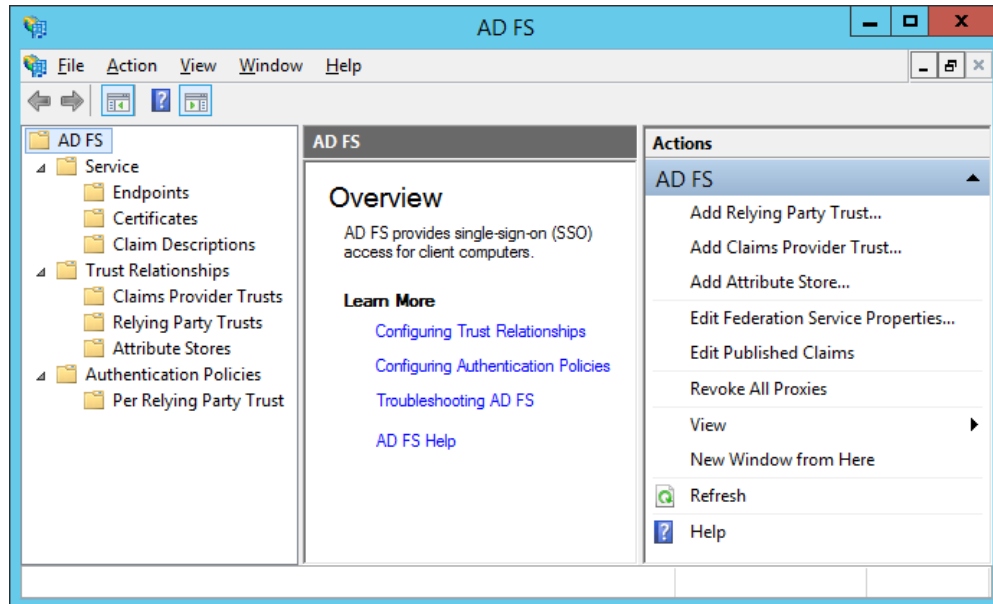


FIGURE 14-4 The AD FS management console in Windows Server 2012 R2

Service node

You use the Service node in the AD FS management console for the following general purposes:

- Managing the services provided through endpoints
- Configuring the certificates used for issuing and receiving tokens and publishing metadata
- Configuring claims types

ENDPOINTS

In AD FS, endpoints are URLs that provide access to various services, such as token issuance and the publishing of federation metadata. You use the Endpoints node to enable or disable various endpoints, and to control whether the endpoint is published to federation server proxies.

As shown in Figure 14-5, when the Endpoints node is selected, the details pane reveals a Token Issuance section and a Metadata section. The Metadata section includes the path to the local federation server's metadata XML file. This metadata file defines the data format for communicating configuration information between a claims provider (such as AD DS) and a relying party. You might need to provide this address or this XML file to a resource partner when you are establishing a federated trust.



EXAM TIP

You need to know that the Endpoint node contains a path to the federation metadata of your organization. You would provide this data to a partner organization to configure a federated trust.

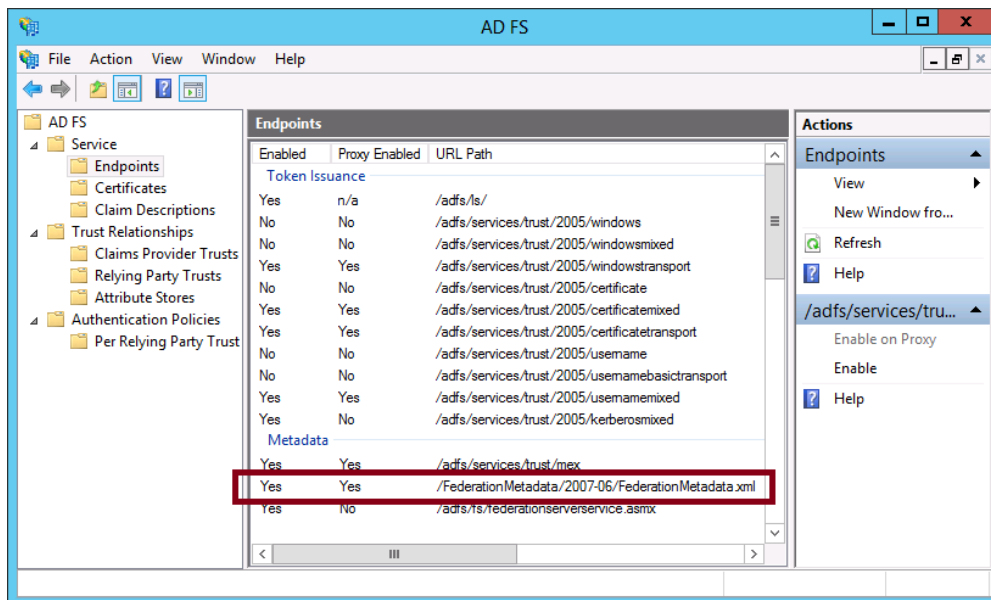


FIGURE 14-5 The Endpoints node and Federation Metadata file

CERTIFICATES

The Certificates node in the AD FS management console shows the three certificates you need to run AD FS. All three of these certificates are based on the Web Server template.

- **Service Communications certificate** This is the certificate that you need to install manually in the Personal store of the local computer account before you can configure AD FS. It is used for Windows Communication Foundation (WCF) message security.
- **Token-Decrypting certificate** This certificate is used to decrypt tokens that AD FS receives. It is created by default as a self-signed certificate, but you can replace it with a certificate based issued by an enterprise certificate authority (CA).
- **Token-Signing certificate** This certificate is used to sign tokens that AD FS issues. As with the Token-decrypting certificate, this certificate is created by default as a self-signed certificate, but you can replace it with one issued by an enterprise CA.

The Certificates node and the three AD FS certificates are shown in Figure 14-6.

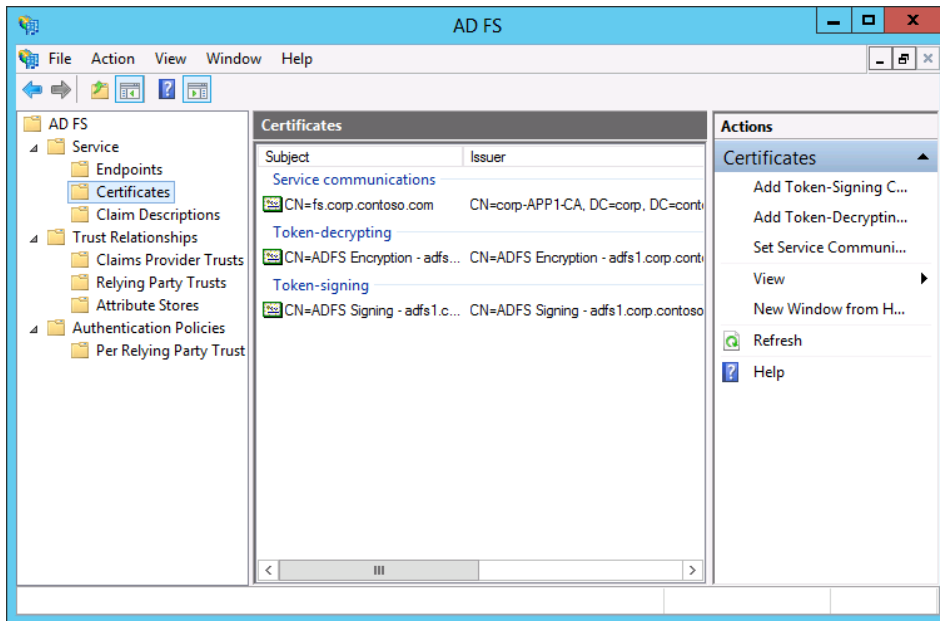


FIGURE 14-6 The three AD FS server certificates

CLAIM DESCRIPTIONS

Through this node you can find information about the claims the local server is currently able to assert about identities. The description of each claim is found at a specified URL. More claims can be added by clicking Add Claim Description in the Actions pane.

Trust Relationships node

The Trust Relationships node in the AD FS management console allows you to manage the trust relationships of AD FS. The three child nodes allow you to add and configure policies for claims providers, relying parties, and attribute stores.

CLAIMS PROVIDER TRUSTS

By default, AD FS automatically trusts Active Directory as a claims provider. To add another claims provider, you can select the Claims Provider Trusts node in the AD FS console tree and then click Add Claims Provider Trust in the Actions pane.

You can also opt to add or edit the claim rules that are configured for Active Directory or another claims provider. These claim rules are called acceptance transform rules; their function is to take claims from the claims provider (usually AD DS) and repackage them for an external relying party.



EXAM TIP

Remember that acceptance transform rules apply to claims provider trusts.

ADD A RELYING PARTY TRUST

After you complete the AD FS Configuration Wizard, you should add one or more relying party trusts by right-clicking the Trust Relationships node and selecting Add Relying Party Trust from the shortcut menu. A relying party refers to an application or other agent that “consumes” (requests and uses) security tokens issued by the local Active Directory Federation Service service. When you click this option and run the Add Relying Party Trust Wizard that opens, you configure AD FS to issue these tokens for a specific application. The relying party trust determines the form of the claims and tokens issued to the application.

To complete the wizard, you need to enter data about the application, either manually or by specifying or pointing to a federation metadata file provided to you or the owner of the application. (Specifying or pointing to a federation metadata file is recommended, as opposed to entering the data manually.) The federation metadata contains the data format used to communicate between AD FS and the relying party.

As part of the Add Relying Party Trust Wizard, you can also choose to require multifactor authentication for the trust. You also have the option of configuring multifactor authentication in the AD FS management console, as part of Authentication Policies. Through Authentication Policies you have the additional option of configuring multifactor authentication universally for all relying party trusts, not just for one specific relying party trust.

CONFIGURING CLAIMS RULES FOR RELYING PARTY TRUSTS

After you complete the Add Relying Party Trust Wizard, you are given an opportunity to configure claims rules for the trust. Claims rules occur in three types: Issuance Transform Rules, Issuance Authorization Rules, and Delegation Authorization Rules.

- **Issuance Transform Rules** These rules specify claims that will be sent to the relying party.
- **Issuance Authorization Rules** These rules specify which users are permitted access to the relying party.

- **Delegation Authorization Rules** These rules specify which users can act as delegates for other users to the relying party.

Claims rules are actually statements written in a special Claim Rule Language. Here's an example of the syntax of this language:

```
c:[Type == "http://contoso.com/department"]
=>issue(Type = "http://adatum.com/department", Value = c.Value);
```

These lines take a specific claim at *http://contoso.com/department* and issue the same claim unchanged to *http://adatum.com/department*.



EXAM TIP

Although it's unlikely that you'll need to understand Claim Rule Language syntax for the 70-417 exam, you should at least know and understand the difference among the three claim rule types for relying party trusts.

MORE INFO For more information about Claim Rule Language, see "Understanding Claim Rule Language in AD FS 2.0 & Higher" at <http://social.technet.microsoft.com/wiki/contents/articles/4792.understanding-claim-rule-language-in-ad-fs-2-0-higher.aspx>.

Authentication Policies node

The Authentication Policies node is new in Windows Server 2012 R2. You use the Authentication Policies node to configure user authentication methods either globally for all relying parties, or specifically to an individual relying party. *Primary authentication* determines which single-factor authentication methods are allowed for intranet or extranet users. You can also require *multi-factor authentication* for specific users or groups who connect from certain locations or from certain types of devices. With multi-factor authentication, a combination of authentication methods is required.

GLOBAL POLICIES

You configure global primary and multi-factor authentication through different tabs of the Edit Global Authentication Policy dialog box. You can open this dialog box in a number of ways when Authentication Policies is selected, such as by clicking the Edit link in the Primary Authentication area. This option is shown in Figure 14-7.

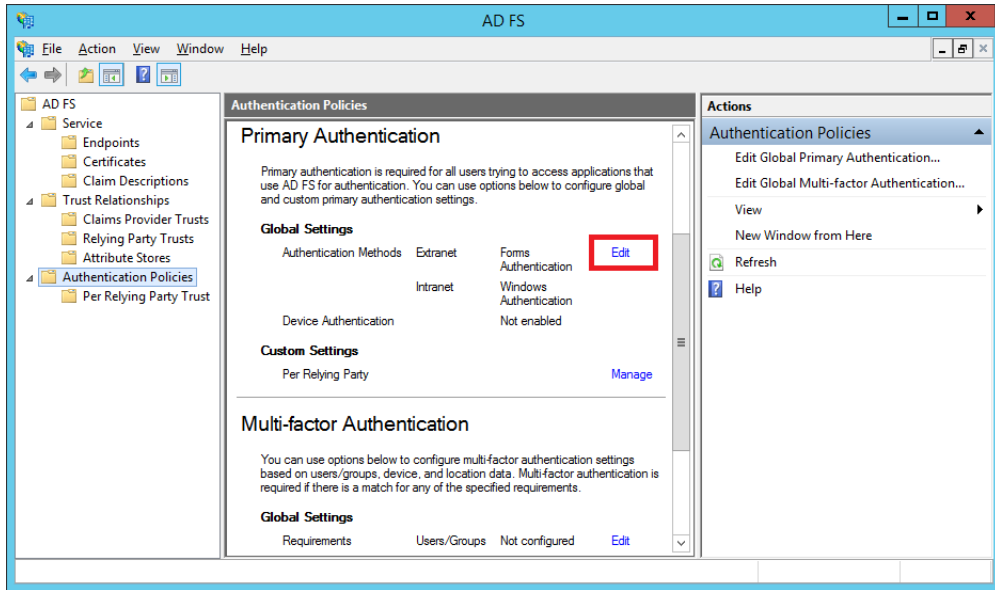


FIGURE 14-7 Editing global primary authentication settings

The Primary tab of the Edit Global Authentication Policies dialog box is shown with the default options in Figure 14-8. By default, users connecting from the extranet are authenticated through Forms Authentication (in other words, a form on a webpage), and users connecting from the intranet are authenticated through Windows authentication.

There's also a third authentication option that is not selected by default for either extranet or intranet users: Certificate Authentication. This option would allow a user to authenticate through a user certificate, typically provided in the form of a smart card. You can also extend the authentication options through third-party products or in-house software development.

If you select more than one option on the Primary tab for either the extranet location or the intranet location, you provide users with a choice of authenticating through any of those selected methods. (In other words, it does not require users to authenticate through all selected methods.)

Finally, the Enable Device Authentication option at the bottom of the Primary tab relates to Workplace Join, a new feature in Windows Server 2012 R2 that is described in more detail later in this chapter.




FIGURE 14-8 The Primary tab of the Edit Global Authentication Policy dialog box

The Multi-Factor tab of the Edit Global Authentication Policy dialog box is shown in Figure 14-9. This tab allows you to require certain users to provide more than one form of authentication. As the tab shows, you can apply the multi-factor authentication requirement to specific users or groups, to devices that are either registered or unregistered (workplace-joined or not), and to connections that originate from either the extranet or the intranet. By default, the only secondary authentication method available is Certificate Authentication. However, these options may be extended through third-party options or in-house development.

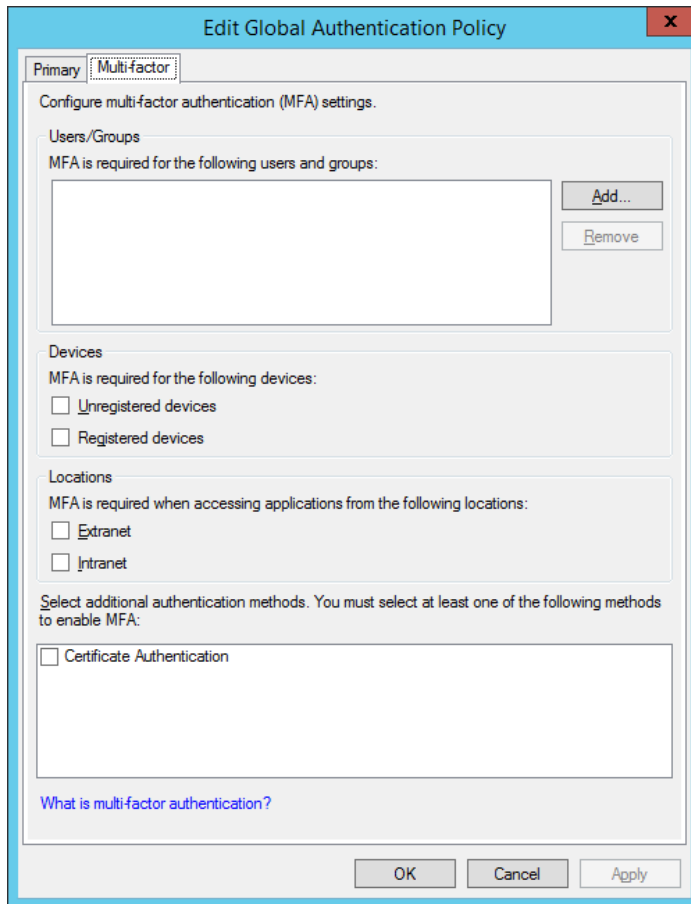


FIGURE 14-9 Configuring global multi-factor authentication triggers

PER RELYING PARTY TRUST POLICIES

You can use the Per Relying Party Trust node beneath Authentication Policies to configure additional primary or multi-factor authentication settings for a particular relying party trust. To set additional primary authentication methods, select the Per Relying Party Trust node, right-click the desired relying party in the center pane, and then select Edit Custom Primary Authentication. The primary authentication methods you set on a per-trust basis are used as additional primary authentication options to any that you have set at the global level.

You can also set additional triggers for multi-factor authentication on a per-relying-trust basis. For example, at the global level, you might require multi-factor authentication for the Finance group. Then, specifically for a relying party called RP1, you might require multi-factor authentication for the Sales group. In that case, whenever a user from *either* Finance or Sales is redirected from RP1 to obtain a user token, that user must provide multi-factor authentication. To set additional multi-factor authentication triggers for a particular relying party,

select the Per Relying Party Trust node, right-click the desired relying party in the center pane, and then select Edit Custom Multi-factor Authentication.



EXAM TIP

Extended Protection for Authentication is a feature that reduces the possibility of man-in-the-middle attacks between clients and servers over HTTP. By default, AD FS in Windows Server 2012 and Windows Server 2012 R2 allows this more secure form of communication when it is initiated by clients, but AD FS does not require it.

Although Extended Protection for Authentication improves security, there's an important drawback you need to know about it: Of the major web browsers, only Internet Explorer 8 and later supports the feature by default. As a result, if an AD FS server is configured to require Extended Protection for Authentication, users connecting to it through third-party browsers will normally experience access denied messages. To revert to the default setting on the AD FS server, type the following command at an elevated Windows PowerShell prompt:

```
Set-ADFSProperties -ExtendedProtectionTokenCheck Allow
```

To turn off Extended Protection for Authentication completely, type the following at an elevated Windows PowerShell prompt:

```
Set-ADFSProperties -ExtendedProtectionTokenCheck None
```

For more information about this issue related to Extended Protection for Authentication and third-party browsers, visit "Plan for single sign-on for Office 365 with third-party browsers" at <http://technet.microsoft.com/en-us/library/hh852537.aspx>.

Workplace Join

The bring-your-own-device (BYOD) trend reflects a major development that's been underway in IT since the rise of personal mobile devices. Before BYOD, employees typically did their work either on desktops or on laptops, both of which were company-owned. And if the company's private network included Active Directory, these desktops and laptops were almost always domain-joined.

Personal mobile phones and tablets have since exploded in popularity. Many employees—including executives—prefer to have the option to work on these personal devices instead of on domain-joined, company-owned computers. With this BYOD trend, IT departments have to find a way to provide secure access to company resources from unmanaged personal devices.

Windows Server 2012 R2 introduces a new feature called Workplace Join whose purpose is to help organizations handle BYOD more securely. Workplace Join allows users to register personal devices in Active Directory without joining the devices to the domain. A registered device cannot be managed or controlled by IT. Group Policy does not apply to workplace-joined devices, and no account exists in Active Directory for the device. However, the registered status of a user's personal device can be used to determine whether that user can access given resources.

As shown in Figure 14-9, you can require users connecting only from unregistered devices to provide secondary authentication. This requirement effectively means that Workplace Join can act as a transparent, secondary authentication factor, depending on how you configure authorization policies. If no other secondary authentication factor is even available, then Workplace Join can also effectively act as an authorization *requirement* for specified users, such as extranet users.

At the time of this writing, Workplace Join supports computers and devices running Windows 8.1, Windows RT 8.1, and iOS. Devices and computers running these operating systems can all be registered or “workplace-joined” to a domain. (Android devices are not yet supported.)

How Workplace Join works

Workplace Join relies on a service called the Device Registration Service (DRS), which is a part of AD FS in Windows Server 2012 R2. Users begin the process of workplace-joining their devices by signing in to their workplace through PC Settings in Windows 8.1 or through a special website on the corporate network. At this point, DRS creates a device object in Active Directory used to represent the device and sets a certificate on the consumer device.

Configuring Workplace Join

As mentioned earlier in this chapter, if you want to allow AD FS to support the Workplace Join feature, you first have to configure the main AD FS certificate in a particular way. Specifically, you need to assign the certificate a DNS-type alternative name of `enterpriseregistration` followed by the UPN suffix of the domain, as in `enterpriseregistration.contoso.com`. The configuration of the request for such a certificate is shown in Figure 14-3.

Assuming you have properly configured this certificate before you first configured AD FS, you then need to perform the following five steps to configure Workplace Join:

1. In internal DNS, add a CNAME (alias) named `enterpriseregistration` that points to the AD FS server.
2. Enable device registration in Active Directory by typing the following two lines at an elevated Windows PowerShell prompt on the AD FS server:

```
Initialize-ADDeviceRegistration -ServiceAccountName Domain\gMSAname$  
Enable-AdfsDeviceRegistration
```

where *Domain* is the NetBIOS name of the domain and *gMSAname\$* is the name of the gMSA you have assigned to AD FS followed by a dollar sign (\$). For example, if the Active Directory domain name is `Contoso.com` and the gMSA is `FsGmsa`, enter the following two lines, one at a time:

```
Initialize-ADDeviceRegistration -ServiceAccountName Contoso\FsGmsa$  
Enable-AdfsDeviceRegistration
```

3. Enable device registration in AD FS Authentication Policies:

- A. In AD FS management console, select Authentication Policies and then click Edit Global Primary Authentication in the Actions pane.
- B. In the Edit Global Primary Authentication dialog box, on the Primary tab, check Enable Device Authentication, as shown in Figure 14-10.

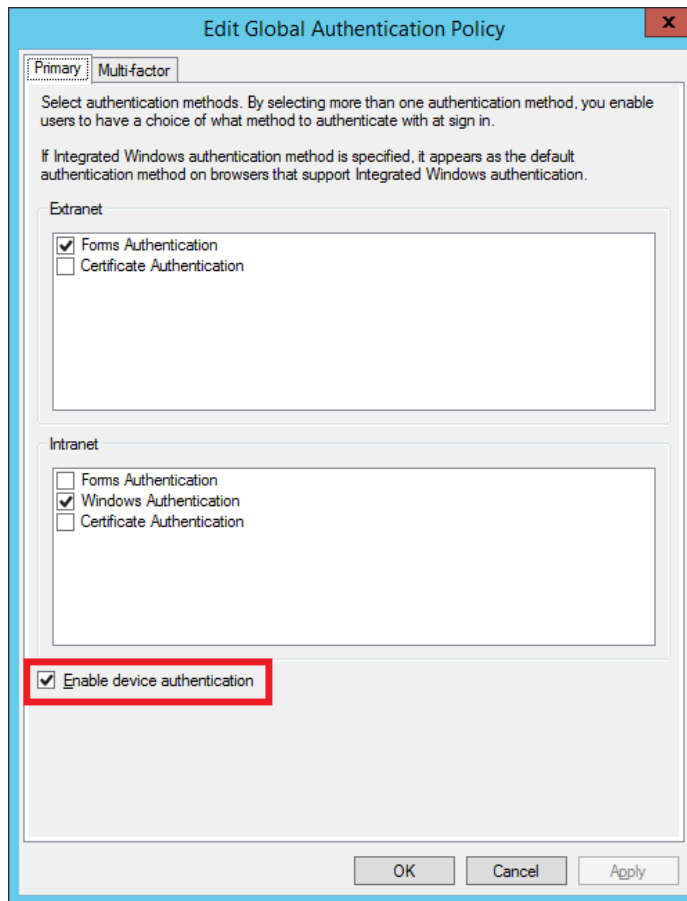


FIGURE 14-10 Enabling device registration in AD FS authentication policies

4. Configure every client that will be workplace-joined to trust the CA root certificate of the CA that has issued the AD FS service communication certificate. (If this AD FS certificate was issued by a trusted public CA, this step is not necessary.)
5. Perform the workplace-join procedure on each client.

If a client is running iOS, navigate to <https://adfsservername/enrollmentserver/otaprofile>, where *adfsservername* is the full external DNS name of the AD FS server or proxy. Then type a user's ID in email address format ("user@contoso.com") and click Join.

If the client is running Windows 8.1, perform the following steps:

- A. On the Charms bar, click Settings and then Change PC Settings.
- B. In PC Settings, click Network and then click Workplace.
- C. In Workplace, type a user's ID in email address format ("user@contoso.com") and then click Join.

After you complete this last step, a new object appears in Active Directory within a RegisteredDevices container, as shown in Figure 14-11.

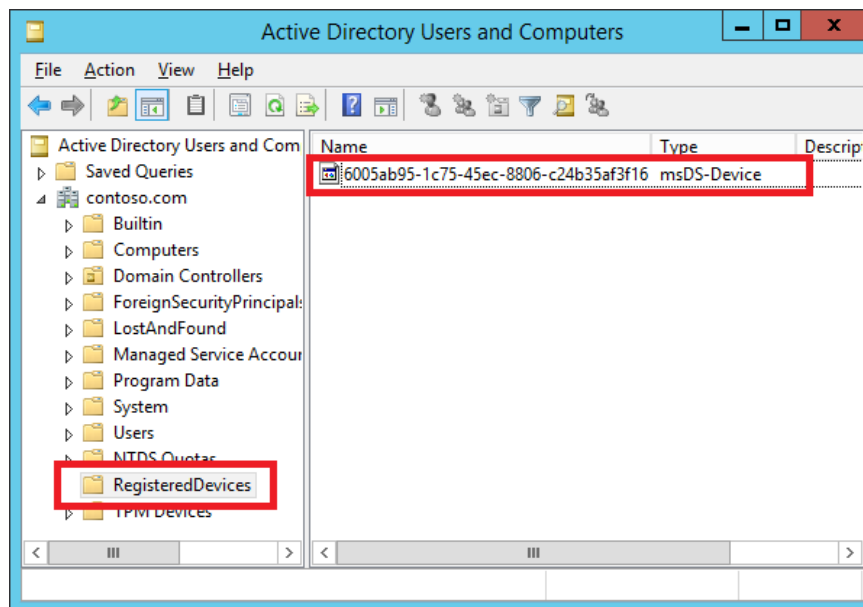


FIGURE 14-11 A workplace-joined device in Active Directory



EXAM TIP

Expect to see more than one question about Workplace Join on the exam. You need to remember the exact requirements on the AD FS server certificate, along with all 5 steps mentioned above, including the names of required cmdlets and the specific dialog box required. You should also understand how you would configure multifactor authentication policies in Figure 14-9 if you wanted to *require* devices connecting from the extranet to be workplace-joined.

Windows PowerShell cmdlets for AD FS

Windows Server 2012 and Windows Server 2012 R2 include an ADFS module for Windows PowerShell that includes over 100 cmdlets. These cmdlets were not available out of the box in Windows Server 2008 or Windows Server 2008 R2. It's possible that you could see a question about AD FS that includes one or more of these Windows PowerShell cmdlets. It's

therefore worth browsing the list of available AD FS cmdlets by using the following command:
Get-Command -Module ADFS.

IMPORTANT It's highly recommended that you enrich your study of AD FS by performing labs and walkthroughs. For example, you can use the following link to connect to a Tech-Net Virtual Lab that demonstrates the new features in AD FS in Windows Server 2012 R2: <http://go.microsoft.com/?linkid=9842896>.

An even better option is to perform an AD FS lab in your own virtual environment. Navigate to the following address to find instructions for a lab environment setup along with various AD FS walkthroughs: <http://technet.microsoft.com/en-us/library/dn280939.aspx>. (Microsoft's virtual labs require Internet Explorer.)

Objective summary

- The general purpose of Active Directory Federation Services (AD FS) is to authenticate Active Directory users outside of Active Directory. AD FS works through claims-based authentication. If a claims-aware resource is configured to trust an AD FS server, that resource directs a user to obtain from AD FS a token that includes claims about that user. The user then returns to the resource with the token and the user is granted or denied access based on the claims contained in the token.
- Before you can configure AD FS by using the AD FS Configuration Wizard, you need to install an SSL certificate in the Personal store of the computer account on the AD FS server. You also need to create the KDS root key for the domain by using the Add-KdsRootKey cmdlet in Windows PowerShell.
- A relying party is a claims-aware application or resource that is configured to use AD FS for authentication. To configure AD FS properly for a particular relying party, you need to create a relying party trust. You can do this by using the Add Relying Party Trust Wizard. When configuring the relying party trust, it's recommended that you specify a federation metadata file that contains all the information AD FS needs about the format of the tokens and claims required by the relying party.
- Authentication policies are used to set the primary (single-factor) authentication options and any multi-factor authentication requirements. Global authentication policies apply to all relying party trusts. Besides these global authentication policies, you can also configure additional primary authentication options and additional multi-factor authentication triggers for a specific relying party.
- Workplace Join is a new feature in Windows Server 2012 R2 that allows users to register their devices in Active Directory without joining a domain. The device is not managed or controlled centrally. However, the registration status of a device can act as a trigger for multi-factor authentication. "Workplace-joining" a device can effectively

serve as a prerequisite for authentication, depending on how you configure authentication policies.

- To prepare for the 70-417 exam, it's recommended that you perform hands-on exercises and walkthroughs that demonstrate new AD FS features.

Objective review

Answer the following questions to test your knowledge of the information in this objective. You can find the answers to these questions and explanations of why each answer choice is correct or incorrect in the "Answers" section at the end of the chapter.

1. You are establishing a federated trust with a partner organization. An IT administrator at the partner administration asks you to send her your federation metadata XML file. Your AD FS is running Windows Server 2012. In which of the following nodes in the AD FS management console would you be able to determine the location of the metadata file?
 - A. Endpoints
 - B. Claims Provider Trusts
 - C. Relying Party Trusts
 - D. Attribute Stores
2. You want to replace the default Token-decrypting certificate in AD FS with one issued by your enterprise CA. You request a certificate and are prompted to select a certificate template. Which certificate template do you choose?
 - A. Computer
 - B. Domain Controller Authentication
 - C. Kerberos Authentication
 - D. Web Server
3. You want to replace the default Token-signing certificate in AD FS with one issued by your enterprise CA. Into which certificate store should you import the certificate?
 - A. Current User\Personal
 - B. Current User\Enterprise Trust
 - C. Local Computer\Personal
 - D. Local Computer\Enterprise Trust

4. Which of the following are prerequisites for running the Active Directory Federation Services Configuration Wizard in Windows Server 2012 R2? (Choose all that apply.)
- A. Pre-installing a Web server certificate on the AD FS server.
 - B. Pre-creating a group Managed Service Account (gMSA) to be used by AD FS.
 - C. Creating a CNAME record with the name `enterpriseregistration` in local DNS that points to the AD FS server.
 - D. Running the `Add-KdsRootKey` cmdlet in the domain at least 10 hours earlier.
5. Your AD FS server has three configured relying party trusts, named RPT1, RPT2, and RPT3. Currently, global authentication policies on the server are configured to require multi-factor authentication (MFA) for members of the Sales group only. No custom authentication policies are configured for any of the relying party trusts.
- You want to configure authentication policies so that for RPT1, only members of the Accounting group are required to provide MFA. For RPT2, you want only members of the Sales group to be required to provide MFA. For RPT3, you want to require members of both the Sales and Accounting groups to provide MFA.
- You want to achieve these desired results without configuring any superfluous settings. What should you do? (Choose all that apply.)
- A. Add the Accounting group as a condition for MFA in the global authentication policy settings.
 - B. Remove the Sales group as a condition for MFA in the global authentication policy settings.
 - C. Add the Accounting group as a condition for MFA in the per-relying party trust authentication policy for RPT1.
 - D. Add the Sales group as a condition for MFA in the per-relying party trust authentication policy for RPT2.
 - E. Add both the Sales and Accounting groups as a condition for multi-factor authentication in the per-relying-trust authentication policy for RPT3.
6. Which of the following statements is *not* true regarding configuration of the Workplace Join feature?
- A. To configure Workplace Join, you must enable device registration in Windows PowerShell.
 - B. To configure Workplace Join, you must enable device registration in global primary authentication policy.
 - C. The AD FS server certificate must be configured with a DNS-type alternative name of `enterpriseregistration` followed by the DNS suffix of the Active Directory domain.
 - D. On the AD FS server, you must install an SSL certificate that is issued by the enterprise CA corresponding to the AD FS server's domain.



Thought experiment

Configuring identity and access solutions at Cpandl.com

You are an administrator for Cpandl.com. The Cpandl.com company network includes 60 servers running Windows Server 2012 R2.

Cpandl.com employs approximately 300 field workers who connect to an in-house web application while visiting various sites. Currently, the web application is supported by two servers that are hosted on the company premises and are running Windows Server 2012 R2.

Your IT manager has expressed an interest in moving the in-house web application to a cloud service to reduce the administrative overhead associated with hosting the application on the company premises. Your manager wants your Active Directory users to be able to sign in to the cloud-hosted application by using their existing company credentials, but he doesn't want to store a copy of the Active Directory database on the premises of the cloud hosting company. To explore this scenario further on behalf of your organization, your manager has asked you to research what is needed to implement AD FS on your company network.

With this information in mind, answer the following questions. You can find the answers to these questions in the "Answers" section.

- 1.** You want employees to use a standard Internet connection (not a VPN) to authenticate and connect to the application wherever they are. However, you don't want to expose the AD FS server directly to the Internet. How can you meet both of these goals?
- 2.** The in-house web application has been modified so that it can authorize through claims and industry-standard tokens. Which wizard should you run to configure the AD FS server to put claims in a format that can be read by this web application?
- 3.** How can you configure authentication policies so that all users connecting from the Internet are able to authenticate through a smart card?

Answers

This section contains the answers to the Objective Review and the Thought Experiment.

Objective 14.1: Review

1. Correct answer: A

- A. Correct:** The Endpoints node includes paths to services and metadata, including the federation metadata file.
- B. Incorrect:** This node shows claims provider trusts, not federation metadata. Claims provider trusts are trust objects typically representing a resource partner that is providing claims information.
- C. Incorrect:** This node shows relying party trusts, not federation metadata. Relying party trusts are trust objects typically created in partner organizations that consume claims information from the local instance of AD FS.
- D. Incorrect:** This node shows attribute stores, not federation metadata. Attribute stores are directories or databases that an organization uses to store its user accounts and their associated attribute values.

2. Correct answer: D

- A. Incorrect:** The Web Server template is needed. The Computer template allows a computer to authenticate itself on the network
- B. Incorrect:** The Web Server template is needed. The Domain Controller Authentication template is used to authenticate Active Directory computers and users.
- C. Incorrect:** The Web Server template is needed. The Kerberos Authentication template offers enhanced security capabilities for domain controllers authenticating Active Directory users and computers.
- D. Correct:** AD FS communicates with other parties through a local web server. The Web Server template is used to prove the identity of this web server.

3. Correct answer: C

- A. Incorrect:** You want to import the certificate into a store for the local computer because the certificate will be used by the computer in general, not by your user account.
- B. Incorrect:** You want to import the certificate into a store for the local computer because the certificate will be used by the computer in general, not by your user account.
- C. Correct:** The Local Computer\Personal certificate store is used to store certificates issued to the local computer and associated with a private key. A private key is needed to sign objects such as tokens.
- D. Incorrect:** The Local Computer\Enterprise Trusts certificate store is used to store certificate trust lists. A certificate trust list provides a way to trust self-signed root certificates from other organizations in a limited way.

4. Correct answers: A, D

- A. Correct:** To be able to complete the AD FS Configuration Wizard, you first need to install an SSL certificate in the Personal store of the local computer account on the AD FS server.
- B. Incorrect:** You don't need to pre-create a gMSA before running the AD FS Configuration Wizard. You can create the gMSA as part of the wizard.
- C. Incorrect:** This step is required to configure Workplace Join, but it is not required to run the AD FS Configuration Wizard.
- D. Correct:** In a production environment, you need to create the KDS root key for the domain at least 10 hours before you can create any gMSA.

5. Correct answers: B, C, D, E

- A. Incorrect:** This step would make membership in the Accounting group a condition for MFA for RPT2, which is not desired.
- B. Correct:** You don't want membership in Sales to be a condition for MFA for RPT1. To prevent this possibility from happening, you must remove Sales as a condition for MFA from the global authentication settings.
- C. Correct:** You need to take this step if you want membership in the Accounting group to serve as a condition for MFA for RPT1 but not some other relying parties.
- D. Correct:** You need to take this step if you want membership in the Sales group to serve as a condition for MFA for RPT2 but not some other relying parties.
- E. Correct:** You need to take this step if you want membership in either the Sales or Accounting groups to serve as a condition for MFA for RPT3 but not some other relying parties.

6. Correct answer: D

- A. Incorrect:** You must enable the Device Registration Service by using the Initialize-ADDeviceRegistration and Enable-DeviceRegistration cmdlets.
- B. Incorrect:** You must select the Enable Device Registration option in the global primary authentication policy.
- C. Incorrect:** The service communications certificate on the AD FS server must be configured with this DNS-type alternative name.
- D. Correct:** The certificate does not need to be issued by an enterprise CA. It can also be issued by a public CA.

Thought experiment

1. Implement a Web Application Proxy to receive authentication requests from users on the Internet. The Web Application Proxy can contact the AD FS server on behalf of the user.
2. The Add Relying Party Trust Wizard.
3. In Global Primary Authentication Policy, beneath Extranet, select the Certificate Authentication check box.

This page intentionally left blank

Index

Numbers & Symbols

6to4 transition technology, 145
802.1x enforcement, 179

A

acceptance transform rules, 398
access control, role-based access control for IPAM (IP Address Management) in Windows Server 2012 R2, 381–382
access control lists (ACLs), 287
access-denied assistance, file classification, 300–301
Access-Denied Assistance tab (FSRM Options dialog box), 300
access policies, configuring, 302–307
 central access rules, 302–306
 deploying central access policies to file servers, 306–307
access rules, 286
ACLs (access control lists), 287
Active Directory
 cloning domain controllers, 193–200
 adding source controllers to Cloneable Domain Controllers group, 194–195
 exporting VMs, 199–200
 importing VMs, 199–200
 Get-ADDCCloningExcludedApplicationList cmdlet, 195–198
 New-ADDCCloneConfigFile cmdlet, 198–199
 prerequisites, 194
 installing domain controllers, 107–127
 restoring deleted objects, 202–209
 AD Administrative Center, 205–207
 deleted object lifetimes, 208
 enabling AD Recycle Bin, 204–205
 Windows PowerShell, 207–208
 updating files, 293–294
 updating folders, 293–294
Active Directory Administrative Center, 205–207
Active Directory Certificate Services (AD CS), 389
Active Directory-detached clusters, 246–247
Active Directory Domain Services Configuration Wizard, installing domain controllers, 107–127
 GUI, 108–112
 IFM option, 119–120
 Windows Azure, 120–127
 Windows PowerShell, 113–118
Active Directory Federation Services. *See* AD FS (Active Directory Federation Services)
Active Directory Recycle Bin, 203–208
Active Directory Rights Management Services (AD RMS), 389
AD CS (Active Directory Certificate Services), 389
Add-ADDSSReadOnlyDomainControllerAccount cmdlet, 113, 117
Add-ADGroupMember cmdlet, 195
Add Authorization Entry dialog box, 332–333
Add-DAAppServer cmdlet, 166
Add-DACLient cmdlet, 157
Add-KdsRootKey cmdlet, 394
/Add-Package option (DISM utility), 20
Add-PswaAuthorizationRule cmdlet, 17
Add Relying Party Trust Wizard, 398
Add-RemoteAccessLoadBalancerNode cmdlet, 152
Add Roles and Features Wizard, 9–12, 361
ADDSSdeployment module, 113
Add Servers option (Server Manager), 47
Add-VMFibreChannelHba cmdlet, 88
Add-VMNetworkAdapterAcl cmdlet, 93–94
AD FS (Active Directory Federation Services), 389–407
 authenticating users, 391–392
 Configuration Wizard, 392–394

- Implementing, 391
- management console, 395–403
 - Authentication Policies node, 399–403
 - Service node, 395–397
 - Trust Relationships node, 397–399
- scenarios, 390–391
- Windows PowerShell cmdlets, 406–407
- Workplace Join, 403–406
- AD FS Proxy, 390
- Adprep /domainprep, 112
- Adprep /forestprep, 112
- AD RMS (Active Directory Rights Management Services), 389
- advanced configuration options
 - (DirectAccess), 167–168
- Advanced Security Settings for Permissions dialog box, 304
- affinity groups, 121–122
- /All argument (DISM utility), 19
- Allow Replication from Any Authenticated Server security option, 332
- Allow Replication from the Specified Servers security option, 332
- All Servers section (Server Manager), 49–51
- All (Write) and Local (Read) mode, 248
- application-consistent recovery points, 338
- application monitoring, virtual machines, 254–260
- Application Server Setup page, configuring
 - DirectAccess, 166–167
- assigning
 - role startup priority, 253
 - SHV configurations to health policies, 186–188
- asynchronous GPO processing, 224–225
- authenticating users, AD FS, 390–392
- Authentication and Ports, Hyper-V Replica, 331
- Authentication page (Remote Access Server Setup Wizard), 161
- Authentication Policies node (AD FS management console), 399–403
 - global policies, 399–402
 - Per Relying Party Trust node, 402–403
- Authorization and Storage, Hyper-V Replica, 332–334
- Automatic allocation, storage pools, 34
- automatic file classification, 295–300
- automatic node drain on shutdown, 242

B

- Backup-GPO cmdlet, 223
- Back Up Now option (Windows Azure Backup), 321–322
- Backup Operators group, 324
- backups, 313–326
 - Back Up Now option, 321–322
 - bandwidth throttling, 322–324
 - certificate requirements, 314
 - creating online backup schedules, 318–321
 - creating backup vaults in management portals, 315–316
 - creating self-signed certificates with Makecert.exe utility, 314–315
 - downloading Windows Azure Backup Agent, 316–317
 - installing Windows Azure Backup Agent, 316–317
 - performing backups in Windows PowerShell, 324–326
 - recovering data, 322
 - registering servers, 317
 - uploading certificates, 316
- bandwidth management, virtual networks, 97–98
- bandwidth throttling, 322–324
- Basic Input Output System (BIOS), 70
- Before You Begin page (Add Roles And Features Wizard), 11
- Behind an Edge Device option (Network Topology page), 160
- Best Practices Analyzer section (Server Manager All Servers page), 50
- bidirectional access (DirectAccess), 144
- BIOS (Basic Input Output System), 70
- blocks (IP addresses), adding to IPM databases, 373–374
- bring-your-own-device (BYOD) trend, Workplace Join, 403
- business continuity
 - backups, configuring, 313–326
 - Back Up Now option, 321–322
 - bandwidth throttling, 322–324
 - certificate requirements, 314
 - creating online backup schedules, 318–321
 - creating backup vaults in management portals, 315–316

- creating self-signed certificates with
 - Makecert.exe utility, 314–315
 - downloading Windows Azure Backup Agent, 316–317
 - installing Windows Azure Backup Agent, 316–317
 - performing backups in Windows PowerShell, 324–326
 - recovering data, 322
 - registering servers, 317
 - uploading certificates, 316
 - site-level fault tolerance, configuring, 329–351
 - extending replication to a third site, 347–348
 - Hyper-V physical host servers, 330–333
 - Hyper-V Replica failover, 342–346
 - Hyper-V Replica in failover clusters, 348–352
 - virtual machines, 333–342
 - BYOD (bring-your-own-device) trend, Workplace Join, 403
- ## C
- caching, Group Policy, 224–225
 - Cancel Failover (unplanned failovers), 344
 - capturing virtual switch extensions, 91
 - case-sensitive string expressions, 298
 - CAU (cluster-aware updating), 242–246
 - Central Access Policies Configuration dialog box, 306
 - central access rules, creating, 302–306
 - .cer files, 315
 - Certificate Authentication, 400
 - Certificate-Based Authentication (HTTPS), 331
 - certificate requirements, Windows Azure Backup, 314
 - certificates, SSL (Secure sockets Layer), 392–394
 - Certificate Services, 394
 - Certificates node (AD FS management console), 396–397
 - checkpoints, 77
 - Choose Initial Replication Method page (Enable Replication Wizard), 339
 - Choose Replication VHDs page (Enable Replication Wizard), 335
 - Claim Rule Language, 399
 - claims
 - authentication, configuring, 287–291
 - defined, 286–287
 - defining user types, 288–290
 - defining device types, 288–290
 - rules for relying party trusts, 398–399
 - Claims Description node (AD FS management console), 397
 - Claims Provider Trusts node (AD FS management console), 398
 - Classification Parameters dialog box, 298–299
 - classifications, defined, 286
 - Classification tab (files and folders), 294
 - Clear-EventLog cmdlet, 136
 - Client Experience settings (DirectAccess), 170
 - Client Setup Wizard, configuring DirectAccess, 156–159
 - Cloneable Domain Controllers group, adding source domain controllers to, 194–195
 - cloning domain controllers, 193–200
 - adding source controller to Cloneable Domain Controllers group, 194–195
 - exporting/importing the VM of the source domain controller, 199–200
 - Get-ADDCCloningExcludedApplicationList cmdlet, 195–198
 - New-ADDCCloneConfigFile cmdlet, 198–199
 - prerequisites, 194
 - cluster-aware updating (CAU), 242–246
 - Cluster-Aware Updating dialog box, 244–245
 - cluster properties, configuring, 248–249
 - Cluster Shared Volume File System (CSVFS), 236
 - cluster shared volumes (CSVs), 236–239
 - cmdlets
 - Add-ADDSReadOnlyDomainControllerAccount, 113, 117
 - Add-ADGroupMember, 195
 - Add-DAAppServer, 166
 - Add-DAClient, 157
 - Add-KdsRootKey, 394
 - Add-PswaAuthorizationRule, 17
 - Add-RemoteAccessLoadBalancerNode, 152
 - Add-VMFibreChannelHba, 88
 - Add-VMNetworkAdapterAcl, 93, 94
 - AD FS, 406–407
 - Clear-EventLog, 136
 - Configure-SMRemoting.exe, 55
 - Convert-VHD, 83
 - Disable-VMEventing, 136
 - Disable-VMResourceMetering, 78
 - Disable-VMSwitchExtension, 93

- ComputerName option (Add-PswaAuthorizationRule cmdlet), 17
- Configuration Friendly Name dialog box, 184
- Configuration function (Windows PowerShell ISE tool), 13
- ConfigurationName option (Add-PswaAuthorizationRule cmdlet), 17
- Configuration Wizard (AD FS), 392–394
- Configure Additional Recovery Points page (Enable Replication Wizard), 336
- Configure and Manage High Availability domain. *See* high availability
- Configure Custom Fields dialog box, 376
- Configure File and Storage Solutions domain. *See* Dynamic Access Control
- Configure NAP Wizard, 181
- Configure Network Services and Access domain. *See* DirectAccess, configuring
- Configure Replication Frequency page (Enable Replication Wizard), 335
- Configure Self-Updating Options Wizard, 245
- Configure-SMRemoting.exe cmdlet, 55
- configuring
 - access policies, 302–307
 - central access rules, 302–306
 - deploying central access policies to file servers, 306–307
 - backups. *See* Windows Azure Backup
 - claims-based authentication, 287–291
 - defining user and device claim types, 288–290
 - enabling Kerberos support, 290–291
 - cluster properties, 248–249
 - constrained delegation, 263
 - custom fields, 376
 - Data Collector Sets, 138–139
 - DirectAccess, 154–168
 - advanced configuration options, 167–168
 - Application Server Setup Wizard, 166–167
 - Client Setup Wizard, 156–159
 - Infrastructure Server Setup Wizard, 162–166
 - Remote Access Server Setup Wizard, 159–162
 - verifying configuration, 168–170
 - domain controllers, 193–200
 - adding source controllers to Cloneable Domain Controllers group, 194–195
 - exporting VMs of the source domain controllers, 199–200
 - Get-ADDCCloningExcludedApplicationList cmdlet, 195–198
 - importing VMs of the source domain controllers, 199–200
 - New-ADDCCloneConfigFile cmdlet, 198–199
 - prerequisites, 194
 - file classification, 291–301
 - access-denied assistance, 300–301
 - adding resource properties to resource property list, 293
 - automatic classification, 295–300
 - creating selected resource properties, 292
 - enabling selected resource properties, 292
 - manual classification, 294–295
 - updating Active Directory files, 293–294
 - Group Policy, 215–225
 - Remote Group Policy update, 216–222
 - Windows PowerShell cmdlets, 222–224
 - high availability
 - failover clustering, 234–248
 - failover clustering roles, 251–259
 - virtual machine movement, 260–277
 - Hyper-V
 - virtual machine settings, 67–80
 - virtual machine storage, 82–89
 - virtual networks, 90–100
 - Hyper-V Settings, 261
 - IPAM (IP Address Management), 361–371
 - adding servers, 367–371
 - connecting to servers, 363
 - manual configuration steps, 371
 - provisioning servers, 363–365
 - selecting servers, 367–371
 - Server Discovery, 366
 - Server Manager, 361–380
 - IP utilization thresholds, 378–379
 - local storage, 30–37
 - creating storage pools, 31–33
 - creating virtual disks, 34–37
 - installing Storage Spaces, 31
 - Minimal Server Interface, 24
 - NAP (Network Access Protection), 177–188
 - SHV multi-configuration, 183–188
 - PSWA, 17
 - servers, 7–28
 - Group Policy Based provisioning method, 368–370
 - installing roles and features, 7–20

Configuring Server Roles and Features domain

- Minimal Server Interface, 22–24
 - NIC teaming, 24–28
 - remote management, 45–60
 - Server Core-GUI convertibility, 20–22
 - site-level fault tolerance, 329–351
 - extending replication to third sites, 347–348
 - Hyper-V physical host servers, 330–333
 - Hyper-V Replica failover, 342–346
 - Hyper-V Replica in failover clusters, 348–352
 - virtual machines, 333–342
 - Configuring Server Roles and Features domain, 45–60
 - Group Policy, 57–59
 - managing multiple servers with Server Manager, 46–57
 - Remote Server Administration Tools, 59–60
 - Confirm Installation Selections page (Add Roles and Features Wizard), 11
 - connection processes (DirectAccess), 146–147
 - connection request policies (NAP processing), 179
 - constrained delegation, configuring, 263
 - Content Classifier classification method, 297
 - Convert-VHD cmdlet, 83
 - Copy-GPO cmdlet, 223
 - Create Additional Hourly Recovery Points option, 337
 - Create Central Access Rule page, 302–303
 - Create Claim Type page, 289
 - Create Classification Rule dialog box, 295–296
 - Create Full NoDefrag %s (IFM menu), 119
 - Create Sysvol Full NoDefrag %s (IFM menu), 119
 - creating
 - central access rules, 302–306
 - CSVs (cluster shared volumes), 236–238
 - custom fields for IPM, 374–376
 - Data Collector Sets manually, 137–138
 - IP address range groups, 377–378
 - KDS root keys, 394
 - self-signed certificates, 314–315
 - SHV configurations, 184–186
 - SoFS (Scale-Out File Server), 251–252
 - storage pools, 31–33
 - virtual disks, 34–37
 - virtual machine resource pools, 135
 - virtual machine settings, 67–80
 - Dynamic Memory, 74–77
 - enhanced session mode, 72–74
 - generation 1 and 2 VMs, 69–71
 - Hyper-V Module in Windows PowerShell, 68–69
 - NUMA (non-uniform memory access), 79
 - RemoteFX, 79–80
 - Resource Metering, 77–78
 - virtual networks, 90–100
 - advanced features for virtual network adapters, 99–100
 - bandwidth management, 97–98
 - network isolation, 93–95
 - SR-IOV, 95–98
 - virtual switch extensions, 91–93
 - CredSSP (Credential Security Support Provider), 262
 - CSVFS (Cluster Shared Volume File System), 236
 - CSVs (cluster shared volumes), 236–239
 - live migration in failover clusters, 264–267
 - CustomDCCloneAllowList.xml file, 196–197
 - custom fields
 - applying to addresses and ranges, 376
 - configuring, 376
 - creating for IPM, 374–376
- ## D
- Data Collector Sets, 137–138
 - DCCloneConfig.xml files, 198
 - Dcgpofix command, 224
 - DC/NPS
 - manual configuration steps for managed infrastructure servers in IPAM, 371
 - DCOM (Distributed Component Object Model), 52–53
 - declarative management, DSC, 14
 - Default Web Site (IIS), 17
 - Delegation Authorization Rules, 399
 - deleted object lifetime, 205, 208
 - deleted objects (Active Directory), restoring, 202–209
 - AD Administrative Center, 205–207
 - deleted object lifetime, 208
 - enabling AD Recycle Bin, 204–205
 - Windows PowerShell, 207–208
 - Deleted Objects container (Active Directory Administrative Center), 205–207
 - deployment, servers, 133–139
 - reviewing older features, 137–139
 - virtual machine resource pools, 134–135
 - Windows PowerShell, 136–137
 - Deployment Image Servicing and Management (DISM) utility, 17–20
 - Deployment Scenario page (DirectAccess Client Setup Wizard), 156

- Description property, 220
- Desired State Configuration (DSC), 13–17
- Device Registration Service (DRS), 404
- DFSR (Distributed File System Replication), 209
- DHCP
 - manual configuration steps for managed infrastructure servers in IPAM, 371
 - server logs, searching for leases by client names/addresses, 373–374
- DHCP Guard, 99
- diagnosing performance problems, 138
- dialog boxes
 - Add Authorization Entry, 332–333
 - Advanced Security Settings for Permissions, 304
 - Central Access Policies Configuration, 306
 - Classification Parameters, 298–299
 - Cluster-Aware Updating, 244–245
 - Configuration Friendly Name, 184
 - Configure Custom Fields, 376
 - Create Classification Rule, 297
 - Edit Global Authentication Policy, 399
 - File Server Resource Manager Options, 299
 - Hyper-V Settings, 262–263, 330
 - IPAM Settings, 375
 - Move Virtual Machine Storage, 264–265
 - New Team, 26
 - Permission Entry For Permissions, 304
 - Properties, 258
 - Select Items, 319
 - Select Resource Properties, 293
 - Select Services, 255
 - Settings (VMs), 341
- DirectAccess
 - benefits, 144
 - Client Experience Settings, 170
 - configuring, 154–168
 - advanced configuration options, 167–168
 - Application Server Setup Wizard, 166–167
 - Client Setup Wizard, 156–159
 - Infrastructure Server Setup Wizard, 162–166
 - Remote Access Server Setup Wizard, 159–162
 - verifying configuration, 168–170
 - connection process, 146–147
 - infrastructure options, 147–152
 - complex, 151–152
 - multidomains, 149–151
 - multisites, 149–151
 - server behind NAT, 148–149
 - simple, 147–148
 - installing, 153–154
 - IPv6 communication, 144–146
- DirectAccessClientComponents module, 153
- Directory Services Restore Mode (DSRM) password, 114
- Directory tab, DCS, 138
- /Disable-Feature switch (DISM utility), 19
- Disable-VMEventing cmdlet, 136
- Disable-VMResourceMetering cmdlet, 78
- Disable-VMSwitchExtension cmdlet, 93
- disabling Group Policy caching, 225
- disaster recovery
 - backups, configuring, 313–326
 - Back Up Now option, 321–322
 - bandwidth throttling, 322–324
 - certificate requirements, 314
 - creating online backup schedules, 318–321
 - creating backup vaults in management portals, 315–316
 - creating self-signed certificates with Makecert.exe utility, 314–315
 - downloading Windows Azure Backup Agent, 316–317
 - installing Windows Azure Backup Agent, 316–317
 - performing backups in Windows PowerShell, 324–326
 - recovering data, 322
 - registering servers, 317
 - uploading certificates, 316
 - site-level fault tolerance, configuring, 329–351
 - extending replication to a third site, 347–348
 - Hyper-V physical host servers, 330–333
 - Hyper-V Replica failover, 342–346
 - Hyper-V Replica in failover clusters, 348–352
 - virtual machines, 333–342
- Disk Quotas Group Policy setting, 225
- disk space requirements, 2
- DISM (Deployment Image Servicing and Management) utility, 17–20
- Distributed Component Object Model (DCOM), 52–53
- Distributed File System Replication (DFSR), 209
- Djoin.exe tool, 119
- DNS
 - manual configuration steps for managed infrastructure servers in IPAM, 371
 - servers, registering, 123
- DNS page (Infrastructure Server Setup Wizard), 163

DNS Suffix Search List page (Infrastructure Server Setup Wizard), 164–165

domain controllers

configuring/cloning, 193–200

adding source controller to Cloneable Domain Controllers group, 194–195

exporting source domain controller VMs, 199–200

Get-ADDCCloningExcludedApplicationList cmdlet, 195–198

importing source domain controller VMs, 199–200

New-ADDCCloneConfigFile cmdlet, 198–199

prerequisites, 194

installing, 107–127

GUI, 108–112

IFM option, 119–120

Windows Azure, 120–127

Windows PowerShell, 113–118

remote management options, 49

Domain Naming Master, 108

downloading Windows Azure Backup Agent, 316–317

Drive Mapping, 225

DRS (Device Registration Service), 404

Dsamain tool, 208

DSC (Desired State Configuration), 13–17

DSRM (Directory Services Restore Mode) password, 114

Dynamic Access Control, 285–307

access policies, 302–307

central access rules, 302–306

deploying central access policies to file servers, 306–307

claims-based authentication, 287–291

defining user claim types, 288–290

defining device claim types, 288–290

enabling Kerberos support, 290–291

file classifications, 286–287, 291–301

access-denied assistance, 300–301

adding resource properties to resource property list, 293

automatic classification, 295–300

creating selected resource properties, 292

enabling selected resource properties, 292

manual classifications, 294–295

updating Active Directory files/folders, 293–294

Dynamic Memory, 74–77

DynamicQuorum cluster property, 248–249

dynamic quorum configuration, 240–241

dynamic witness, configuring failover clustering, 241

E

Edge option (Network Topology page), 160

Edit Global Authentication Policy dialog box, 399

Edit Virtual Hard Disk Wizard, 84–85

EKU (Enhanced Key Usage), 331

Enable-ADOptionalFeature cmdlet, 204

Enable Corporate Compliance for DirectAccess Clients with NAP setting (Authentication page), 162

Enable Internet Bandwidth Usage Throttling For Backup Operations check box, 322

Enable-NetFirewallRule cmdlet, 52

Enable-NetFirewallRule-DisplayName command, 331

Enable Replication Wizard, 334–340

Enable This Computer as a Replica Server check box, 331

Enable-VMEventing cmdlet, 136

Enable-VMResourceMetering cmdlet, 77, 134

Enable-VMSwitchExtension cmdlet, 93

Enable Windows 7 Client Computers to Connect Via DirectAccess setting (Authentication page), 162

enabling

access-denied assistance, 300

Active Directory Recycle Bin, 204–205

bandwidth throttling, 322–324

Kerberos support for claims-based authentication, 290–291

replication on clustered VMs, 351

Endpoints node (AD FS management console), 396

enforcement points (NAP), 178

enforcement types (NAP), 179

Enhanced Key Usage (EKU), 331

Enhanced session mode

policy, 72

virtual machines, 72–74

Evaluation Type tab (Create Classification Rule dialog box), 298

EVENT CATALOG page, searching DHCP server logs for leases by client names/addresses, 373–374

Events section (All Servers page, Server Manager), 50

Event Trace Data option, manually creating DCS, 137

excluding items from backups, 320

Export Configuration Settings (Add Roles and Features Wizard), 12

Export-Counter cmdlet, 136

exporting source domain controller VMs, 199–200

Export-VM cmdlet, 199

Extended Protection for Authentication feature, 403

extending replication to third sites, 347

F

failover clustering

configuring, 234–248

Active Directory-detached clusters, 246–247

CAU (cluster-aware updating), 242–246

CSVs (cluster shared volumes), 236–239

dynamic quorum configuration, 240–241

dynamic witness, 241

node drain, 241

roles, 251–259

storage pools, 234–235

virtual hard disk sharing, 239–240

defined, 233

Hyper-V Replica, 348–351

Failover Cluster Manager, 234, 348

failovers

Hyper-V Replica, 342–346

TCP/IP settings, 340–341

feature files

reinstalling, 4–5

removing, 3

Features on Demand, 2–5

Federation Metadata file, 396

Fibre Channel adapters, 86–88

fields, custom fields

applying to addresses and ranges, 376

configuring, 376

creating for IPM, 374–376

file attributes, 286

file classifications, configuring, 291–301

access-denied assistance, 300–301

adding resource properties to resource property list, 293

automatic classification, 295–300

creating selected resource properties, 292

enabling selected resource properties, 292

manual classifications, 294–295

updating Active Directory files/folders, 293–294

File DSC resource, 15–16

File Replication Service (FRS), 209

files

.cer, 315

CustomDCCloneAllowList.xml, 196–197

DCCloneConfig.xml, 198

.pfx, 315

VHD (virtual hard disk), 236

File Server Resource Manager Options dialog box, 299

File Server role service, 31

file solutions. *See* Dynamic Access Control

filtering virtual switch extensions, 91

firewall rules

Remote Group Policy update, 221–222

remote management, 59

First Failure response, 255

fixed provisioned spaces, 36

Folder Classifier classification method, 297

Folder Redirection Group Policy setting, 225

folders, configuring access-denied assistance, 301

-Force parameter (Invoke-GpUpdate cmdlet), 219

forwarding virtual switch extensions, 91

FRS (File Replication Service), 209

full installation, 20

G

generation 1 virtual machines, 69–71

generation 2 virtual machines, 69–71

Geo-Redundant Replication, 122

Get-ADComputer cmdlet, 219–220

Get-ADDCCloneExcludedApplicationList cmdlet, 195–198

Get-ADForest cmdlet, 204

Get-ADObject cmdlet, 207

Get-ADUser cmdlet, 220

(Get-Cluster).DatabaseReadWriteMode cmdlet, 249

(Get-Cluster).DynamicQuorum cmdlet, 248

(Get-ClusterNode Node2).NodeWeight cmdlet, 249

(Get-Cluster).WitnessDynamicWeight cmdlet, 248

Get-Command, 27, 68

Get-Command-Module AD FS command, 407

Get-Counter cmdlet, 136

Get-DAConnectionStatus cmdlet, 169

Get-Event cmdlet, 136

- Get-EventLog cmdlet, 136
- /Get-Features switch (DISM utility), 19
- Get-GPInheritance cmdlet, 223
- Get-GPO cmdlet, 223
- Get-GPOReport cmdlet, 223
- Get-GPPermission cmdlet, 223
- Get-GPPrefRegistryValue cmdlet, 223
- Get-GPRegistryValue cmdlet, 223
- Get-GPResultantSetOfPolicy cmdlet, 223
- Get-GPStarterGPO cmdlet, 223
- Get-Help cmdlet, 27
- /Get-ImageInfo switch (DISM utility), 18
- Get-NetLbfoTeam cmdlet, 27
- Get-OBPolicy cmdlet, 326
- Get-OBSchedule cmdlet, 326
- Getting Started Wizard, 154–155
- Get-VMFibreChannelHba cmdlet, 88
- Get-VMNetworkAdapterACL cmdlet, 94–95
- Get-VMResourcePool cmdlet, 134
- Get-VMSwitchExtension cmdlet, 93
- Get-WindowsFeature cmdlet, 3, 8
- global policies, AD FS management console, 399–402
- gMSAs (group Managed Service Accounts), 394
- Gpfixup command, 224
- GPT (GUID Partition Table) partition style, 32
- Gpupdate /sync command, 225
- Graphical Management Tools and Infrastructure feature, 21
- Group DSC resource, 15
- group Managed Service Accounts (gMSAs), 394
- Group Policy
 - caching, 224–225
 - configuring, 215–225
 - Remote Group Policy update, 216–222
 - Windows PowerShell cmdlets, 222–224
 - deploying central access policies to file servers, 306–307
 - enabling Kerberos support for claims, 290–291
 - enabling remote management, 57–59
- Group Policy Based provisioning method, configuring servers, 368
- Group Policy Management Console, 218
- Group Policy Remote Update firewall ports, 221
- groups, security, IPAM server, 380
- guest clusters, 239
- GUIs, installing domain controllers, 108–112
- GUID Partition Table (GPT) partition style, 32

H

- hardware requirements, server installation, 2
- health checks, SHV configuration, 185
- health policies (NAP processing), 179, 186–188
- Health Policy condition, network policies, 179
- Health Registration Authority (HRA), 147
- Health Registration Authority (HRA) servers, 179
- high availability
 - failover clustering, 234–248
 - Active Directory-detached clusters, 246–247
 - CAU (cluster-aware updating), 242–246
 - CSVs (cluster shared volumes), 236–239
 - dynamic quorum configuration, 240–241
 - dynamic witness, 241
 - node drain, 241
 - roles, 251–259
 - storage pools, 234–235
 - virtual hard disk sharing, 239–240
 - virtual machine movement, 260–277
 - live migration, 261–273
 - storage migration, 274–276
- High Availability Wizard, 251, 348
- host clusters, 239
- Hot Spare allocation, storage pools, 34
- HRA (Health Registration Authority), 147, 179
- Hyper-V
 - configuring, 261
 - virtual machine settings, 67–80
 - virtual machine storage, 82–89
 - virtual networks, 90–100
 - host clusters, 239
- Hyper-V Module, Windows PowerShell, 68–69
- Hyper-V Replica, 329–351
 - configuring physical host servers, 330–333
 - configuring virtual machines, 333–342
 - failover TCP/IP settings, 340–341
 - resynchronizing primary and replica VMs, 341–342
 - extending replication to third sites, 347–348
 - in failover clusters, 348–352
 - failovers, 342–346
- Hyper-V Replica Broker role, 348
- Hyper-V Replica HTTP Listener, 331
- Hyper-V Settings dialog box, 262–263, 330

- I
- laaS (infrastructure-as-a-service), 120
- identity solutions. *See* AD FS (Active Directory Federation Services)
- IFM (Install from Media) option, 119–120
- /IgnoreCheck option (DISM utility), 20
- Import-Counter cmdlet, 136
- Import-GPO cmdlet, 223
- Importing source domain controller VMs, 199–200
- Import-VM cmdlet, 199
- infrastructure
 - DirectAccess, 147–152
 - complex, 151–152
 - multidomains, 149–151
 - multisites, 149–151
 - server behind NAT, 148–149
 - simple, 147–148
 - NAP, 178–180
- infrastructure-as-a-service (laaS), 120
- Infrastructure Server Setup Wizard, configuring
 - DirectAccess, 162–166
- Initialize-ADDeviceRegistration cmdlet, 404
- Install-ADDSDomain cmdlet, 113, 116–117
- Install-ADDSDomainController cmdlet, 113, 116
- Install-ADDSEForest cmdlet, 113, 115
- Installation Progress page (Add Roles and Features Wizard), 12
- Install from Media (IFM) option, 119–120
- installing
 - Active Directory domain controllers, 107–127
 - DirectAccess, 153–154
 - domain controllers, 107–127
 - GUI, 108–112
 - IFM option, 119–120
 - Windows Azure, 120–127
 - Windows PowerShell, 113–118
 - IPAM (IP Address Management), 361–362
 - Makecert utility, 314
 - roles, 7–20
 - DISM (Deployment Image Servicing and Management), 17–20
 - DSC (Desired State Configuration), 13–17
 - PSWA (Windows PowerShell Web Access), 16–17
 - Server Manager, 9–12
 - Windows PowerShell, 8–9
 - servers, 1–5
 - Features on Demand, 2–5
 - hardware requirements, 2
 - SSL (Secure Sockets Layer) certificates, 392–394
 - Storage Spaces, 31
 - Windows Azure Backup Agent, 316–317
- Install-PswaWebApplication cmdlet, 17
- Install-WindowsFeature cmdlet, 3, 9
- Integration Services, 79
- Invoke-GPUUpdate cmdlet, 218–220, 223
- IP addresses
 - adding to IPAM databases, 372–374
 - IPAM (IP Address Management), 359–360
 - administrative solutions, 360
 - configuring, 361–380
 - installing, 361–380
 - limitations, 360
 - managing space, 372–380
 - range groups, creating, 377–378
 - IP Address Management. *See* IPAM (IP Address Management)
 - IPAM Administrators group, 380
 - IPAM ASM Administrators group, 380
 - IPAM (IP Address Management), 360
 - administrative solutions, 360
 - configuring, 361–371
 - adding servers, 367–371
 - connecting to IPAM servers, 363
 - manual configuration steps, 371
 - provisioning IPAM servers, 363–365
 - selecting servers, 367–371
 - Server Discovery, 366
 - Server Manager, 361–380
 - starting Server Discovery, 366
 - installing, 361–362
 - limitations, 360
 - managing IP address space, 372–380
 - adding IP addresses to IPAM databases, 372–374
 - applying custom fields to addresses and ranges, 376
 - creating custom fields, 374–376
 - creating IP address range groups, 377–378
 - delegating administration, 380
 - viewing/configuring IP utilization thresholds, 378–379
 - role-based access control in Windows Server 2012 R2, 381–382
- IPAM IP Audit Administrators group, 380
- IPAM MSM Administrators group, 380
- IPAM Overview page preconfiguration, 362–381

IPAM servers

- IPAM servers
 - connecting to, 363
 - local servers, manual configuration steps for managed infrastructure servers in IPAM, 371
 - provisioning, 363
- IPAM Settings dialog box, 375
- IPAM Users group, 380
- IP-HTTPS transition technology, 145
- IPsec enforcement, 179
- IP utilization thresholds
 - configuring, 378–379
 - viewing, 378–379
- IPv6 communication, DirectAccess, 144–146
- iSCSI Initiator, 307
- iSCSI Target, 307
- iSNS, 307
- Issuance Authorization Rules, 398
- Issuance Transform Rules, 398

K

- KDS root keys, creating, 394
- Kerberos, 262, 290–291
- Knowledge Base (KB) article 2682011 performance update, 56

L

- LastLogonDate property, 220
- LBFO (Load Balancing and Failover), 24–28
- LDP utility, 203
- legacy emulated hardware devices, generation 2 VMs, 69
- limited operating system support, generation2 VMs, 71
- live migration, virtual machines, 261–273
 - failover clusters, 264–267
 - nonclustered environments, 267–270
 - processor compatibility, 270–272
 - virtual switch name matching, 272–273
- load balancing, DirectAccess, 151–152
- Load Balancing and Failover (LBFO), 24–28
- Locally Redundant Replication, 122
- local storage, configuring, 30–37
 - creating storage pools, 31–33
 - creating virtual disks, 34–37

- installing Storage Spaces, 31
- logical unit numbers (LUNs), 31, 236
- Logman.exe cmdlet, 136
- LUNs (logical unit numbers), 31, 236

M

- maintenance, Active Directory, 202–209
 - Administrative Center, 205–207
 - deleted object lifetime, 208
 - enabling Recycle Bin, 204–205
 - Windows PowerShell, 207–208
- Majority (Read And Write) mode, 248
- Makecert.exe utility, 314–315
- Manage Certificate option, Windows Azure, 316
- Managed Object Format (MOF) files, 13
- management
 - backups. *See* Windows Azure Backup Data Collector Sets, 138–139
 - Group Policy, 215–225
 - caching, 224–225
 - Remote Group Policy update, 216–222
 - Windows PowerShell cmdlets, 222–224
- IP address space, 372–380
 - adding IP addresses to IPAM databases, 372–374
 - applying custom fields to addresses and ranges, 376
 - creating custom fields, 374–376
 - creating IP address range groups, 377–378
 - delegating administration, 380
 - viewing/configuring IP utilization thresholds, 378–379
- servers, 133–139
 - review of older features, 137–139
 - virtual machine resource pools, 134–135
 - Windows PowerShell, 136–137
- virtual machine movement, 260–277
 - live migration, 261–273
 - storage migration, 274–276
 - VM network health protection, 276–277
- management console (AD FS), 395–403
 - Authentication Policies node, 399–403
 - global policies, 399–402
 - Per Relying Party Trust node, 402–403
 - Service node, 395–397
 - certificates, 396–397

- claim descriptions, 397
- endpoints, 396
- Trust Relationships node, 397–399
 - claims provider trusts, 398
 - claims rules for relying party trusts, 398–399
- Management page (Infrastructure Server Setup Wizard), 165–166
- manual configurations, IPAM (IP Address Management), 371
- manual file classification, 294–295
- Master Boot Record (MBR) partition style, 32
- Maximum RAM setting, Dynamic Memory, 76
- MBR (Master Boot Record) partition style, 32
- Measure-VM cmdlet, 78
- Measure-VMResourcePool cmdlet, 135
- Memory Buffer setting (Dynamic Memory), 76
- Memory Weight setting (Dynamic Memory), 76
- Metadata section, Endpoints node of AD FS
 - management console, 396
- metering virtual machine resource pools, 134–135
- Microsoft Management Console (MMC) snap-ins, 52
- Migrate a Cluster Wizard, 259
- Minimal Server Interface, 22–24
- Minimum RAM setting (Dynamic Memory), 75
- mirror virtual disks, 35
- MMC (Microsoft Management Console) snap-ins, 52
- MOF (Managed Object Format) files, 13
- monitoring servers, 133–139
 - review of older features, 137–139
 - virtual machine resource pools, 134–135
 - Windows PowerShell, 136–137
- Move Virtual Machine Storage dialog box, 264–265
- Move Wizard, 274
- msDS-DeletedObjectLifetime attribute, 208
- MSONlineBackup module, 324
- MS-Service class condition, network policies, 181
- multidomain DirectAccess infrastructure, 149–151
- multi-factor authentication, 399–402
- Multi-Factor tab (Edit Global Authentication Policy dialog box), 401
- multiple servers, managing with Server Manager, 46–57
 - All Servers section, 49–51
 - DCOM and WinRM, 51–53
 - non-domain-joined servers, 47–49
 - re-enabling Windows Server 2012 for remote management, 54–57
- multisite DirectAccess infrastructure, 149–151

N

- Name property, 220
- Name Resolution Policy, 164
- Name Resolution Policy Table (NRPT), 163
- NAP-Capable condition, network policies, 179
- NAP (Network Access Protection), configuring, 177–188
 - .NET Framework 4 updates, 56
- NetLbfo module, 27
- Network Adapters page (Remote Access Server Setup Wizard), 160–161
- Network Connectivity Assistant page (DirectAccess Client Setup Wizard), 158
- network health protection, virtual machines, 276–277
- network isolation, 93–95
- Network Load Balancing (NLB), 249, 151
- Network Location Server page (Infrastructure Server Setup Wizard), 162–163
- network policies (NAP processing), 179
- Network Policy Server (NPS), 178
- Network Topology page (Remote Access Server Setup Wizard), 159
- New-ADDCCloneConfigFile cmdlet, 198–199
- New-Cluster cmdlet, 247
- New-Event cmdlet, 136
- New-GPLink cmdlet, 223
- New-GPO cmdlet, 223
- New-GPStarterGPO cmdlet, 223
- New Inbound Rule Wizard, 58
- New-ItemProperty-Name cmdlet, 48
- New-NetLbfoTeam cmdlet, 27
- New-OBPolicy cmdlet, 325–326
- New-OBSchedule cmdlet, 325–326
- New-StoragePool cmdlet, 32
- New Storage Pool Wizard, 33, 234
- New Team dialog box, 26
- New-VHD cmdlet, 83
- New-VirtualDisk cmdlet, 34
- New Virtual Disk Wizard, 34–35
- New Virtual Machine Wizard, 69
- New-VMResourcePool cmdlet, 135
- New-VMSwitch cmdlet, 96
- NIC teaming, 24–28, 99
- NLB (Network Load Balancing), 249, 151
- Node and Disk Majority quorum configuration, 240
- Node and File Share Majority quorum configuration, 240

node drain, configuring failover clustering

- node drain, configuring failover clustering, 241
- Node Majority quorum configuration, 240
- NodeWeight cluster property, 249
- nonclustered environments, 267–270
- non-domain-joined servers, adding to Server Manager, 47–49
- non-uniform memory access (NUMA), 79
- NPS (Network Policy Server), 178
- NRPT (Name Resolution Policy Table), 163
- Ntdsutil command-line utility, 203, 206
- Ntdsutil.exe tool, 119
- NUMA (non-uniform memory access), 79

O

- OBPolicy objects, 325
- Offline Domain Join, 119
- one-time passwords (OTPs), 151
- OperatingSystem property, 220
- Operations Status item (Remote Access Management Console), 168
- OTPs (one-time passwords), 151

P

- parity virtual disks, 35
- PEAP (Protected Extensible Authentication Protocol), 179
- Performance Counter Alert option, manually creating DCS, 138
- Performance Counter option, manually creating DCS, 137
- performance, diagnosing problems, 138
- Performance section (All Servers page, Server Manager), 50
- Permission Entry For Permissions dialog box, 304
- Per Relying Party Trust node, AD FS management console, 402–403
- .pfx files, 315
- PhysicalDisk\%Disk Time (performance counter data), 138
- Physical Hard Disk option, 84
- physical host clusters, 239
- physical host servers, configuring Hyper-V, 330–333
- PKI (public key infrastructure), 149

- planned failovers, 342–343
- Port Mirroring, 99
- predefined resource properties, 292
- prerequisites, cloning domain controllers, 194
- /PreventPending option (DISM utility), 20
- Previous Versions tab, 324
- primary servers, Hyper-V Replica, 330
- Primary tab (Edit Global Authentication Policies dialog box), 400
- primary VMs, resynchronizing, 341–342
- primordial pools, 31–32
- principals, 304
- processor compatibility, virtual machine live migration, 270–272
- Processor\%Processor Time (performance counter data), 138
- processor requirements, 2
- Properties dialog box, 258
- properties tabs (DCS), 138
- Protected Extensible Authentication Protocol (PEAP), 179
- Protected Network option (Network Adapter settings), 276
- Provision IPAM Wizard, 364
- PSWA (Windows PowerShell Web Access), 16–17
- public key infrastructure (PKI), 149
- PXE boot-compatible network adapters, generation 2 VMs, 71

Q

- Quick Migration, 267

R

- RADIUS protocol, 178
- RAM requirements, 2
- RandomDelayInMinutes, Invoke-GpUpdate cmdlet, 219
- ranges (IP addresses), adding to IPM databases, 373–374
- Read Access - Geo Redundant Replication, 122
- recovering data, backups, 322
- recovery points, 337
- redundancy, 122–123

- registering servers to enable backups, 317
- Register Server Wizard, 317
- regular expressions, 298
- reinstalling feature files, 4–5
- relying party trusts, AD FS management console, 398
- Remote Access Management console, installing
 - DirectAccess, 153–154
- Remote Access Server Setup Wizard, configuring
 - DirectAccess, 159–162
- Remote Access Setup Wizard, 154
- Remote Desktop Services Installation option, 11
- RemoteFX, 79–80
- Remote Group Policy update, 216–222
 - firewall rules, 221–222
 - Task Scheduler, 220–221
 - updating GP in organizational units, 216–218
 - updating GP with Invoke-GpUpdate cmdlet, 218–220
- remote management, configuring servers for, 45–60
 - Group Policy, 57–59
 - managing multiple servers with Server Manager, 46–57
 - Remote Server Administration Tools, 59–60
- Remote Server Administration Tools, 59–60
- remote servers, deploying roles and features
 - Server Manager, 9–12
 - Windows PowerShell, 8–10
- remote updating mode, 243
- Remote Volume Management, 52
- Remove-GLink cmdlet, 223
- Remove-GPO cmdlet, 223
- Remove-GPPrefRegistryValue cmdlet, 223
- Remove-GPRegistryValue cmdlet, 223
- Remove-VMFibreChannelHba cmdlet, 88
- Remove-VMNetworkAdapterACL cmdlet, 94
- removing feature files, 3
- Rename-GPO cmdlet, 223
- replica servers, Hyper-V Replica, 330
- replication. *See* Hyper-V Replica
- replica VMs, resynchronizing, 341–342
- requirements
 - failover cluster storage pools, 235
 - server installation, 2
 - Storage Spaces, 31
 - VM Monitoring feature, 254
- Requires keyword, configuring features in
 - sequences, 16
- Reset-VMResourceMetering cmdlet, 78
- Resource Control VM setting, 79
- Resource Metering, 77–78
- ResourceMeteringEnabled status, 134
 - ResourcePoolName parameter (Enable-VMResourceMetering cmdlet), 134
- resource pools, virtual machines, 134–135
 - ResourcePoolType parameter (Enable-VMResourceMetering cmdlet), 134
- resource properties (Dynamic Access Control)
 - adding to resource property lists, 293
 - creating, 292
 - defined, 286
 - enabling, 292
 - predefined, 292
- Resource Property Lists container, 293
- resources, DSC, 14–16
- Restart The Service setting, 257
- Restore-ADObject cmdlet, 207
- Restore-GPO cmdlet, 223
- restoring deleted objects (Active Directory), 202–209
 - Administrative Center, 205–207
 - deleted object lifetime, 208
 - enabling Recycle Bin, 204–205
 - Windows PowerShell, 207–208
- resynchronizing primary and replica VMs, 341–342
- retention setting, backups, 321
- Reverse Replication Wizard, 344
- Review Options page (Active Directory Domain Services Configuration Wizard), 110
- RID Master, 108
- RODC accounts, 117–118
- role-based access control, IPAM (IP Address Management) in Windows Server 2012 R2, 381–382
- roles
 - configuring failover clustering, 251–259
 - assigning role startup priority, 253–254
 - SoFS (Scale-Out File Server), 251–252
 - virtual machine application monitoring, 254–260
 - configuring servers for remote management, 45–60
 - Group Policy, 57–59
 - managing multiple servers with Server Manager, 46–57
 - Remote Server Administration Tools, 59–60
 - defined, 233
 - installing, 7–20

role startup priority, assigning

- DISM (Deployment Image Servicing and Management), 17–20
- DSC (Desired State Configuration), 13–17
- PSWA (Windows PowerShell Web Access), 16–17
- Server Manager, 9–12
- Windows PowerShell, 8–9
- role startup priority, assigning, 253
- Router Guard, 99

S

- SafeModeAdministratorPassword parameter (Test-ADDSTestForestInstallation cmdlet), 114
- SAS (Serial Attached SCSI) disk array, 234
- Save Template option, DCS, 138
- scalability, failover clusters, 233
- Scale-Out File Server role, 238
- Scale-Out File Server (SoFS), 251–252
- scenarios, AD FS, 390–391
- Schedule Backup Wizard, 318–321
- Sconfig configuration tool, 55–57
- scope, classification rules, 296
- SCSI boot, generation 2 VMs, 71
- SDK (Software Development Kit), 314
- seamless connectivity (DirectAccess), 144
- searching DHCP server logs for leases by client names/addresses, 373–374
- Second Failure response, 255
- Secure Sockets Layer (SSL) certificates, 314, 392–394
- security groups (IPAM server), 380
- Security Health Validator, 183
- Select Destination Server page (Add Roles and Features Wizard), 11
- Select Groups page (DirectAccess Client Setup Wizard), 157
- selecting items for backup (Schedule Backup Wizard), 318–319
- Select Installation Type page (Add Roles and Features Wizard), 11
- Select Items dialog box, 319
- Select Provisioning Method page, 367–368
- Select Resource Properties dialog box, 293
- Select Services dialog box, 255
- Select The Storage Layout page (New Virtual Disk Wizard), 34
- self-signed certificates, creating with Makecert.exe utility, 314–315
- self-updates, 243
- Serial Attached SCSI (SAS) disk array, 234
- Server Core-GUI convertibility, 20–22
- Server Core Installation, 20
- Server Discovery
 - adding servers, 367–369
 - configuring IPAM (IP Address Management), 366
 - selecting servers, 367–369
 - starting, 366
- Server Graphical Shell feature, 21
- SERVER INVENTORY page, 367
- Server Manager
 - configuring IPAM, 361–380
 - deploying roles and features on remote servers, 9–12
 - managing multiple servers, 46–57
 - All Servers section, 49–51
 - DCOM and WinRM, 51–53
 - non-domain-joined servers, 47–49
 - re-enabling Windows Server 2012 for remote management, 54–57
- servers
 - adding, 367–369
 - configuring, 7–28
 - Group Policy Based provisioning method, 368–370
 - installing roles and features, 7–20
 - Minimal Server Interface, 22–24
 - NIC teaming, 24–28
 - remote management, 45–60
 - Server Core-GUI convertibility, 20–22
 - installing, 1–5
 - Features on Demand, 2–5
 - hardware requirements, 2
 - monitoring, 133–139
 - review of older features, 137–139
 - virtual machine resource pools, 134–135
 - Windows PowerShell, 136–137
 - registering to enable backups, 317
- Server With A GUI installation, 20
- Service Communications certificate, 397
- Service DSC resource, 15
- service-level agreements (SLAs), 89
- Service node (AD FS management console), 395–397
 - certificates, 396–397
 - claim descriptions, 397
 - endpoints, 396
- Services section (All Servers page, Server Manager), 50

- Set-ADForestMode cmdlet, 204
- Set-ADFSProperties cmdlet, 403
- Set-ADOject cmdlet, 208
- Set-DAClient cmdlet, 158
- Set-DAClientDNSConfiguration cmdlet, 164
- Set-DANetworkLocationServer cmdlet, 162
- Set-DAServer cmdlet, 162
- Set-ExecutionPolicy RemoteSigned cmdlet, 57
- Set-GPInheritance cmdlet, 223
- Set-GPLink cmdlet, 223
- Set-GPPermission cmdlet, 224
- Set-GPPrefRegistryValue cmdlet, 224
- Set-GPRegistryValue cmdlet, 224
- Set-Item wsman:\localhost\Client\TrustedHosts cmdlet, 47
- Set-OBMachineSetting cmdlet, 324, 325
- Set-OBPolicy cmdlet, 326
- Set-OBSchedule cmdlet, 326
- Set-RemoteAccessLoadBalancer cmdlet, 152
- Settings dialog box (VMs), 341
- Set-VM cmdlet, 74
- Set-VMFibreChannelHba cmdlet, 88
- Set-VMNetworkAdapter cmdlet, 98
- Shadow Copies settings, file servers, 324
- shutdown, automatic node drain, 242
- SHV (System Health Validator), 179, 183–188
 - assigning configurations to a health policy, 186–188
 - creating additional SHV configurations, 184–186
 - default configuration, 184
- side-by-side store, 2
- simple DirectAccess infrastructure, 147–148
- simple virtual disks, 34
- single-root I/O virtualization (SR-IOV), 95–98
- single-sign-on (SSO) access, 389–390
- site-level fault tolerance, configuring, 329–351
 - extending replication to third sites, 347–348
 - Hyper-V physical host servers, 330–333
 - Hyper-V Replica failover, 342–346, 348–352
 - virtual machines, 333–342
 - failover TCP/IP settings, 340–341
 - resynchronizing primary and replica VMs, 341–342
- Size Of The Virtual Disk page (New Virtual Disk Wizard), 36
- SLAs (service-level agreements), 89
- Smart Paging, 76–77
- snapshots, 77
- SoFS (Scale-Out File Server), 251–252
- Software Development Kit (SDK), 314
- Software Installation Group Policy setting, 225
- SoH (statement of health), 178
- Specify Connection Parameters page (Enable Replication Wizard), 334
- Specify Generation page (New Virtual Machine Wizard), 69
- Specify Replica Server page (Enable Replication Wizard), 334
- Specify Retention Setting page (Schedule Backup Wizard), 321
- Specify The Provisioning Type page (New Virtual Disk Wizard), 36
- SR-IOV (single-root I/O virtualization), 95–98
- SSL (Secure Sockets Layer) certificates, 314, 392–394
- SSO (single-sign-on) access, 389–390
- Start-DscConfiguration cmdlet, 13
- starting Server Discovery, 366
- Start-OBBackup cmdlet, 326
- Start-OBRegistration cmdlet, 324, 326
- startup priorities, assigning to roles, 253–254
- Startup RAM setting (Dynamic Memory), 75
- statement of health (SoH), 178
- Stop Condition tab (DCS), 138
- storage
 - configuring local storage, 30–37
 - creating storage pools, 31–33
 - creating virtual disks, 34–37
 - installing Storage Spaces, 31
 - virtual machines, 82–89
 - Fibre Channel adapters, 86–88
 - Storage QoS, 88–89
 - VHDX disk format, 82–85
- storage accounts, 122–123
- storage migration, virtual machines, 274–276
- storage pools
 - creating, 31–33
 - failover clustering, 234–235
- storage Quality of Service (QoS) virtual machines, 88–89
- Storage Spaces, 30–37
 - creating storage pools, 31–33
 - creating virtual disks, 34–37
 - installing, 31
- string expressions, 298
- Subsequent Failures response, 255
- Suspend-ClusterNode cmdlet, 241
- switches

synchronous GPO processing

- /Disable-Feature, 19
- /Get-Features, 19
- /Get-ImageInfo, 18
- synchronous GPO processing, 224–225
- System Configuration Information option, manually creating DCS, 137
- System Health Validator, 179, 183–188
 - assigning configurations to a health policy, 186–188
 - creating additional SHV configurations, 184–186
 - default configuration, 184

T

- Take No Action setting, 257
- Task Scheduler, Remote Group Policy update, 220–221
- TCP/IP settings, configuring Hyper-V VMs, 340–341
- Teredo transition technology, 145
- Test-ADDSDomainControllerInstallation cmdlet, 113, 115
- Test-ADDSDomainControllerUninstallation cmdlet, 113, 118–119
- Test-ADDSDomainInstallation cmdlet, 113
- Test-ADDSEnvironmentInstallation cmdlet, 113–114
- Test-ADDSEnvironmentReadonlyDomainControllerAccountCreation cmdlet, 113, 117–118
- test failovers, 345–346
- TGT (ticket-granting ticket), 287
- thin provisioning, 36
- Throttling tab, 322
- ticket-granting ticket (TGT), 287
- Token-Decrypting certificates, 397
- Token Issuance section, Endpoints node of AD FS management console, 396
- Token-Signing certificates, 397
- tokens, user authentication, 389
- trust groups, 332
- Trust Relationships node (AD FS management console), 397–399
 - claims provider trusts, 398
 - claims rules for relying party trusts, 398–399
 - relying party trusts, 398

U

- UEFI (Unified Extensible Firmware Interface), 70
- Uninstall-ADDSDomainController cmdlet, 113, 118

- uninstalling domain controllers, Windows PowerShell, 118
- Uninstall-WindowsFeature cmdlet, 2
- unplanned failovers, 343–345
- Update-FSRMClassificationPropertyDefinition cmdlet, 293–294
- updating
 - Active Directory files/folders, 293–294
 - Group Policy
 - Group Policy Management Console, 216–218
 - Invoke-GpUpdate cmdlet, 218–220
- uploading certificates to Windows Azure, 316
- Use Computer Certificates setting (Authentication page), 161
- Use Force Tunneling option, 158
- Use Hardware Topology button, 79
- user authentication, AD FS, 390–392
- User Authentication setting (Authentication page), 161
- user claims types, configuring claims-based authentication, 288–290
- UserName option (Add-PswaAuthorizationRule cmdlet), 17

V

- verifying DirectAccess configuration, 168–170
- VHD (virtual hard disk) files, 236
 - DISM utility, 18
- VHDX disk format, 82–85
- VHDX files, DISM utility, 18
- VHDX sharing, 239
- viewing IP utilization thresholds, 378–379
- virtual disks, creating, 34–37
- virtual hard disk (VHD) files, 236
- virtual machines
 - application monitoring, 254–260
 - configuring Hyper-V, 333–342
 - failover TCP/IP settings, 340–341
 - resynchronizing primary and replica VMs, 341–342
 - configuring storage, 82–89
 - Fibre Channel adapters, 86–88
 - Storage QoS, 88–89
 - VHDX disk format, 82–85
 - creating and configuring settings, 67–80
 - Dynamic Memory, 74–77
 - enhanced session mode, 72–74

- generation 1 and 2 VMs, 69–71
- Hyper-V Module in Windows PowerShell, 68–69
- NUMA (non-uniform memory access), 79
- RemoteFX, 79–80
- Resource Metering, 77–78
- migration, 260–277
 - live migration, 261–273
 - storage migration, 274–276
 - VM network health protection, 276–277
- resource pools, 134–135
- virtual networks, 90–100, 124
 - advanced features for virtual network adapters, 99–100
 - bandwidth management, 97–98
 - network isolation, 93–95
 - SR-IOV, 95–98
 - virtual switch extensions, 91–93
- virtual private networks. *See* VPNs
- virtual switch extensions, 91–93
- Virtual Switch Manager, 91–92
- virtual switch name matching, 272–273
- VM Monitoring feature, 254
- Volume Shadow Copy Service (VSS) copies, 338
- VPNs (virtual private networks)
 - enforcement, 179
 - VPN icon, 170
- VSSAdmin tool, 324
- VSS (Volume Shadow Copy Service) copies, 338
- uploading certificates, 316
- Windows Deployment Services (WDS), 137
- Windows Firewall Remote Management, 52
- Windows Imaging (WIM), DISM utility, 18
- Windows Management Framework updates, 56
- Windows Management Instrumentation (WMI), 52–54
- Windows PowerShell
 - cmdlets. *See* individual names of cmdlets
 - deploying roles and features on remote servers, 8–9
 - DSC (Desired State Configuration), 13–17
 - Hyper-V Module, 68–69
 - installing domain controllers, 113–118
 - adding to existing domains, 115–116
 - first controller in new domain of existing forests, 116–117
 - new forests, 114–115
 - RODC accounts, 117–118
 - monitoring servers, 136–137
 - performing Windows Azure Backups, 324–326
 - restoring deleted objects, 207–208
 - uninstalling domain controllers, 118
- Windows PowerShell Classifier classification method, 297
- Windows PowerShell Web Access (PSWA), 16–17
- Windows Remote Management (WinRM), 53–54
- Windows Security Health Validator, 179, 183
- Windows Server 2012 R2, role-based access control for IPAM, 381–382
- Windows Server Update Services (WSUS), 137
- Windows SHVs, 179
- Winrm Quickconfig command, 55
- WinRM (Windows Remote Management), 53–54
- WitnessDynamicWeight cluster property, 248
- witnesses, configuring failover clustering, 241
- wizards
 - Active Directory Domain Services Configuration, 107–127
 - Add Relying Party Trust, 398
 - Add Roles And Features, 9–12, 361
 - AD FS Configuration, 392–394
 - Configure NAP, 181
 - Configure Self-Updating Options, 245
 - configuring DirectAccess
 - Application Server Setup, 166–167
 - Client Setup, 156–159
 - Infrastructure Server Setup, 162–166
 - Remote Access Server Setup, 159–162
 - Remote Access Setup, 154

W

- WDS (Windows Deployment Services), 137
- Web Application Proxy, 390
- WIM (Windows Imaging), DISM utility, 18
- Windows Azure, installing domain controllers, 120–127
- Windows Azure Backup, 313–326
 - Back Up Now option, 321–322
 - bandwidth throttling, 322–324
 - certificate requirements, 314
 - creating an online backup schedule, 318–321
 - creating backup vault in management portal, 315–316
 - creating self-signed certificate with Makecert.exe utility, 314–315
 - performing in Windows PowerShell, 324–326
 - recovering data, 322
 - registering servers, 317

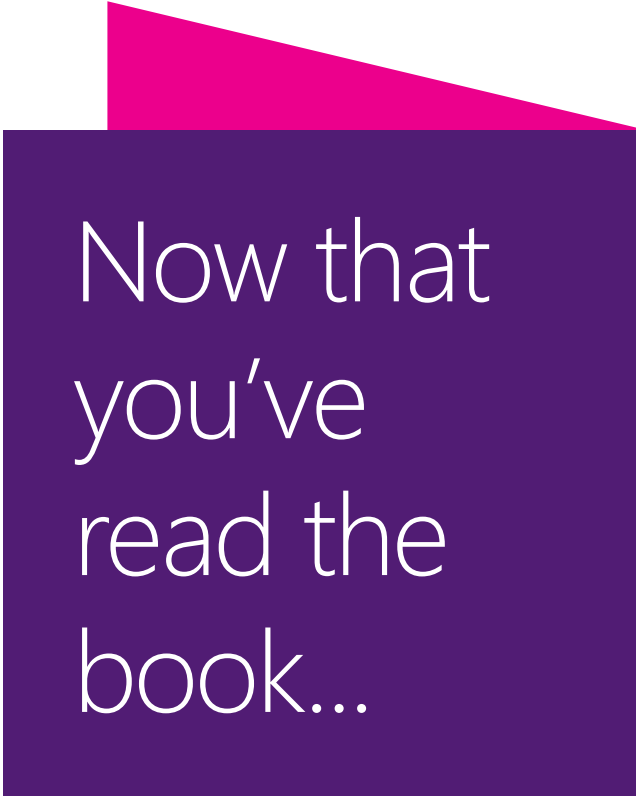
WMI (Windows Management Instrumentation)

- Edit Virtual Hard Disk, 84–85
- Enable Replication, 334–340
- Getting Started, 154–155
- High Availability, 251, 348
- Migrate A Cluster, 259
- Move, 274
- New Inbound Rule, 58
- New Storage Pool, 33, 234
- New Virtual Disk, 34–35
- Provision IPAM, 364
- Register Server, 317
- Reverse Replication, 344
- Schedule Backup, 318–321
- WMI (Windows Management Instrumentation), 52–54
- Workplace Join, AD FS, 403–406
- World Wide Node Name (WWNN), 87
- World Wide Port Name (WWPN), 87
- Write-Output \$UtilizationReport cmdlet, 78
- WS-Management Protocol, 52
- WSUS (Windows Server Update Services), 137
- WWNN (World Wide Node Name), 87
- WWPN (World Wide Port Name), 87

About the Author



J.C. MACKIN (MCSA, MCSE, MCT) is a writer, analyst, and trainer who has specialized in Windows networks since Windows NT 4.0. He has written or co-authored more than 10 books about Windows Server administration and certification. You can follow him on Twitter at @jcmackin.



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!

