Microsoft

MCTS EXAM

70-643

2

SECOND EDITION

Covers Windows Server 2008 R2

Configuring Windows Server® 2008 Applications Infrastructure

J.C. Mackin

SELF-PACED

Training Kit

Microsoft®

MCTS Self-Paced Training Kit (Exam 70-643): Configuring Windows Server® 2008 Applications Infrastructure (2nd Edition)

J.C. Mackin
Anil Desai

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

Library of Congress Control Number: 2011928628
ISBN: 978-0-7356-4878-4

Printed and bound in the United States of America.

3 4 5 6 7 8 9 10 11 QG 7 6 5 4 3 2

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.

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: Jeff Koch
Developmental Editor: Karen Szall
Project Editor: Carol Dillingham
Editorial Production: nSight, Inc.
Technical Reviewer: Mitch Tulloch; Technical Review services provided by Content Master, a member of CM Group, Ltd.
Copyeditor: Kerin Forsyth
Indexer: Lucie Haskins
Cover: Twist Creative • Seattle

[2012-01-27]

Exam 70-643: TS: Windows Server 2008 Applications Infrastructure, Configuring

OBJECTIVE	CHAPTER	LESSON
DEPLOYING SERVERS		
Deploy images by using Windows Deployment Services.	1	1, 2
Configure Microsoft Windows activation.	1	4
Configure Windows Server Hyper-V and virtual machines.	1	3
Configure high availability.	2	2
Configure storage.	2	1
CONFIGURING REMOTE DESKTOP SERVICES		
Configure RemoteApp and Remote Desktop Web Access.	4	3
Configure Remote Desktop Gateway (RD Gateway).	4	2
Configure Remote Desktop Connection Broker.	3	2
Configure and monitor Remote Desktop resources.	4	1
Configure Remote Desktop licensing.	3	1, 2
Configure Remote Desktop Session Host.	3 4	1, 2 1
CONFIGURING A WEB SERVICES INFRASTRUCTURE		
Configure Web applications.	5 6	2 2
Manage Web sites.	5 6	2 1
Configure a File Transfer Protocol (FTP) server.	7	1
Configure Simple Mail Transfer Protocol (SMTP).	7	2
Manage the Web Server (IIS) role.	5	1, 2
Configure SSL security.	6	2
Configure Web site authentication and permissions.	6	1, 2
CONFIGURING NETWORK APPLICATION SERVICES		
Manage the Streaming Media Services role.	8	1
Secure streaming media.	8	1
Configure SharePoint Foundation options.	9	1
Configure SharePoint Foundation integration.	9	1

Exam Objectives The exam objectives listed here are current as of this book's publication date. Exam objectives are subject to change at any time without prior notice and at Microsoft's sole discretion. Please visit the Microsoft Learning website for the most current listing of exam objectives: http://www.microsoft.com/learning/en/us/exams/70-643.mspx.

Contents at a Glance

Contents

Introduction

This training kit is designed for information technology (IT) professionals who support or plan to support Windows Server 2008 R2 networks and who also plan to take the Microsoft Certified Technology Specialist (MCTS) 70-643 exam. It is assumed that before you begin using this kit, you have a solid, foundation-level understanding of Windows client and server operating systems and common Internet technologies.

The material covered in this training kit and on the 70-643 exam relates to the technologies in a Windows Server 2008 R2 network that support remote access to operating systems, web content, media, and applications. The topics in this training kit cover what you need to know for the exam as described on the Skills Measured tab for the exam, which is available at *http://www.microsoft.com/learning/en/us/Exam.aspx?ID=70-643&locale=en-us#tab2*.

By using this training kit, you learn how to do the following:

- Deploy Windows servers and clients across a network by using Windows Deployment Services and the Windows Automated Installation Kit (AIK) for Windows 7
- Configure Hyper-V
- Configure an activation infrastructure
- Configure high-availability storage solutions for servers
- Configure and manage Remote Desktop Services in Windows Server 2008 R2
- Configure and manage Internet Information Services 7.5
- Configure Windows Media Services
- Configure Microsoft SharePoint Foundation 2010

Refer to the objective mapping page in the front of this book to see where in the book each exam objective is covered.

Hardware Requirements (Hyper-V)

You should use the Hyper-V virtualization platform on a single physical computer as the computer lab environment for this training kit. Hyper-V is a feature of Windows Server 2008 and Windows Server 2008 R2 and is a topic that is covered on the 70-643 exam. Because the task of setting up this computer lab environment in Hyper-V helps you develop knowledge and skills required for the 70-643 exam, using an alternative virtualization software platform such as VirtualBox is not recommended.

You need only one physical computer to perform the exercises in this book. This physical host computer should meet the following minimum hardware requirements:

- An x64-based processor that includes both hardware-assisted virtualization (AMD-V or Intel VT) and hardware data execution protection. (On AMD systems, the data execution protection feature is called the No Execute or NX bit. On Intel systems, this feature is called the Execute Disable or XD bit.) These features must also be enabled in the BIOS.

- 4.0 GB of RAM (more is recommended).
- 90 GB of available hard disk space.
- DVD-ROM drive.
- Internet connectivity.

Software Requirements

The following software is required to complete the practice exercises:

- Windows Server 2008 R2. You can download an evaluation edition of Windows Server 2008 R2 at the Microsoft Download Center at *http://www.microsoft.com/downloads*.
- The Windows Automated Installation Kit (AIK) for Windows 7. You can download the Windows AIK for Windows 7 at the Microsoft Download Center at *http://www.microsoft.com/downloads*. This kit is installed on top of Windows Server 2008 R2 on the Server1 computer.

Practice Setup Instructions

After you have installed Windows Server 2008 R2 on your physical computer and have established an Internet connection, you are ready to begin the practice setup. This setup occurs in five phases. In the first phase, you install the Hyper-V server role and create a virtual network named contoso.local. In the second phase, you create three virtual machines. Figure I-1 shows the virtual hardware configuration of the virtual machines as they appear after this second phase.

IMPORTANT **DOWNLOAD REQUIRED SOFTWARE**

Before you begin preparing the practice computers, you must have a copy of Windows Server 2008 R2 (either as an .iso file or as a DVD) and the Windows Automated Installation Kit (either as an .iso file or as a DVD).

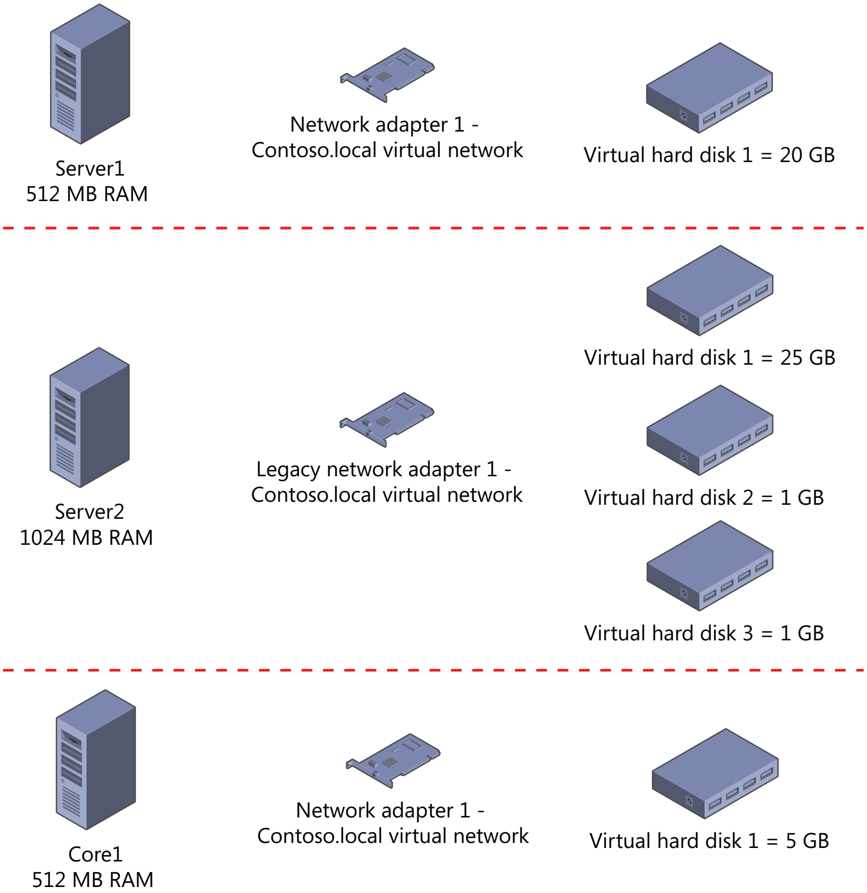
Server1
512 MB RAM
Network adapter 1 -
Contoso.local virtual network
Virtual hard disk 1 = 20 GB
Virtual hard disk 1 = 25 GB
Server2
1024 MB RAM
Legacy network adapter 1 -
Contoso.local virtual network
Virtual hard disk 2 = 1 GB
Virtual hard disk 3 = 1 GB
Core1
512 MB RAM
Network adapter 1 -
Contoso.local virtual network
Virtual hard disk 1 = 5 GB

FIGURE I-1 Hardware configuration for the three computers in Hyper-V.

In the third phase of the practice setup, you configure the software for the Server1 and Core1 machines. (No software configuration is necessary for Server2 because this computer must be left as a virtual bare-metal machine.)

The fourth phase of practice setup describes the configuration necessary to provide an Internet connection for all computers. By performing these steps, you add a second virtual network adapter to Server1 and configure Network Address Translation (NAT) across its two adapters, as shown in Figure I-2.

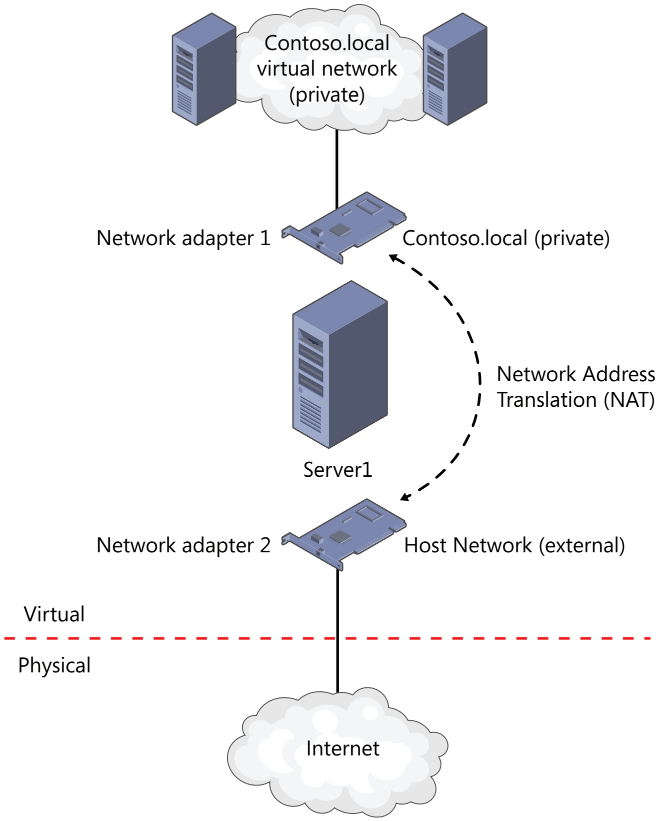

FIGURE I-2 Providing an Internet connection for all three computers in Hyper-V.

In the fifth and final phase of the practice setup, you activate the Server1 and Core1 servers over the Internet.

Phase 1: Install and Configure the Hyper-V Server Role on the Host Server

Perform the following steps to create and prepare the virtual environment for the lab computers.

Add the Hyper-V Server Role

On the host computer running Windows Server 2008 R2, click Add Roles in the Initial Configuration Tasks window and then add the Hyper-V server role.

After Hyper-V is installed, use the Hyper-V Manager administrative tool to open Virtual Network Manager and add the following two networks:

- Virtual Network 1
  - Type = Private
  - Name = Contoso.local
- Virtual Network 2
  - Type = External
  - Name = Host Network
  - Select an external adapter with Internet access

Phase 2: Create the Virtual Machines

Perform the following steps to create the virtual machines for this training kit.

Create the Server1 Virtual Machine

In Hyper-V Manager, right-click the local server node in the console tree, click New, and then click Virtual Machine to launch the New Virtual Machine Wizard. Use the New Virtual Machine Wizard to specify the following settings:

- Name: Server1
- RAM: 512 MB
- Connection: Contoso.local
- Create a virtual hard disk
- Virtual hard disk size: 20 GB
- Install an operating system later
- Do not add a second adapter yet

Create the Core1 Virtual Machine

Use the New Virtual Machine Wizard to create a second virtual machine. Configure all settings identically to those of the Server1 machine, except in the following two cases:

- Name: Core1
- Virtual Hard Disk Size: 5 GB

Create the Server2 Virtual Machine

Use the New Virtual Machine Wizard to create the third virtual machine. Configure all settings identically to those of the Server1 machine, except in the following cases:

- Name: Server2
- RAM: 1024 MB
- Hard Disk Name: Server2DiskA.Vhd
- Hard Disk Size: 25 GB
- Connection: Not Connected

IMPORTANT **USE A LEGACY NETWORK ADAPTER FOR SERVER2**

Server2 needs a Legacy Network Adapter for you to perform the exercises related to Windows Deployment Services found in Lesson 2 of Chapter 1, "Implementing and Configuring a Windows Deployment Infrastructure."

Replace the Default Network Adapter with a Legacy Network Adapter

The default network adapter assigned in Hyper-V is incompatible with network-based installations. For this reason, you must replace the default adapter with the Legacy Network Adapter. Right-click Server2 in Hyper-V manager and click Settings. In the Settings For Server2 dialog box, select the hardware named Network Adapter and then click Remove. Click Add Hardware, select Legacy Network Adapter, and click Add. You should see a new Legacy Network Adapter in the list of hardware. Select the new Legacy Network Adapter, assign it to the Contoso.localnetwork from the drop-down list, and click Apply.

Attach a Second and Third Hard Disk to Server2

In Hyper-V Manager, right-click Server2 and then click Settings. In the Settings For Server2 window, select IDE Controller 0 in the list of hardware, select Hard Drive, and then click Add. Click New to start the New Virtual Hard Disk Wizard.

In the New Virtual Hard Disk Wizard, specify the following:

- Dynamically Expanding
- Name: Server2DiskB.vhd
- Location: Browse To The \Virtual Machines\Server2\ folder
- Create A New Blank Virtual Hard Disk
- Size: 1 GB

In the Settings For Server2 window, select IDE Controller 1. Repeat the same process to create and attach a third 1 GB virtual hard disk named Server2DiskC.vhd. Click Apply to save the changes.

Configure Startup Order in the BIOS Setting

In the Settings For Server2 window, click BIOS in the list of hardware. Change the startup order so that Legacy Network Adapter is listed first. Click Apply to save the changes.

After you have finished creating and configuring the Server2 virtual machine, the Settings window should look like the one shown in Figure I-3.

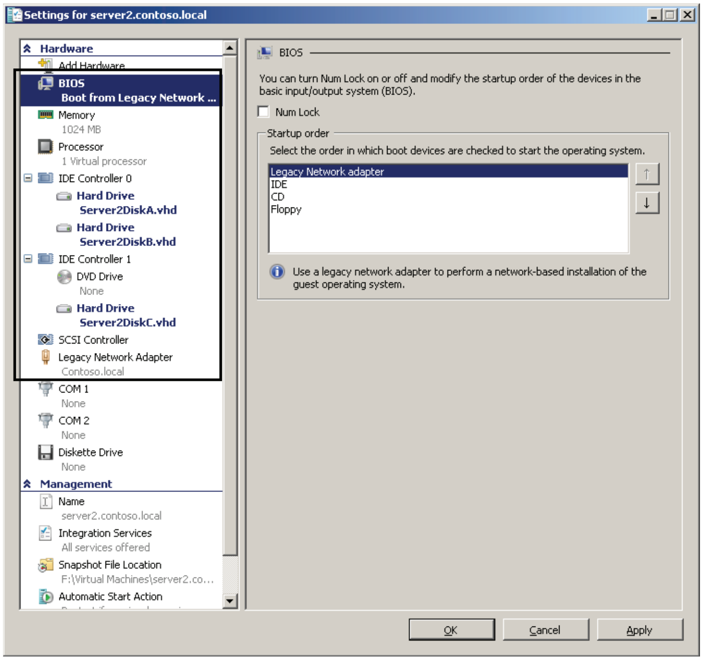

Settings for server2.contoso.local

Hardware
- Add Hardware
- BIOS
  Boot from Legacy Network ...
- Memory
  1024 MB
- Processor
  1 Virtual processor
- IDE Controller 0
  - Hard Drive
    Server2DiskA.vhd
  - Hard Drive
    Server2DiskB.vhd
- IDE Controller 1
  - DVD Drive
    None
  - Hard Drive
    Server2DiskC.vhd
- SCSI Controller
- Legacy Network Adapter
  Contoso.local
- COM 1
  None
- COM 2
  None
- Diskette Drive
  None

Management
- Name
  server2.contoso.local
- Integration Services
  All services offered
- Snapshot File Location
  F:\Virtual Machines\server2.co...
- Automatic Start Action

BIOS

You can turn Num Lock on or off and modify the startup order of the devices in the basic input/output system (BIOS).

□ Num Lock

Startup order

Select the order in which boot devices are checked to start the operating system.

Legacy Network adapter
IDE
CD
Floppy

Use a legacy network adapter to perform a network-based installation of the guest operating system.

FIGURE I-3 The Server2 virtual machine should have three attached virtual hard disks and a Legacy Network Adapter. It should be configured to start from the Legacy Network Adapter.

When you are finished, click OK to save the changes you have made in the Settings For Server2 window.

Phase 3: Configure the Operating Systems on Server1 and Core1

Use the following instructions to configure the Server1 and Core1 computers.

Configure Server1

Server1 will be used as a DHCP server, DNS server, and Active Directory domain controller for the contoso.local domain. Server1 must also have the Windows AIK for Windows 7 installed. Perform the following steps to meet the configuration requirements for the server.

1. In Hyper-V Manager, right-click Server1 and click Settings.
2. In the Settings For Server1 window, click DVD drive.
3. In the Media section, do one of the following:
   - Click Physical CD/DVD Drive and specify the physical CD/DVD-ROM drive on the host machine in which you will place the Windows Server 2008 R2 media.
   - Click Image File and browse to an ISO file that contains a disc image of the Windows Server 2008 R2 media.
4. Click OK to close the Settings For Server1 window.
5. In Hyper-V Manager, right-click Server1 and then click Connect.
6. In the Server1 On Localhost window, click Start on the menu bar. (This button is round and blue-green.)

   The Server1 computer starts, and the Windows Server 2008 R2 installation process begins.
7. Perform a default installation of Windows Server 2008 R2. Use the following guidelines:
   - If desired, choose a language and keyboard corresponding to your region.
   - Do not enter a product key at this time.
   - Choose Windows Server 2008 R2 Enterprise (Full Installation).
   - Install Windows in the default location (Disk 0 Unallocated Space).
   - Use a strong password of your choice when logging on as Administrator for the first time.
8. Configure the Local Area Connection on Server1. You can perform this step by using either the Initial Configuration Tasks window or a command prompt.

   If you prefer to use the Initial Configuration Tasks window, click Configure Networking, open the properties of the Local Area Connection, and then configure the properties of Internet Protocol version 4 (TCP/IPv4) with the following options and values:
   - Select Use The Following IP Address.
     - IP Address: 192.168.10.1
     - Subnet Mask: 255.255.255.0
     - Default Gateway: Leave blank

- Select Use The Following DNS Server Addresses.
    - Preferred DNS server: 192.168.10.1
    - Alternate DNS server: Leave blank

To configure these same IP settings at a command prompt instead, type the following two commands in succession:

```
netsh interface ipv4 set address “local area connection” static 192.168.10.1
255.255.255.0
netsh interface ipv4 set dns “local area connection” static 192.168.10.1
```

9. Configure the computer name. You can perform this step by using either the Initial Configuration Tasks window or a command prompt.

   In Initial Configuration Tasks, click Provide Computer Name And Domain. Click Change and specify the computer name as Server1. Do not specify a domain at this time.

   To set the computer name at the command prompt instead, type the following command:

   ```
   netdom renamecomputer %computername% /newname:Server1 /reboot
   ```

10. Use the Run box from the Start menu to run Dcpromo and configure Server1 as a domain controller in a new Active Directory domain named contoso.local. Specify the following options in the Active Directory Domain Services Installation Wizard:
    - Create a New Domain In A New Forest.
    - FQDN Of The Forest Root: contoso.local.
    - Forest Functional Level: Windows Server 2008 R2.
    - Additional Domain Controller Options: DNS Server (Default).
    - If you are warned that the computer has a dynamically assigned IP address, click Yes.
    - If you are warned that a delegation for this DNS server cannot be created, click Yes.
    - Locations for database, log files, and Sysvol: Leave defaults.
    - Directory Services Restore Mode Administrator Password: Any strong password of your choice.
11. After the Active Directory Domain Services Installation Wizard completes, restart Server1 immediately and then log on to the contoso.local domain from Server1 as CONTOSO\Administrator.

IMPORTANT **HOW DO YOU LOG ON TO A COMPUTER IN HYPER-V?**

Note that in Hyper-V, you must use the *Ctrl+Alt+End* command to enter the Ctrl+Alt+Del keystroke. You can also choose Ctrl+Alt+Delete from the Action menu.

12. Add the DHCP Server role. In the Initial Configuration Tasks window, click Add Roles. Use the Add Roles Wizard to add the DHCP Server role with the following options:
    - Network Connection Bindings: Default. (Leave 192.168.10.1 selected.)
    - IPv4 DNS Server Settings:
        - Parent Domain: contoso.local
        - Preferred DNS Server IPv4 Address: 192.168.10.1
        - Alternate DNS Server IPv4 Address: Leave blank
    - WINS Server Settings: WINS is not required.
    - Add a DHCP scope with the following specifications:
        - Scope Name: Contoso.local
        - Starting IP Address: 192.168.10.2
        - Ending IP Address: 192.168.10.10
        - Subnet Mask: 255.255.255.0
        - Default Gateway: 192.168.10.1 (This assumes an Internet access configuration as described in Phase 3 of the practice setup instructions.)
        - Subnet Type: Wired
        - Activate This Scope: Leave selected
        - DHCPv6 Stateless Mode: Leave default
        - IPv6 DNS Server Settings: Leave default
        - Authorize DHCP Server: Leave default
13. Create and name three domain administrator accounts. To do so, use the following step-by-step instructions.
    a. In the Active Directory Users And Computers administrative tool, expand the *contoso.local* node in the console tree and then select the Users folder.
    b. Right-click the Users folder, point to New on the shortcut menu, and then click User.
    c. In the New Object – User dialog box, type **ContosoAdmin1** in the Full Name and User Logon Name text boxes and then click Next.
    d. Enter a password of your choice, click Next, and then click Finish.
    e. In the Active Directory Users And Computers console, locate the ContosoAdmin1 account you have just created in the details pane. Right-click the account and then click Add To A Group from the shortcut menu.
    f. In the Select Groups dialog box, type **domain admins** and then press Enter. In the Active Directory Domain Services message box, click OK.

   g. Create two additional domain administrator accounts, named ContosoAdmin2 and ContosoAdmin3, respectively, by using steps **b** through **f**.

   h. If desired, create an additional domain administrator account with your name.

14. Enable file sharing on Server1.

    a. In the Search Programs And Files box of the Start menu, type Manage Advanced Sharing Settings and then press Enter.

    b. In the Advanced Sharing Settings window, select the Turn On File And Printer Sharing option and then click Save Changes.

15. Install the Windows Automated Installation Kit for Windows 7 (Windows AIK) by using the Windows AIK DVD or .iso file you have downloaded from the Microsoft Download Center. To do so, use the following step-by-step instructions.

    a. Use the Settings For Server1 window in Hyper-V to mount the Windows AIK for Windows 7 .iso file as a DVD drive.

    b. In the AutoPlay window, use the Windows AIK for Windows 7 Setup link to install the Windows AIK for Windows 7 and any prerequisite components if necessary.

Configure Core1

Core1 will act as a member server in the contoso.local domain. Use the following instructions to configure the Core1 server.

1. In Hyper-V Manager, right-click Core1 and click Settings.
2. In the Settings For Core1 window, select DVD Drive from the list of hardware.
3. In the Media section, do one of the following:
   - Click Physical CD/DVD Drive and specify the physical CD/DVD-ROM drive on the host machine in which you will place the Windows Server 2008 R2 media.
   - Click Image File and browse to an ISO file that contains a disc image of the Windows Server 2008 R2 media.
4. Click OK to close the Settings For Core1 window.
5. In Hyper-V Manager, right-click Core1 and then click Connect.
6. In the Core1 On Localhost window, click Start on the menu bar. (This button is round and blue-green.)

   The Core1 computer starts, and the Windows Server 2008 R2 installation process begins. Perform a default installation by using the following guidelines:

   - If desired, choose a language and keyboard corresponding to your region.
   - Do not enter a product key at this time.
   - Choose Windows Server 2008 R2 Enterprise (Server Core Installation).

- Install Windows in the default location (Disk 0 Unallocated Space).
- To log on for the first time, specify a user of Administrator with a strong password of your choice.

7. Verify the IP configuration. At the command prompt, type **ipconfig /all** to ensure that Core1 has received an IP configuration from Server1.
8. At the command prompt, type **sconfig**.
9. Use option 2 in the Server Configuration utility to change the computer name to Core1 and then agree to restart the computer.
10. Log on to Core1 as Administrator and type **sconfig** at the command prompt to start the Server Configuration utility again.
11. Use option 1 in the Server Configuration utility to join Core1 to the Contoso.local domain and then agree to restart the computer.

Phase 4: Configure Internet Access for the Contoso.local Network

In this phase, you add to Server1 a second adapter that is bound to a physical network adapter on the physical host machine. You then configure network address translation (NAT) on Server1.

Add and Configure a Second Virtual Adapter on Server1

Complete the following steps to add and configure a second virtual adapter on Server1.

1. Shut down Server1. Open Server1 settings in Hyper-V Manager.
2. In the Settings For Server1 dialog box, select Add Hardware, select Network Adapter, and then click Add. When the new Network Adapter appears in the list of hardware, assign it to the network named Host Network.

   The physical adapter on the host computer should already have its own IP address and be able to communicate with the Internet.
3. Start and log on to Server1.
4. In Server Manager, click Add Roles. Use the following information to complete the Add Roles Wizard:
   - Select Server Roles: Network Policy And Access Services
   - Select Role Services: Routing And Remote Access Services (Do not select any other role services at this time.)

Configure NAT on Server1

Use the following step-by-step instructions to configure NAT on Server1.

1. After you have installed the Routing And Remote Access Services role service, open the Routing And Remote Access administrative tool through the Start menu.
2. In the Routing And Remote Access console tree, right-click the *Server1* node and then click Configure And Enable Routing And Remote Access.
3. Specify the following settings in the Routing And Remote Access Server Setup Wizard:
   - On the Configuration page, click Network Address Translation (NAT).
   - On the NAT Internet Connection page, select Local Area Connection 2 as the public interface to connect to the Internet.
4. In Server Manager, select the *Server Manager* node. In the Security Information area of the details pane, click Configure IE ESC. Select the option to turn IE ESC off for Administrators.
5. Open Internet Explorer and select Internet Options from the Tools menu. Set the home page to an Internet-based webpage of your choice.
6. Verify Internet connectivity in Internet Explorer by clicking the Home icon.

Phase 5: Activate the Servers (Recommended)

Perform the following steps if you have product keys for both Server1 and Core1.

1. Activate Server1. Open the System Control Panel and select the option to change the product key. Type the product key when prompted and click Next.

   Windows automatically activates over the Internet.
2. Activate Core1 by using the following step-by-step procedure:
   a. Log on to contoso.local from Core1 as a domain administrator and then type the following command to install the new product key, where *productkey* is your product key (with dashes):

      ```
      slmgr -ipk productkey
      ```

   b. When you receive a message indicating that the product key was installed successfully, type the following command to activate Windows:

      ```
      slmgr –ato
      ```

EXAM TIP

You need to know these last two commands for the 70-643 exam.

c. After you receive a message indicating that the product has been activated successfully, you can shut down Core1 by typing the following command:

```
shutdown /s /t 0
```

Using the Companion CD

A companion CD is included with this training kit. The companion CD contains the following:

- **Practice tests** You can reinforce your understanding of how to configure Windows Server 2008 R2 by using electronic practice tests you customize to meet your needs from the pool of Lesson Review questions in this book. Alternatively, you can practice for the 70-643 certification exam by using tests created from a pool of 200 realistic exam questions, which give you many practice exams to ensure that you are prepared.
- **Webcasts and Videos** To supplement your learning, the CD includes Microsoft-sponsored webcasts and videos from experts. These webcasts and videos are lectures and demonstrations that provide additional information about subjects covered in the book.
- **An eBook** An electronic version (eBook) of this book is included for when you do not want to carry the printed book with you. The eBook is in Portable Document Format (PDF), which is viewable by using Adobe Acrobat or Adobe Reader, and in XML Paper Specification (XML).

NOTE **COMPANION CONTENT FOR DIGITAL BOOK READERS**

If you bought a digital-only edition of this book, you can enjoy select content from the print edition's companion CD. Visit *http://go.microsoft.com/FWLink/?Linkid=220878* to get your downloadable content.

How to Install the Practice Tests

To install the practice test software from the companion CD to your hard disk, do the following:

1. Insert the companion CD into your CD drive and accept the license agreement.

   A CD menu appears.

NOTE **IF THE CD MENU DOES NOT APPEAR**

If the CD menu or the license agreement does not appear, AutoRun might be disabled on your computer. Refer to the Readme.txt file on the CD-ROM for alternate installation instructions.

2. Click Practice Tests and follow the instructions on the screen.

How to Use the Practice Tests

To start the practice test software, follow these steps.

1. Click Start\All Programs\Microsoft Press Training Kit Exam Prep.

   A window appears that shows all the Microsoft Press training kit exam prep suites installed on your computer.

2. Double-click the lesson review or practice test you want to use.

> ***NOTE*** **LESSON REVIEWS VS. PRACTICE TESTS**
>
> **Select the (70-643) Configuring Windows Server 2008 Applications Infrastructure (2nd Edition) lesson review to use the questions from the "Lesson Review" sections of this book. Select the (70-643) Configuring Windows Server 2008 Applications Infrastructure (2nd Edition) practice test to use a pool of 200 questions similar to those that appear on the 70-643 certification exam.**

Lesson Review Options

When you start a lesson review, the Custom Mode dialog box appears so that you can configure your test. You can click OK to accept the defaults, or you can customize the number of questions you want, how the practice test software works, the exam objectives to which you want the questions to relate, and whether you want your lesson review to be timed. If you are retaking a test, you can select whether you want to see all the questions again or only the questions you missed or did not answer.

After you click OK, your lesson review starts.

- To take the test, answer the questions and use the *Next* and *Previous* buttons to move from question to question.
- After you answer an individual question, if you want to see which answers are correct—along with an explanation of each correct answer—click Explanation.
- If you prefer to wait until the end of the test to see how you did, answer all the questions and then click Score Test. You will see a summary of the exam objectives you chose and the percentage of questions you got right overall and per objective. You can print a copy of your test, review your answers, or retake the test.

Practice Test Options

When you start a practice test, you choose whether to take the test in Certification Mode, Study Mode, or Custom Mode.

- **Certification Mode** Closely resembles the experience of taking a certification exam. The test has a set number of questions. It is timed, and you cannot pause and restart the timer.
- **Study Mode** Creates an untimed test in which you can review the correct answers and the explanations after you answer each question.
- **Custom Mode** Gives you full control over the test options so that you can customize them as you like.

In all modes, the user interface when you are taking the test is basically the same but with different options enabled or disabled, depending on the mode. The main options are discussed in the previous section, "Lesson Review Options."

When you review your answer to an individual practice test question, a "References" section is provided that lists where in the training kit you can find the information that relates to that question and provides links to other sources of information. After you click Test Results to score your entire practice test, you can click the Learning Plan tab to see a list of references for every objective.

How to Uninstall the Practice Tests

To uninstall the practice test software for a training kit, use the Programs And Features option in Windows Control Panel.

Support & Feedback

The following sections provide information on errata, book support, feedback, and contact information.

Errata

We have made every effort to ensure the accuracy of this book and its companion content. Any errors that have been reported since this book was published are listed on our Microsoft Press site:

http://go.microsoft.com/FWLink/?Linkid=220879

If you find an error that is not already listed, you can report it to us through the same page.

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

Please note that product support for Microsoft software is not offered through the addresses above.

We Want to Hear from You

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

http://www.microsoft.com/learning/booksurvey

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

Stay in Touch

Let us keep the conversation going! We are on Twitter: *http://twitter.com/MicrosoftPress*.

Preparing for the Exam

Microsoft certification exams are a great way to build your résumé 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 publicly 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.

Microsoft®
CERTIFIED
Technology Specialist

CHAPTER 2

Configuring Server Storage and Clusters

Storage area networks (SANs), host bus adapters (HBAs), and logical unit numbers (LUNs) were once the sole domain of storage specialists, far removed from the expertise of your average Windows administrator. However, the arrival of new technologies, such as the Windows Virtual Disk service and Internet SCSI (iSCSI), along with the increasingly complex realities of enterprise storage, has brought these once-specialized topics into the realm of Windows Server 2008 administration. To be an effective Windows server administrator today, you still need to know the difference between the various RAID levels, but you also need to know quite a bit more about advanced server storage technologies.

This chapter introduces you to the basics of disk management in Windows Server 2008 R2, along with more advanced storage technologies such as SANs. The chapter then builds upon this storage information to introduce the various clustering technologies available in Windows Server 2008 R2.

Exam objectives in this chapter:

- Configure storage.
- Configure high availability.

Lessons in this chapter:

Before You Begin

To complete the lessons in this chapter, you must have:

- A computer named Server2 that is running Windows Server 2008 R2. Beyond the disk on which the operating system is installed, Server2 must be equipped with two additional hard disks of equal size.
- A basic understanding of Windows administration.

Lesson 1: Configuring Server Storage

A variety of server storage solutions is available for corporate networks, and Windows Server 2008 R2 connects to these technologies in new ways. This lesson introduces you to the major server storage types and the tools built into Windows Server 2008 R2 you can use to manage them.

After this lesson, you will be able to:

- Understand the basic features of direct-attached storage, network-attached storage, and storage-area networks.
- Know the function of the Virtual Disk service.
- Understand the features of simple, spanned, striped, mirrored, and RAID-5 volumes.
- Use the Disk Management console to create the various volume types.

Estimated lesson time: 80 minutes

REAL WORLD

J.C. Mackin

Although you cannot create them in Windows, the RAID levels known as RAID 0+1 and RAID 1+0 are becoming increasingly common in the real world. A RAID 0+1 (or 01) is a *mirror of stripes*, essentially twin copies of a striped volume. This type of RAID is constructed by creating RAID 0 sets and then mirroring them. A RAID 1+0 (or 10), alternatively, is a *stripe of mirrors* in which the data is striped across multiple mirrored sets. You construct this type of RAID by first creating a series of mirror sets and then building a RAID 0 set across the mirror sets.

Both of these solutions allocate 50 percent of the disks for fault tolerance, and both offer excellent read and write performance. RAID 1+0, however, offers a better chance for recoverability if more than one disk fails.

Note also that the naming conventions for these two RAID levels are not firmly established. Some companies (including Microsoft) might generally refer to both RAID 01 and 10 as 0+1. If you need to clarify your requirements to vendors, you are better off specifying either a *mirror of stripes* or a *stripe of mirrors*.

Understanding Server Storage Technologies

As the demand for server storage has grown, so too has the number of new storage technologies. Over the years, the range of server storage options has broadened from simple direct-attached storage (DAS) to network-attached storage (NAS) and, most recently, to Fibre Channel (FC) and iSCSI SANs.

Direct-Attached Storage

DAS is storage attached to one server only. Examples of DAS solutions are a set of internal hard disks within a server or a rack-mounted RAID connected to a server through a SCSI or FC controller. The main feature of DAS is that it provides a single server with fast, *block-based* data access to storage directly through an internal or external bus. (Block-based, as opposed to file-based, means that data is moved in unformatted blocks rather than in formatted files.) DAS is an affordable solution for servers that need good performance and do not need enormous amounts of storage. For example, DAS is often suitable for infrastructure servers, such as DNS, WINS and DHCP servers, and domain controllers. File servers and web servers can also run well on a server with DAS.

The main limitation of DAS is that it is directly accessible from a single server only, which leads to inefficient storage management. For example, Figure 2-1 shows a LAN in which all storage is attached directly to servers. Despite the web and App2 servers having excess storage, there is no easy way for these resources to be redeployed to either the Mail or App1 server, which need more storage space.

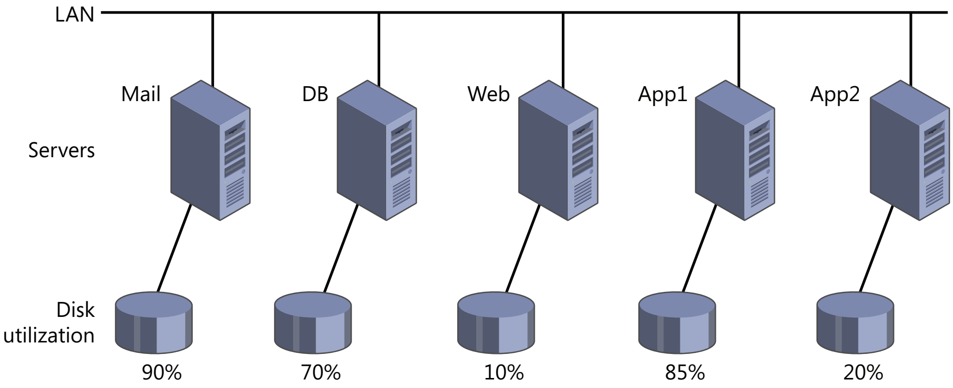

FIGURE 2-1 A network with only a DAS solution.

The main tool used for managing DAS in Windows is the Disk Management console. This tool, which you can access in Server Manager, enables you to partition disks and format volume sets. You can also use the Diskpart.exe command-line utility to perform the same functions available in Disk Management and to perform additional functions as well.

Network-Attached Storage

NAS is self-contained storage that other servers and clients can easily access over the network. A NAS device or appliance is a preconfigured server that runs an operating system specifically designed for handling file services. The main advantage of NAS is that it is simple to implement and can provide a large amount of storage space to clients and servers on a LAN. The downside of NAS is that, because your servers and clients access a NAS device over the LAN as opposed to over a local bus, access to data is slower and file-based as opposed to block-based. NAS performance is, therefore, almost always slower than that of DAS.

Because of its features and limitations, NAS is often a good fit for file servers, web servers, and other servers that don't need extremely fast access to data. In addition, NAS appliances come with their own management tools, which are typically web-based.

Figure 2-2 shows a network in which clients use a NAS appliance as a file server.

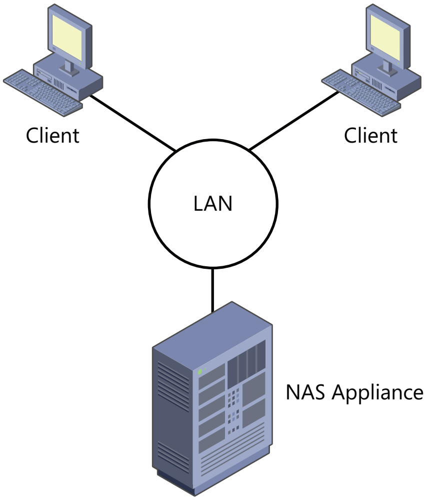

FIGURE 2-2 A LAN with a NAS appliance.

Storage-Area Networks

SANs are high-performance networks dedicated to delivering block data between servers and storage subsystems. From the point of view of the operating system, SAN storage appears as if it were installed locally. The most important characteristic that distinguishes a SAN from DAS is that in a SAN, the storage is not restricted to one server but is, in fact, available to any of a number of servers. (SAN storage can be moved from server to server, but outside of clustered file system environments, it is not accessible by more than one server at a time.)

NOTE **SAN VS. DAS**

Although DAS data transfer rates are typically faster than those of a SAN, the performance gap between DAS and SAN technologies is constantly shrinking. Despite the bus speed advantage offered by DAS, SANs are still considered preferable because the advantage SANs offer of shared storage outweighs the shortcoming of slightly lower access speeds.

A SAN is made up of special devices, including SAN network adapters, called HBAs, on the host servers, cables and switches that help route storage traffic, disk storage subsystems, and tape libraries. These hardware devices that connect servers and storage in a SAN are called the *SAN fabric*. All these devices are interconnected by fiber or copper. When connected to the fabric, the available storage is divided up into virtual partitions called logical unit numbers (LUNs), which then appear to servers as local disks.

SANs are designed to enable centralization of storage resources while eliminating the distance and connectivity limitations posed by DAS. For example, parallel SCSI bus architecture limits DAS to 16 devices at a maximum (including the controller) distance of 25 meters. Fibre Channel SANs extend this distance limitation to 10 km or more and enable an essentially unlimited number of devices to attach to the network. These advantages enable SANs to separate storage from individual servers and to pool unlimited storage on a network where that storage can be shared.

SANs are a good solution for servers that require fast access to very large amounts of data (especially block-based data). Such servers can include mail servers, backup servers, streaming media servers, application servers, and database servers. The use of SANs also enables efficient long-distance data replication, which is typically part of a disaster recovery (DR) solution.

Figure 2-3 illustrates a simple SAN.

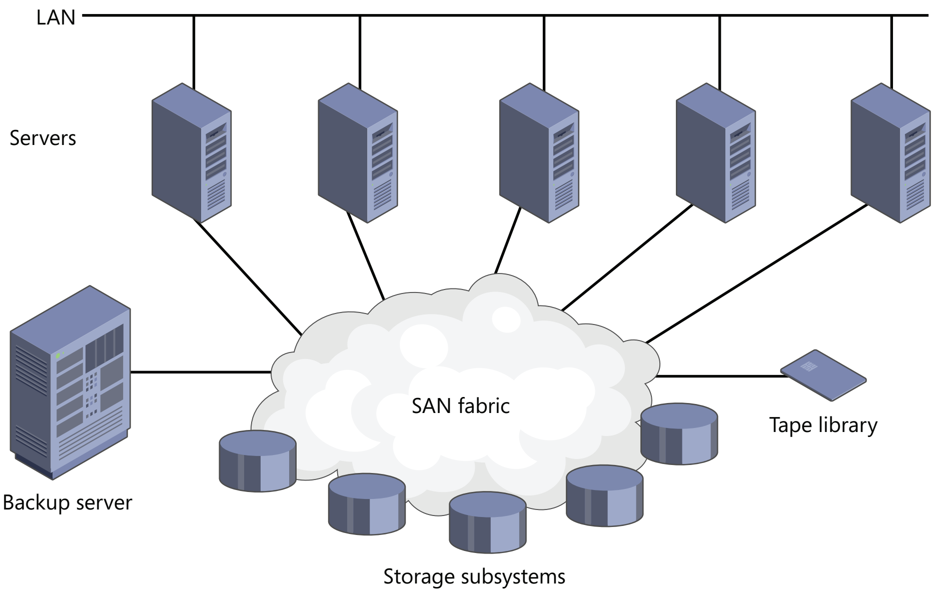

FIGURE 2-3 A sample storage area network (SAN).

SANs generally occur in two varieties: Fibre Channel and iSCSI.

FIBRE CHANNEL SANS

Fibre Channel (FC) delivers high-performance block input/output (I/O) to storage devices. Based on serial SCSI, FC is the oldest and most widely adopted SAN interconnect technology. Unlike parallel SCSI devices, FC devices do not need to arbitrate (or contend) for a shared bus. Instead, FC uses special switches to transmit information between multiple servers and storage devices at the same time.

The main advantage of FC is that it is the most widely implemented SAN technology and has, at least until recently, offered the best performance. The disadvantages of FC technology are the cost of its hardware and the complexity of its implementation. Fibre Channel network components include server HBAs, cabling, and switches. All these components are specialized for FC, lack interoperability among vendors, are relatively expensive, and require special expertise.

ISCSI SANS

Internet SCSI (iSCSI) is an industry standard developed to enable transmission of SCSI block commands over an Ethernet network by using the TCP/IP protocol. Servers communicate with iSCSI devices through a locally installed software agent known as an *iSCSI initiator.* The iSCSI initiator executes requests and receives responses from an *iSCSI target*, which itself can be the end-node storage device or an intermediary device such as a switch. For iSCSI fabrics, the network also includes one or more Internet Storage Name Service (iSNS) servers that, much like DNS servers on a LAN, provide discoverability and zoning of SAN resources.

By relying on TCP/IP, iSCSI SANs take advantage of networking devices and expertise that are widely available, a fact that makes iSCSI SANs generally simpler and less expensive to implement than FC SANs.

Aside from lower cost and greater ease of implementation, other advantages of iSCSI over FC include:

- **Connectivity over long distances** Organizations distributed over wide areas might have a series of unlinked SAN islands that the current FC connectivity limitation of 10 km cannot bridge. (There are new means of extending Fibre Channel connectivity up to several hundred kilometers, but these methods are both complex and costly.) In contrast, iSCSI can connect SANs in distant offices by using in-place metropolitan area networks (MANs) and wide-area networks (WANs).
- **Built-in security** No security measures are built into the Fibre Channel protocol. Instead, security is implemented primarily through limiting physical access to the SAN. In contrast to FC, the Microsoft implementation of the iSCSI protocol provides security for devices on the network by using the Challenge Handshake Authentication Protocol (CHAP) for authentication, and the Internet Protocol security (IPsec) standard for encryption. Because these methods of securing communications already exist in Windows networks, they can be readily extended from LANs to SANs.

NOTE **ISCSI SAN FABRIC**

An iSCSI SAN can use dedicated devices for its fabric, or it can rely on an organization's existing LAN, MAN, or WAN infrastructure. For both security and performance, a dedicated iSCSI network separating network traffic from storage traffic is recommended.

The main disadvantage of an iSCSI SAN is that, unless it is built with dedicated (and expensive) 10 Gbps Ethernet cabling and switches, the I/O transfer of iSCSI is slower than an FC-based SAN can deliver. And if you do choose to use 10-GB equipment for your iSCSI SAN instead of the much more common choice of gigabit Ethernet, the high cost of such a 10-GB solution would eliminate the price advantage of iSCSI relative to FC.

EXAM TIP

Vocabulary terms you should understand for the exam include LUNs, HBA, iSCSI initiator, iSCSI target, SAN fabric, and iSNS.

Configuring a SAN Connection with iSCSI Initiator

You can use the iSCSI Initiator built into Windows Server 2008 and Windows Server 2008 R2 to connect to an iSCSI SAN, configure the features of this iSCSI connection, and provision storage. To configure a SAN connection with iSCSI Initiator, select the tool from the Administrative Tools group in the Start menu. This step opens the Targets tab of the iSCSI Initiator Properties dialog box, as shown in Figure 2-4.

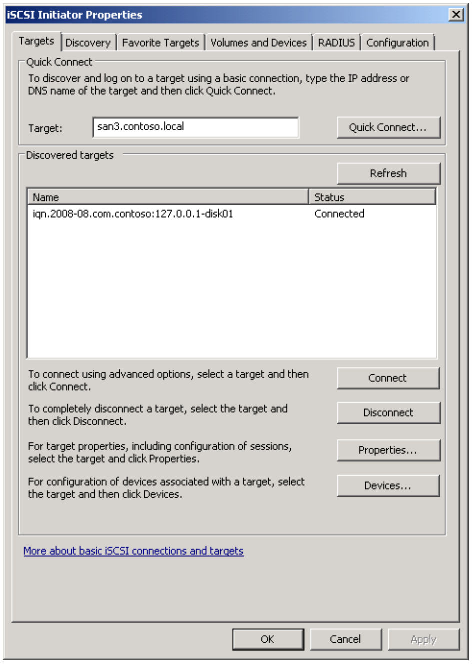

To discover and log on to a target using a basic connection, type the IP address or DNS name of the target and then click Quick Connect.

Target: san3.contoso.local

Name	Status
iqn.2008-08.com.contoso:127.0.0.1-disk01	Connected

To connect using advanced options, select a target and then click Connect.

To completely disconnect a target, select the target and then click Disconnect.

For target properties, including configuration of sessions, select the target and click Properties.

For configuration of devices associated with a target, select the target and then click Devices.

More about basic iSCSI connections and targets

FIGURE 2-4 iSCSI Initiator Properties.

To connect to an iSCSI SAN, specify an iSCSI target by name in the Target text box and then click Quick Connect. (Quick Connect is a new feature in Windows Server 2008 R2.) The Targets tab also provides access to Multipath I/O (MPIO) settings through the *Devices* and *Connect* buttons. MPIO enables you to configure multiple simultaneous connections to an iSCSI target so that if one adapter fails, another connection can continue processing I/O without any interruption of service. To enable MPIO, use the Add Features Wizard to add the Multipath I/O feature.

After you establish a connection to an iSCSI target, you can use the following tabs to configure the connection:

- Discovery

  On this tab, you can discover targets on specified portals and choose iSNS servers.

- Favorite Targets

  Use this tab to ensure that connections to selected iSCSI targets are restored every time the local computer restarts.

- Volumes And Devices

  This tab enables you to provision volumes and devices on targets and bind to them so they are readily available on system restart.

- RADIUS

This tab enables you to specify a RADIUS server and shared secret for the authentication of the iSCSI connection.

- Configuration

  This tab enables you to require negotiation of the CHAP authentication protocol and IPsec encryption for all connections to the local iSCSI Initiator. The tab also provides a unique identification number for the iSCSI Initator, which you can specify on a remote iSCSI target to configure a connection to the local machine.

EXAM TIP

For the 70-643 exam, you need to be familiar with the various configuration options available for iSCSI Initiator, including MPIO, favorites, RADIUS, and IPsec.

Other Tools for Managing SANs

Windows Server 2008 and Windows Server 2008 R2 include the Virtual Disk service (VDS), an application programming interface (API) that enables FC and iSCSI SAN hardware vendors to expose disk subsystems and SAN hardware to administrative tools in Windows. When vendor hardware includes the VDS hardware provider, you can manage that hardware within Windows Server 2008 and Windows Server 2008 R2 by using iSCSI Initiator and other tools, such as Disk Management, Storage Manager for SANs (SMfS), Storage Explorer, or the command-line tool, DiskRAID.exe.

- **Storage Manager for SANs** SMfS is available in Windows Server 2008 and Windows Server 2008 R2 as a feature you can add by using the Add Features Wizard. You can use SMfS to manage SANs by provisioning disks, creating LUNs, and assigning LUNs to different servers in the SAN.
- **Storage Explorer** Storage Explorer is available by default in Windows Server 2008 and Windows Server 2008 R2 through the Administrative Tools program group. You can use Storage Explorer to display detailed information about servers connected to the SAN and about fabric components such as HBAs, FC switches, and iSCSI initiators and targets. You can also use Storage Explorer to perform administrative tasks on an iSCSI fabric.
- **DiskRAID** DiskRAID is a command-line tool that enables you to manage LUNs in a VDS-enabled hardware RAID.

Managing Disks, Volumes, and Partitions in Windows Server 2008 R2

The main tool you can use to manage disks, volumes, and partitions in Windows Server 2008 and Windows Server 2008 R2 is Disk Management. With Disk Management, you can initialize disks, bring disks online or offline, create volumes within disks, format volumes, change disk partition styles, extend and shrink volumes, and create fault-tolerant disk sets.

To access Disk Management, you can type **Diskmgmt.msc** in the Run box, select Disk Management beneath the *Storage* node in Server Manager, or select the *Disk Management* node in the Computer Management console (accessible through Administrative Tools).

Disk Management is shown in Figure 2-5.

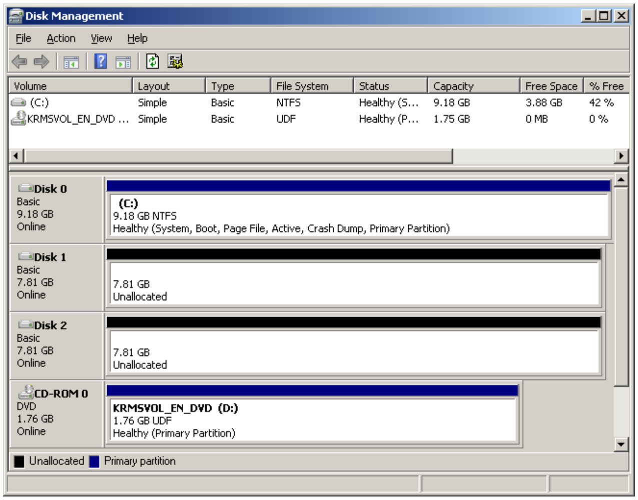

Volume	Layout	Type	File System	Status	Capacity	Free Space	% Free
(C:)	Simple	Basic	NTFS	Healthy (S...	9.18 GB	3.88 GB	42 %
KRMSVOL_EN_DVD ...	Simple	Basic	UDF	Healthy (P...	1.75 GB	0 MB	0 %

Disk 0 Basic 9.18 GB Online	**(C:)** 9.18 GB NTFS Healthy (System, Boot, Page File, Active, Crash Dump, Primary Partition)
Disk 1 Basic 7.81 GB Online	7.81 GB Unallocated
Disk 2 Basic 7.81 GB Online	7.81 GB Unallocated
CD-ROM 0 DVD 1.76 GB Online	**KRMSVOL_EN_DVD (D:)** 1.76 GB UDF Healthy (Primary Partition)

■ Unallocated ■ Primary partition

FIGURE 2-5 Disk Management in Windows Server 2008 R2.

Understanding Basic and Dynamic Disks

Disk Management enables you to manage both basic and dynamic disks.

By default, all disks are basic disks. A basic disk is a physical disk that contains primary partitions, extended partitions, or logical drives. The number of partitions you can create on a basic disk depends on the disk's *partition style*. On disks that use the master boot record (MBR) partition style, you can create up to four primary partitions per basic disk, or you can create up to three primary partitions and one extended partition. Within the one extended partition, you can then create unlimited logical drives. On basic disks that use the GUID partition table (GPT) partition style, you can create up to 128 primary partitions. Because GPT disks do not limit you to four partitions, you do not need to create extended partitions or logical drives. GPT disks are recommended for disks larger than 2 terabytes (TB) and for disks on 64-bit systems.

NOTE **PARTITION STYLES**

Partition styles refer to the most elemental disk structure visible to the operating system. Partition styles do not affect file formats within partitions, such as NTFS or FAT32. Basic and dynamic disks can occur on either partition style.

Dynamic disks provide advanced features that basic disks do not: features such as the ability to create an unlimited number of volumes, volumes that span multiple disks (spanned and striped volumes), and fault-tolerant volumes (mirrored and RAID-5 volumes). There are five types of dynamic volumes: simple, spanned, striped, mirrored, and RAID-5.

In previous versions of Windows, you needed to convert a basic disk to a dynamic disk before you could create any of these volume types. When you use Disk Management in Windows Server 2008 R2 to create any of these volume types, however, basic disks are automatically converted to dynamic disks during the process. As a result, the question of whether a disk is basic or dynamic has become less important from an administrative point of view. Despite this development, it is still important to know for dual-boot configurations that many earlier versions of Windows (such as Windows NT, Windows 98, and Windows ME) cannot access dynamic disks. Also relevant for dual-boot configurations is the fact that dynamic disks are only compatible with Windows operating systems.

EXAM TIP

Even though basic disks are automatically converted to dynamic disks when necessary, you still need to know which volume types require dynamic disks for the 70-643 exam.

Creating Volumes

You can use Disk Management or the Diskpart command-line utility to create the following volume types in Windows Server 2008 R2.

- **Simple or basic volumes** Simple volumes are basic drives that are not fault tolerant. A simple volume can consist of a single region on a disk or multiple regions that are on the same disk and linked together.

  To create a simple volume in Disk Management, right-click unallocated space on a disk and then click New Simple Volume, as shown in Figure 2-6. (This process is identical whether you are creating the volume on a basic or dynamic disk, even though on a basic disk, the new volume is technically called a partition or basic volume.) Note that first you might need to right-click the disk and select Online.

  To create a simple volume by using the Diskpart utility, use the utility to select the disk and then, on a dynamic disk, type the **create volume simple** command. To create a new volume (partition) on a basic disk, type **create partition**. You can use **create volume ?** or **create partition ?** to learn the specific syntax associated with these commands.

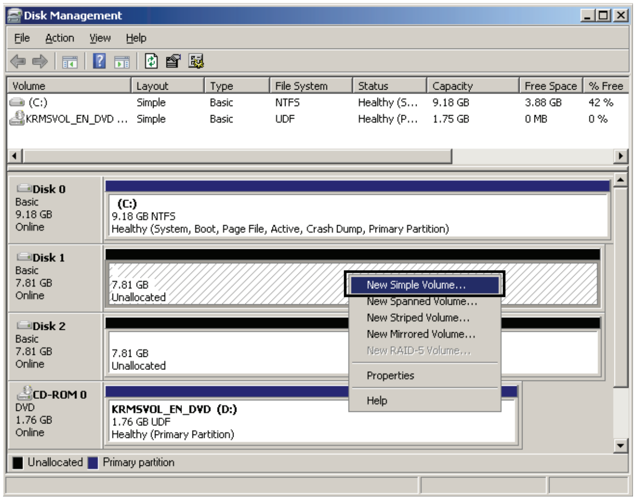

Volume	Layout	Type	File System	Status	Capacity	Free Space	% Free
(C:)	Simple	Basic	NTFS	Healthy (S...	9.18 GB	3.88 GB	42 %
KRMSVOL_EN_DVD ...	Simple	Basic	UDF	Healthy (P...	1.75 GB	0 MB	0 %

Disk 0
Basic
9.18 GB
Online

(C:)
9.18 GB NTFS
Healthy (System, Boot, Page File, Active, Crash Dump, Primary Partition)

Disk 1
Basic
7.81 GB
Online

7.81 GB
Unallocated

Disk 2
Basic
7.81 GB
Online

7.81 GB
Unallocated

CD-ROM 0
DVD
1.76 GB
Online

KRMSVOL_EN_DVD (D:)
1.76 GB UDF
Healthy (Primary Partition)

FIGURE 2-6 Creating a simple volume.

- **Spanned volumes** A spanned volume is a dynamic volume consisting of disk space on more than one disk. If a simple volume is not a system volume or boot volume, you can extend it across additional disks to create a spanned volume, or you can create a new volume as a spanned volume by using unallocated space on more than one disk.

  To create a new spanned volume, in Disk Management, right-click unallocated space on one of the disks where you want to create the spanned volume and then click New Spanned Volume. This step opens the New Spanned Volume Wizard, which allows you you to add space to the spanned volume from the disks available.

  Figure 2-7 shows a spanned volume, assigned drive letter E. Notice how the drive uses space from Disk 1 and Disk 2, but appears as only a single volume with a capacity of 7.32 GB.

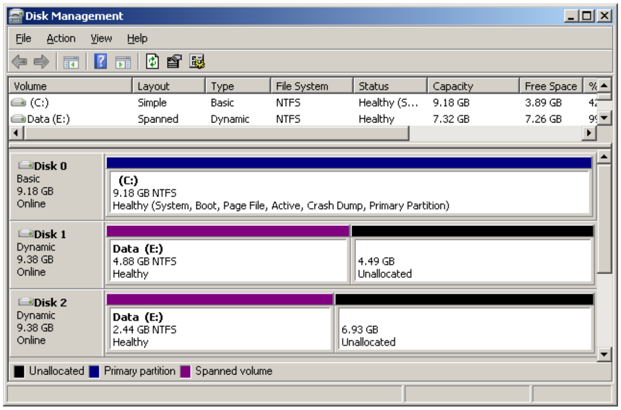

Volume	Layout	Type	File System	Status	Capacity	Free Space	%
(C:)	Simple	Basic	NTFS	Healthy (S...	9.18 GB	3.89 GB	4
Data (E:)	Spanned	Dynamic	NTFS	Healthy	7.32 GB	7.26 GB	9

Disk 0
Basic
9.18 GB
Online

(C:)
9.18 GB NTFS
Healthy (System, Boot, Page File, Active, Crash Dump, Primary Partition)

Disk 1
Dynamic
9.38 GB
Online

Data (E:)
4.88 GB NTFS
Healthy

4.49 GB
Unallocated

Disk 2
Dynamic
9.38 GB
Online

Data (E:)
2.44 GB NTFS
Healthy

6.93 GB
Unallocated

FIGURE 2-7 A spanned volume in Disk Management.

- **Striped volumes** A striped volume, which is also known as RAID 0, is a dynamic volume that stores data in stripes across two or more physical disks. Striped volumes offer the best performance of all the volumes available in Windows, but they do not provide fault tolerance. If a disk in a striped volume fails, the data in the entire volume is lost.

  Figure 2-8 shows how data in a striped volume is written across a set of disks.

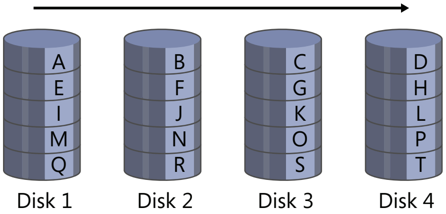

FIGURE 2-8 A RAID 0, or striped volume, stripes data across disks.

When should you use a striped volume? A striped volume is the best storage solution for temporary data that does not need fault tolerance but does require high performance. Examples of such temporary data include page files and Temp folders. To create a new striped volume in Disk Management, right-click unallocated space on a disk and then click New Striped Volume.

A striped volume in Disk Management is shown in Figure 2-9. The volume uses 1.46 GB of space from both Disk 1 and Disk 2 and appears as a single volume E with a total capacity of 2.93 GB. Note how the volume is being used to store temporary data (the Page File).

> ***NOTE*** **RAID DISKS**
>
> **As with all RAID solutions, a striped volume is built with disks of equal size.**

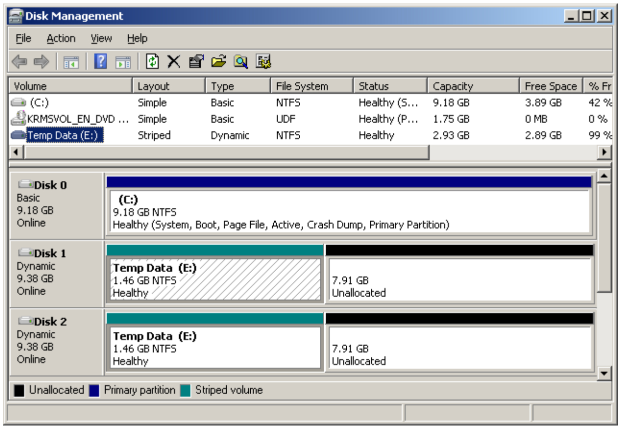

Volume	Layout	Type	File System	Status	Capacity	Free Space	% Fr
(C:)	Simple	Basic	NTFS	Healthy (S...	9.18 GB	3.89 GB	42 %
KRMSVOL_EN_DVD ...	Simple	Basic	UDF	Healthy (P...	1.75 GB	0 MB	0 %
Temp Data (E:)	Striped	Dynamic	NTFS	Healthy	2.93 GB	2.89 GB	99 %

Disk 0
Basic
9.18 GB
Online

(C:)
9.18 GB NTFS
Healthy (System, Boot, Page File, Active, Crash Dump, Primary Partition)

Disk 1
Dynamic
9.38 GB
Online

Temp Data (E:)
1.46 GB NTFS
Healthy

7.91 GB
Unallocated

Disk 2
Dynamic
9.38 GB
Online

Temp Data (E:)
1.46 GB NTFS
Healthy

7.91 GB
Unallocated

■ Unallocated ■ Primary partition ■ Striped volume

FIGURE 2-9 A RAID 0, or striped volume, in Disk Management.

- **Mirrored volumes** Also known as a RAID 1, a mirrored volume is a fault-tolerant volume that provides data redundancy by using two copies, or mirrors, of the same volume. All data written to the mirrored volume is written to both volumes, which are located on separate physical disks. If one of the physical disks fails, the data on the failed disk becomes unavailable, but the system continues to operate by using the unaffected disk.

  Figure 2-10 illustrates how data is stored on a mirrored volume. Because data is duplicated, no data is lost if either disk fails.

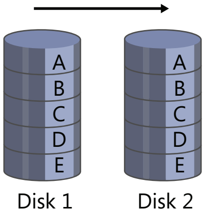

FIGURE 2-10 A RAID 1, or mirrored volume, copies all data onto a second disk.

NOTE **TRIPLE MIRRORING AND BEYOND**

Although mirrored volumes configured in Windows Server 2008 and Windows Server 2008 R2 are limited to two disks, mirrors created through third-party solutions can be created out of three disks or more. In a triple mirror configuration, for example, the contents of one disk are duplicated on two additional disks. Multiple mirrors degrade write performance but improve fault tolerance. They are good solutions for mission-critical data.

As a fault-tolerant solution, a mirrored volume has advantages and disadvantages. One advantage of a mirrored volume is that it offers very good read performance as well as fairly good write performance. In addition, mirroring requires only two disks, and almost any volume can be mirrored, including the system and boot volumes. The disadvantage of a mirrored volume is that it requires 50 percent of a disk's total storage capacity to be reserved for fault tolerance. Overall, if you need a fault-tolerant storage solution, a mirror is a good choice if you have only two disks; if you need good read and write performance; or if you need to provide fault tolerance for the system volume, the boot volume, or other mission-critical data.

To create a mirrored volume, you can either add a mirror to an existing volume or create a new mirrored volume. To add a mirror to an existing volume in Disk Management, right-click the existing volume and then click Add Mirror, as shown in Figure 2-11.

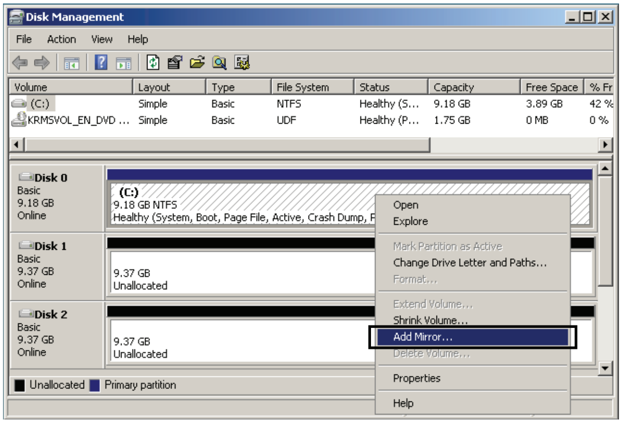

Volume	Layout	Type	File System	Status	Capacity	Free Space	% Fr
(C:)	Simple	Basic	NTFS	Healthy (S…	9.18 GB	3.89 GB	42 %
KRMSVOL_EN_DVD …	Simple	Basic	UDF	Healthy (P…	1.75 GB	0 MB	0 %

Disk 0
Basic
9.18 GB
Online
(C:)
9.18 GB NTFS
Healthy (System, Boot, Page File, Active, Crash Dump, F

Disk 1
Basic
9.37 GB
Online
9.37 GB
Unallocated

Disk 2
Basic
9.37 GB
Online
9.37 GB
Unallocated

■ Unallocated ■ Primary partition

FIGURE 2-11 Adding a mirror to the System partition.

To create a new mirrored volume in Disk Management, right-click unallocated space on a disk and then click New Mirrored Volume. A new mirrored volume is shown in Figure 2-12. The drive uses 5.86 GB of space from both Disk 1 and Disk 2 and appears as a single volume E with a total capacity of 5.86 GB.

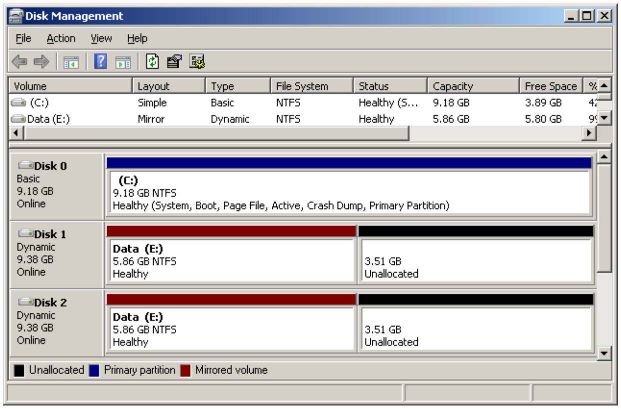

Volume	Layout	Type	File System	Status	Capacity	Free Space	%
(C:)	Simple	Basic	NTFS	Healthy (S…	9.18 GB	3.89 GB	4:
Data (E:)	Mirror	Dynamic	NTFS	Healthy	5.86 GB	5.80 GB	9!

Disk 0
Basic
9.18 GB
Online
(C:)
9.18 GB NTFS
Healthy (System, Boot, Page File, Active, Crash Dump, Primary Partition)

Disk 1
Dynamic
9.38 GB
Online
Data (E:)
5.86 GB NTFS
Healthy
3.51 GB
Unallocated

Disk 2
Dynamic
9.38 GB
Online
Data (E:)
5.86 GB NTFS
Healthy
3.51 GB
Unallocated

■ Unallocated ■ Primary partition ■ Mirrored volume

FIGURE 2-12 A RAID 1, or mirrored, volume.

- **Raid-5 volumes** A RAID-5 volume is a fault-tolerant volume that combines areas of free space from at least three physical hard disks into one logical volume. RAID-5 volumes stripe data along with *parity* (evenness or oddness) *information* across a set of disks. When a single disk fails, Windows uses this parity information to re-create the data on the failed disk. RAID-5 volumes can accept the loss of only a single disk in the set.

EXAM TIP

On the 70-643 exam, you might see a RAID-5 volume referred to as a striped volume with parity.

Figure 2-13 shows a RAID-5 volume made up of four disks. Data written to the volume is striped across these disks from left to right. For each stripe across the set of disks, one disk is used to hold parity information about the evenness or oddness of the other data in the stripe. In the simplified example shown in Figure 2-13, parity is set to 1 when the sum of the values in the stripe is odd, and parity is set to 0 when the sum of the remaining values is even. By using this parity information along with other disk data, if any one (and only one) disk fails, Windows can reconstruct the complete contents of that failed disk. The data of the failed drive can be re-created in real time as users request it. The party information can also be re-created live on a new disk after the failed disk has been replaced.

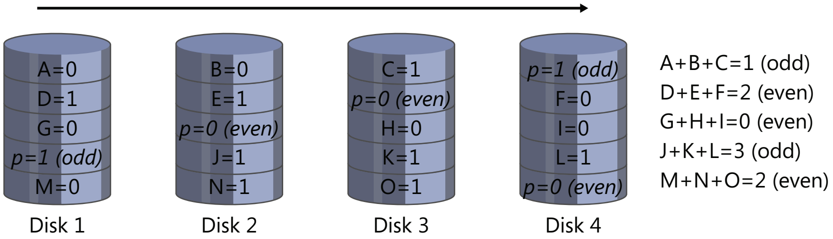

FIGURE 2-13 A RAID-5 volume calculates parity (evenness or oddness) for fault tolerance.

Space approximately equivalent to one disk is always used for fault tolerance in a RAID-5 volume. For example, if you create a RAID-5 out of four 120-GB disks, the total storage space available in that RAID-5 is 360 GB.

When should you use a RAID-5 volume? A RAID-5 volume is characterized by very good read performance, relatively poor write performance, and optimal use of storage space in a fault-tolerant solution. Therefore, consider using a RAID-5 volume when good write performance is not a priority, or when you need a fault-tolerant storage solution that makes the best use of available storage. Note also that you cannot assign the system or boot partition to a RAID-5 volume created in Windows Server 2008 or Windows Server 2008 R2.

NOTE **SOFTWARE AND HARDWARE RAIDS**

A RAID-5 volume created in Disk Management is an example of a software RAID because the RAID is created by the operating system. Some vendors, however, sell disk enclosures that include their own built-in RAID setup utility. If you configure a RAID-5 with this vendor software, the storage appears to Windows as a single local volume. A RAID configuration such as this, which is transparent to the operating system, is known

as a hardware RAID. Although software RAID has lower performance than hardware RAID does, software RAID is inexpensive and easy to configure because it has no special hardware requirements other than multiple disks. If cost is more important than performance, software RAID is an appropriate solution.

To create a RAID-5 volume in Disk Management, right-click unallocated space on one of the dynamic disks on which you want to create the RAID-5 volume and then click New RAID-5 Volume and follow the instructions in the New RAID-5 Volume Wizard.

To create a RAID-5 volume by using the Diskpart utility, use the *create volume raid* command. You can use the *help create volume raid* command to learn the exact syntax.

EXAM TIP

For the 70-643 exam, make sure you understand RAID levels and the different volume types.

Extending a Volume

You can add more space to existing simple or spanned volumes by extending them into unallocated space on the same disk or on a different disk. To extend a volume, it must either be formatted with the NTFS file system or unformatted. To extend a volume in Disk Management, right-click the simple or spanned volume you want to extend and then click Extend Volume.

NOTE **EXTENDING BOOT AND SYSTEM VOLUMES**

You cannot extend a boot or system volume onto another disk.

Shrinking a Volume

You can decrease the space used by simple or spanned volumes by shrinking them into contiguous free space at the end of the volume. For example, if you need to increase the amount of unallocated space on a disk to make room for a new partition or volume, you can attempt to shrink the existing volumes on the disk. When you shrink a partition, any ordinary files are automatically relocated on the disk to create the new unallocated space. There is no need to reformat the disk to shrink the partition.

The amount of space you can gain from shrinking a volume varies greatly. In general, the greater the percentage of unused space on the volume and the fewer the bad clusters, the more you can shrink the volume. If, however, the number of bad clusters detected by dynamic bad-cluster remapping is too great, you will not be able to shrink the volume at all. If this occurs, consider moving the data and replacing the disk.

CAUTION **DO NOT SHRINK RAW PARTITIONS THAT CONTAIN DATA**

If a partition is not formatted with a file system but still contains data (such as a database file), shrinking the partition can actually destroy the data.

To shrink a volume in Disk Management, right-click the simple or spanned volume that you want to shrink and then click Shrink Volume, as shown in Figure 2-14.

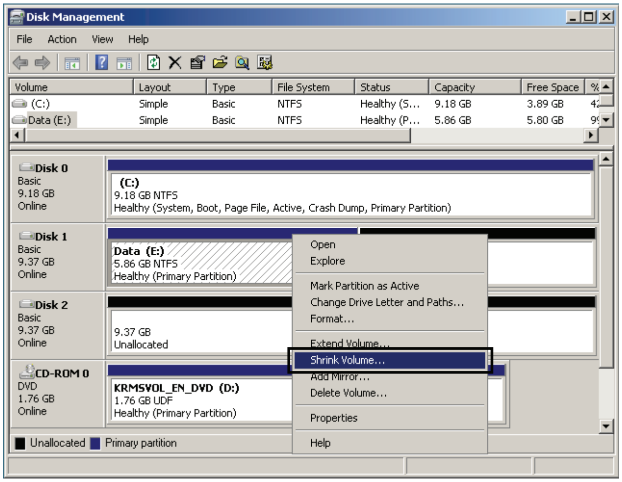

Volume	Layout	Type	File System	Status	Capacity	Free Space	%
(C:)	Simple	Basic	NTFS	Healthy (S...	9.18 GB	3.89 GB	42
Data (E:)	Simple	Basic	NTFS	Healthy (P...	5.86 GB	5.80 GB	9

Disk 0
Basic
9.18 GB
Online

(C:)
9.18 GB NTFS
Healthy (System, Boot, Page File, Active, Crash Dump, Primary Partition)

Disk 1
Basic
9.37 GB
Online

Data (E:)
5.86 GB NTFS
Healthy (Primary Partition)

Disk 2
Basic
9.37 GB
Online

9.37 GB
Unallocated

CD-ROM 0
DVD
1.76 GB
Online

KRMSVOL_EN_DVD (D:)
1.76 GB UDF
Healthy (Primary Partition)

FIGURE 2-14 Shrinking a volume in Disk Management.

EXAM TIP

Shrinking is a new feature in Windows Server 2008 and Windows Server 2008 R2. Expect to see a question on this topic on the 70-643 exam.

Configuring a Mount Point

A mount point is a folder in a volume that acts as a pointer to the root directory of another volume. For example, if you need to make more storage space available to the system or boot disk, you can create a new volume on another disk and then mount that volume in a folder in the system volume.

This arrangement is illustrated in Figure 2-15. In this scenario, the original disk capacity of the C drive is 9.18 GB. By mounting a 3.51-GB volume in a folder named MountedVolume in C, you are able to access more disk space through C even though you have not changed the capacity of the disk.

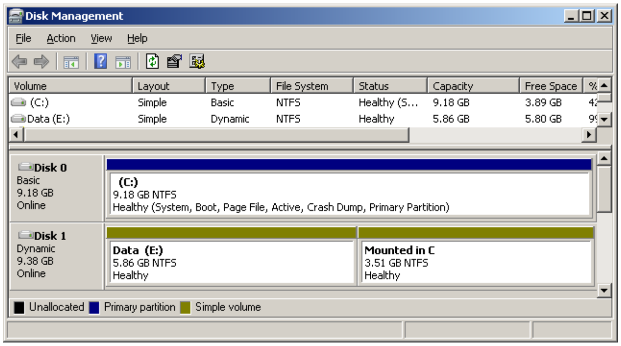

Disk Management

Volume	Layout	Type	File System	Status	Capacity	Free Space	%
(C:)	Simple	Basic	NTFS	Healthy (S...	9.18 GB	3.89 GB	4
Data (E:)	Simple	Dynamic	NTFS	Healthy	5.86 GB	5.80 GB	9

Disk 0
Basic
9.18 GB
Online

(C:)
9.18 GB NTFS
Healthy (System, Boot, Page File, Active, Crash Dump, Primary Partition)

Disk 1
Dynamic
9.38 GB
Online

Data (E:)
5.86 GB NTFS
Healthy

Mounted in C
3.51 GB NTFS
Healthy

■ Unallocated ■ Primary partition ■ Simple volume

FIGURE 2-15 A new volume mounted in the system volume.

You can create a mount point in Disk Management by creating a new volume and then choosing the option to mount the volume in an empty NTFS folder, as shown in Figure 2-16.

NOTE **EXTENDING THE SYSTEM OR BOOT VOLUME**

Because you cannot extend a system volume onto another disk, mount points are the only way you can make more space available to the system volume without replacing hardware.

You can also create a mount point for an existing volume by right-clicking the volume and then selecting Change Drive Letter And Paths. In the Change Drive Letter And Paths dialog box, click Change and then choose the option to mount the volume in an empty NTFS folder.

EXAM TIP

Understand mount points for the 70-643 exam.

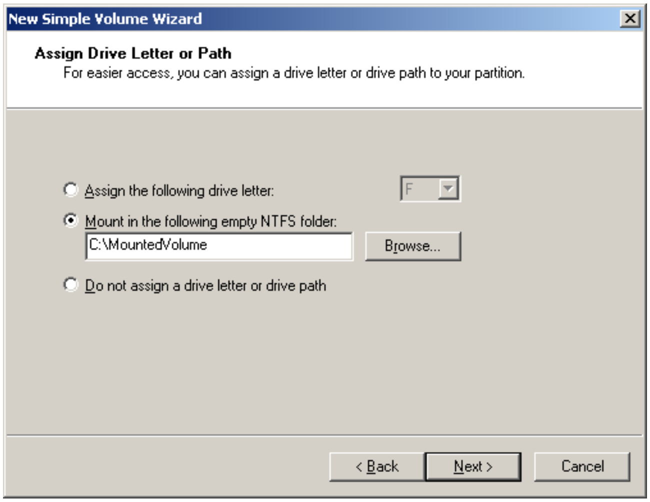

New Simple Volume Wizard

Assign Drive Letter or Path
For easier access, you can assign a drive letter or drive path to your partition.

○ Assign the following drive letter: F

● Mount in the following empty NTFS folder:
C:\MountedVolume

Browse...

○ Do not assign a drive letter or drive path

FIGURE 2-16 Mounting a new volume in an empty NTFS folder.

✓ **Quick Check**

1. Can you extend a mirrored volume?
2. True or False: A hardware RAID-5 volume cannot act as a system volume in Windows Server 2008 R2.

Quick Check Answers

1. No
2. False. Although a *software* RAID-5 cannot be used as a system volume in Windows Server 2008 R2, a *hardware* RAID-5 volume can indeed be used as the system volume because a hardware RAID would be transparent to the operating system. The limitation for RAID-5 volumes affects what you can configure from within the Windows operating system. You cannot add the system or boot partition to a software RAID-5 volume, and you cannot install an operating system on a RAID-5 volume that you create in Windows. (Note also that installing an operating system on a hardware RAID-5 volume, although possible, is discouraged because of the poor write performance associated with RAID-5.)

PRACTICE Work with Disk Sets

In this practice, you create various volume types in Disk Management.

NOTE **HOW MANY DISKS DO YOU NEED FOR THESE EXERCISES?**

These exercises require Server2 to have two unpartitioned disks (Disk 1 and Disk 2) of equal size. These disks must be separate from the system disk (Disk 0) on which you have installed Windows Server 2008 R2. If you have not done so already, you should now create these two new virtual disks by using Hyper-V Manager. (You can specify a size of 1 GB each.) You should then use the Settings For Server2.contoso.local dialog box in Hyper-V Manager to add one of these virtual disks to IDE Controller 0 and the other to IDE Controller 1.

EXERCISE 1 Working with Disks and Simple Volumes

In this exercise, which you perform on Server2, you create simple volumes on Disk 1 while switching first between dynamic and basic disks and then between MBR and GPT disks.

1. Log on to Contoso.com from Server2 as a domain administrator.
2. In the Run box, type **diskmgmt.msc** and then press Enter.
3. If the Initialize Disk dialog box appears, select MBR (Master Boot Record) and then click OK to initialize Disk 1 and Disk 2.

4. In Disk Management, in the top pane, ensure that C is the only lettered volume that is visible. If necessary, back up data and then delete all other volumes. (You may ignore any volume named "System Reserved.")

   In the bottom pane of Disk Management, three disks should be displayed: Disk 0, Disk 1, and Disk 2.

5. Right-click the unallocated space on Disk 1 and then click New Simple Volume.

   The New Simple Volume Wizard opens.

6. On the Welcome page of the New Simple Volume Wizard, click Next.
7. On the Specify Volume Size page, read all the text on the page and then click Next.
8. On the Assign Drive Letter Or Path page, read all the text on the page and then click Next.
9. On the Format Partition page, read all the text on the page, verify that Perform A Quick Format is selected, and then click Next.
10. On the Completing The New Simple Volume Wizard page, click Finish.

    After the creation and formatting are complete, the new volume appears in Disk Management.

11. In Disk Management, in the bottom pane, right-click the Disk 1 tile and then click Convert To Dynamic Disk.
12. In the Convert To Dynamic Disk dialog box, verify that Disk 1 is selected and then click OK.
13. In the Disks To Convert Dialog box, click Convert.
14. In the Disk Management dialog box, read all the text and then click Yes.

    After several moments, the new volume changes from blue to green.

15. Right-click the Disk 1 tile and then answer the following questions:

    Question: Can you convert Disk 1 back to a basic disk?

    Answer: No, because the option is grayed out.

    Question: Can you convert Disk 1 to the GPT partition style?

    Answer: No, because the disk contains volumes, so the option is unavailable.

16. Right-click the new volume you have created on Disk 1 and then click Delete Volume. Click Yes when prompted to confirm.
17. After the volume has deleted, answer the following question:

    Question: Is Disk 1 now listed as Basic or Dynamic?

    Answer: Basic. By default, disks with no volumes are basic.

18. Right-click Disk 1 and choose the option to convert Disk 1 to a dynamic disk. Then, after the conversion has completed, right-click Disk 1 to convert it back to a basic disk.

When a disk contains no volumes, you can convert it freely between basic and dynamic. However, when a disk contains volumes, you can convert only from basic to dynamic.

19. Right-click Disk 1 and then click Convert To GPT Disk.

20. After a few moments, right-click Disk 1 again and then click Convert To MBR Disk.

    When a disk contains no volumes, you can convert it freely between MBR and GPT partition styles. However, you cannot convert the partition style of a disk when it contains any volumes.

21. Leave Disk Management open and proceed to Exercise 2.

EXERCISE 2 Creating Mount Points

In this exercise, which you perform on Server2, you mount two volumes as folders in volume C.

1. While you are logged on to Contoso.com from Server2 as a domain administrator, in the root of volume C, create two new folders named **MountVol1** and **MountVol2**, respectively.

2. In Disk Management, right-click the unallocated space in Disk 1 and then click New Simple Volume.

3. On the Welcome To The New Simple Volume Wizard page, click Next.

4. On the Specify Volume Size page, in the Simple Volume Size In MB text box, type a value that represents approximately half of the available space. For example, if 1021 MB are available, type 500 and then click Next.

5. On the Assign Drive Letter Or Path page, select Mount In The Following Empty NTFS Folder. Type **C:\MountVol1** in the associated text box (or use the *Browse* button to select that folder) and then click Next.

6. On the Format Partition page, in the Volume Label text box, replace the “New Volume” text by typing **Mounted in C**.

7. On the Format Partition page, verify that Perform A Quick Format is selected and then click Next.

8. On the Completing the New Simple Volume Wizard page, click Finish.

   After a few moments, the new volume appears in Disk Management. It is not assigned a drive letter, but it is labeled Mounted In C.

9. In the Start menu, select Computer.

   In the Computer window, only the C drive is visible. You cannot directly access the new drive you have just created.

10. Open the C drive.

    In the C drive, the MountVol1 folder is marked by a special icon. It is also associated with a large size, even though the volume is empty.

11. Open the properties of MountVol1.

    In the MountVol1 Properties dialog box, the type is listed as Mounted Volume.

12. In the MountVol1 Properties dialog box, click Properties.

    The Mounted In C (C:\MountVol1) Properties dialog box opens. The dialog box displays the same information you would find in the properties sheet of a volume.

13. Click OK to close the Mounted In C (C:\MountVol1) Properties dialog box and then click OK to close the MountVol1 Properties dialog box.

14. In Disk Management, create a new simple volume on Disk 1 by using the same process described in Exercise 1. Use all the remaining space on Disk 1 for the new volume and do *not* select the option to mount the volume in an NTFS folder. Type **Mounted In C (2)** as the name for the volume and choose the option to perform a quick format.

    After the new volume is created, it appears in Disk Management, assigned a drive letter such as E.

15. In Disk Management, right-click the Mounted In C (2) volume and then click Change Drive Letter And Paths.

16. In the Change Drive Letter And Paths dialog box, click Remove and then click Yes to confirm.

    You can mount an existing volume only if you first remove any drive letter associated with it.

17. In Disk Management, right-click the Mounted In C (2) volume again and, again, click Change Drive Letter And Paths.

18. In the Change Drive Letter And Paths dialog box, click Add.

19. In the Add Drive Letter Or Path dialog box, click Mount In The Following Empty NTFS Folder and then type or browse to C:\MountVol2.

20. In the Add Drive Letter or Path dialog box, click OK.

21. Click Start and then Computer to verify that Mounted In C (2) has been configured as a mount point in the folder named MountVol2 in the C drive.

22. In Disk Management, delete both the Mounted In C and the Mounted In C (2) volumes. Verify that only unallocated space remains on Disk 1.

23. Close all windows except for Disk Management and then proceed to Exercise 3.

EXERCISE 3 Add and Break a Mirror

In this exercise, which you perform on Server2, you create a new volume on Disk 1 and then add a mirror on Disk 2.

1. While you are still logged on to Contoso.com from Server2 as a domain administrator, in Disk Management, create a new simple volume on Disk 1, using all the available space on the disk. Complete the New Simple Volume Wizard by using all the default options.

2. Right-click the new E volume on Disk 1 and then click Add Mirror.

3. In the Add Mirror dialog box, select Disk 2 and then click Add Mirror.
4. In the Disk Management dialog box, read all the text and then click Yes.

   A new volume is created on Disk 2, and then, after both Disk 1 and Disk 2 are converted to dynamic disks, the new volume on Disk 1 is also assigned the drive letter C. The status of the twin volumes is then shown as Resynching while the mirror is created. This process of resynchronization varies, depending on the size of the volumes.
5. After the mirror volume has finished resynchronizing, take a few moments to browse Disk Management, noting the single volume listed in the top pane and the capacity of the drive.
6. On Disk 2, right-click volume E. Use the options available on the shortcut menu to answer the following questions:

   Question: Which option on the shortcut menu should you choose if you want to turn the mirrored volume into two separate volumes?

   Answer: Break Mirrored Volume. You should choose this option when one of the disks fails or becomes corrupted.

   Question: Which option on the shortcut menu should you choose if you want to delete the mirror on Disk 2 immediately?

   Answer: Remove Mirror.
7. On the shortcut menu, click Remove Mirror.
8. In the Remove Mirror dialog box, select Disk 2 and then click Remove Mirror.
9. In the Data Management dialog box, click Yes to confirm.

   In Disk Management, Disk 2 once again appears as a basic disk with only unallocated space.
10. Delete Volume E on Disk 1.
11. Leave Disk Management open and proceed to Exercise 4.

EXERCISE 4 Creating a Spanned Volume

In this exercise, you create a spanned volume on Disk 1 and Disk 2. You need two unpartitioned dynamic disks for this exercise.

1. In Disk Management, right-click the unallocated space in Disk 1 and then click Create New Spanned Volume.

   The New Spanned Volume Wizard opens.
2. On the Welcome To The New Spanned Volume Wizard page, read all the text on the page and then click Next.
3. On the Select Disks page, verify that only Disk 1 is visible in the Selected area.
4. In the Select The Amount Of Space In MB text box, type an amount that is equal to approximately half of the available space. For example, if the default number in the box is 1021, replace that amount by typing **500**.

5. On the Select Disks page, select Disk 2, which is shown in the Available area, and then click Add to move Disk 2 to the Selected area.
6. In the Selected area, click to select Disk 2.
7. In the Select The Amount Of Space In MB text box, type an amount that is equal to approximately 25 percent of the available space. For example, if the default number in the box is 1021, replace that amount by typing **250**.
8. On the Select Disks page, click Next.
9. On the Assign Drive Letter Or Path page, click Next.
10. On the Format Volume page, in the Volume Label text box, replace the text by typing **Spanned Volume**.
11. On the Format Volume page, verify that the Perform A Quick Format check box is selected and then click Next.
12. On the Completing The New Spanned Volume Wizard page, click Finish.
13. If the Disk Management dialog box appears, read all the text and then click Yes.

    After the creation and formatting complete, the new spanned volume appears in Disk Management. The new volume spans disks 1 and 2.
14. Spend a few moments browsing the information related to the new volume in Disk Management. Note, for example, the capacity of the volume and the fact that it is assigned a single drive letter.
15. Leave Disk Management open and proceed to Exercise 5.

EXERCISE 5 Creating a Striped Volume

In this exercise, you create a new striped volume in the remaining space on Disk 1 and Disk 2.

1. While you are logged on to Contoso.com from Server2 as a domain administrator, in Disk Management, right-click the unallocated space in Disk 1 and then click New Striped Volume.

   The New Striped Volume Wizard appears.
2. On the Welcome page of the New Striped Volume Wizard, click Next.
3. On the Select Disks page, note that only Disk 1 appears in the Selected area.
4. On the Select Disks page, select Disk 2 in the Available area and then click Add to move Disk 2 to the Selected area.

   The amount of space associated with Disk 1 and Disk 2 is identical. In a striped volume, all member disks must be the same size.
5. On the Select Disks page, click Next.
6. On the Assign Drive Letter Or Path page, click Next.
7. On the Format Volume page, in the Volume Label text box, replace the text by typing **Striped Volume**.

8. On the Format Volume page, verify that the Perform A Quick Format check box is selected and then click Next.
9. On the Completing The New Striped Volume Wizard page, click Finish.

   After the creation and formatting complete, the new striped volume appears in Disk Management.
10. Spend a few moments browsing the information related to the new striped volume in Disk Management. Note, for example, the capacity of the volume and the fact that it is assigned a single drive letter.
11. Leave Disk Management open and proceed to Exercise 6.

EXERCISE 6 Shrinking and Extending a Volume

In this exercise, you shrink the spanned volume you created in Exercise 5. Then, after deleting the striped volume you created in the same exercise, you extend the spanned volume into the available space on Disk 1.

1. While you are still logged on to Contoso.com from Server2 as a domain administrator, in Disk Management, right-click the Spanned Volume on Disk 2 and then click Shrink Volume.

   The Querying Shrink Space box appears, and then the Shrink [Drive Letter] dialog box appears.
2. In the Shrink dialog box, read all the text.

   In the Enter The Amount Of Space To Shrink In MB text box, the default amount provided is equal to the maximum amount you can shrink the drive.
3. Click Shrink to shrink the volume the maximum allowable amount.

   After several moments, the spanned volume appears in its newer, smaller size. Now, the volume might or might not be limited to Disk 1.
4. In Disk Management, right-click the striped volume (not the spanned volume) and then click Delete Volume.
5. In the Delete Striped Volume dialog box, read all the text and then click Yes to confirm.

   After the volume is deleted, new unallocated space appears on Disk 1.
6. Right-click the spanned volume on Disk 1 and then click Extend Volume.

   The Extend Volume Wizard opens.
7. On the Welcome To The Extend Volume Wizard page, read all the text and then click Next.
8. On the Select Disk page, verify that only Disk 1 is shown in the Selected area.
9. On the Select Disk page, leave the default (full) amount of space to expand on Disk 1 and then click Next.
10. On the Completing the Extend Volume Wizard page, click Finish.

After a few moments, the volume appears in Disk Management, occupying all the space on Disk 1. If the volume is confined to Disk 1, it is now designated as a simple volume. If some portion remains on Disk 2, it is still designated as a spanned volume.

11. Right-click the volume on Disk 1, click Delete Volume, and then click Yes to confirm the deletion.

    After a few moments, Disk Management shows that Disk 1 and Disk 2 have returned to their original state.

12. Log off Server2.

Lesson Summary

- In general, disk storage occurs in three varieties: direct-attached storage (DAS), network-attached storage (NAS), and storage-area networks (SANs). Both DAS and SANs provide block-based access to data storage, and NAS provides file-based access. SANs provide the additional benefit of shared storage that you can easily move from server to server.
- When vendor disk storage subsystems include a hardware provider for Virtual Disk Service (VDS), you can manage that hardware within Windows Server 2008 R2 by using tools such as Disk Management, Storage Manager for SANs (SMfS), Storage Explorer, iSCSI Initiator, or the DiskRAID.exe command-line tool.
- Disk Management is the main tool you can use for managing disks and volumes in Windows Server 2008 R2. Disk Management enables you to create simple, spanned, striped, mirrored, and RAID-5 volumes.
- By using Disk Management, you can extend or shrink a simple or spanned volume.
- By using Disk Management, you can configure a volume as a mount point in another volume.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. The questions are also available on the companion CD if you prefer to review them in electronic form.

NOTE **ANSWERS**

Answers to these questions and explanations of why each answer choice is correct or incorrect are located in the "Answers" section at the end of the book.

1. You work as a network administrator, and your responsibilities include managing server storage. You have been asked to purchase a new disk subsystem for your company's storage-area network (SAN). You are in the process of testing hardware solutions before making purchases, and you attach a new disk subsystem to the network. You

want to provision the new disks and create new logical unit numbers (LUNs) to assign to a server named Server1. You open Storage Manager for SANs, but you can't see the new hardware. However, you can connect to the new hardware by using the software provided by the vendor. You want to be able to manage the new disk subsystem you purchase by using Storage Manager for SANs. What should you do?

- **A.** In Disk Management, choose the Rescan Disks option.
- **B.** Choose a disk subsystem from a vendor that has a Virtual Disk Service hardware provider.
- **C.** On Server1, configure iSCSI Initiator to specify the new hardware as a favorite target.
- **D.** Use Storage Explorer to configure Server1 as an iSNS server.

2. You work as an IT support specialist. Your job responsibilities include managing server storage. You are designing storage for a new application server. The application makes heavy use of temporary storage, and you want to allocate three 20-GB disk drives to that storage. If excellent read and write performance is a high priority, and you also want to use as much available space as possible, which of the following volume types should you create?

- **A.** Simple volume
- **B.** Spanned volume
- **C.** Mirrored volume
- **D.** Striped volume
- **E.** RAID-5 volume

3. Each client desktop in your company is configured with a K drive mapped to a network share on a file server running Windows Server 2008 R2. The shared folder on the file server is named Data, and it contains many personal folders and shared project data. The Data folder is stored on a nonsystem disk assigned the volume letter E, which is running low on storage space.

To address the problem of low storage space available through the K drive, you add another physical disk to the file server and create a new volume on it. Your goals are to maximize the space available through K and cause minimum disruption to users and their procedures for storing and sharing data. What should you do?

- **A.** Mount the new volume as a folder within the Data folder.
- **B.** Copy the Data folder to the new volume and re-create the Data share to point to the new location.
- **C.** Extend the E drive onto the new volume.
- **D.** Create a mirror consisting of the new volume and the E drive.

Lesson 2: Configuring Server Clusters

In enterprise networks, groups of independent servers are often used to provide a common set of services. Different physical computers, for example, can answer requests directed at a common website or database server. Although these server groups are often referred to generally as *clusters*, cluster types can serve very different purposes. This lesson describes the load balancing and high-availability server clusters you can configure in Windows Server 2008 R2.

After this lesson, you will be able to:

- Understand the features and limitations of DNS round-robin.
- Understand the main function and features of Network Load Balancing clusters.
- Know the basic steps to configure a Network Load Balancing cluster.
- Understand the main function and features of failover clusters.
- Understand the requirements for creating a failover cluster.

Estimated lesson time: 90 minutes

Server Cluster Fundamentals

In Windows Server 2008 R2, you can configure three types of server groups for load balancing, scalability, and high availability. First, a *round-robin distribution group* is a set of computers that uses DNS to provide basic load balancing with minimal configuration requirements. Next, a *Network Load Balancing (NLB) cluster* (also called an *NLB farm*) is a group of servers that not only provide load balancing but also increase scalability. Finally, a *failover cluster* can increase the availability of an application or service in the event of a server failure.

NOTE **WHAT IS LOAD BALANCING?**

Load balancing is a means of distributing incoming connection requests to two or more servers in a manner that is transparent to users. Load balancing can be implemented with hardware, software, or a combination of both.

Round-Robin Distribution

Round-robin DNS is a simple method for distributing a workload among multiple servers. In round-robin, a DNS server is configured with more than one record to resolve another server's name to an IP address. When clients query the DNS server to resolve the name (find the address) of the other server, the DNS server responds by cycling through the records one at a time and by pointing each successive client to a different address and different machine.

For example, suppose that a DNS server that is authoritative for the contoso.com DNS domain is configured with two resource records, each resolving the name web.contoso.com by pointing to a different server, as shown in Figure 2-17. When the first client (Client1) queries the DNS server to resolve the web.contoso.com name, the DNS server answers by pointing the client to the server named websrv1 located at the 192.168.3.11 address. This is the information associated with the first DNS record matching "web." When the next client, Client2, queries the DNS server to resolve the same name (web.contoso.com), the DNS server answers the query with the information provided in the second record matching "web." This second record points to a server named websrv2, which is located at the 192.168.3.12 address. If a third client then queries the DNS server for the same name, the server will respond with information in the first record again.

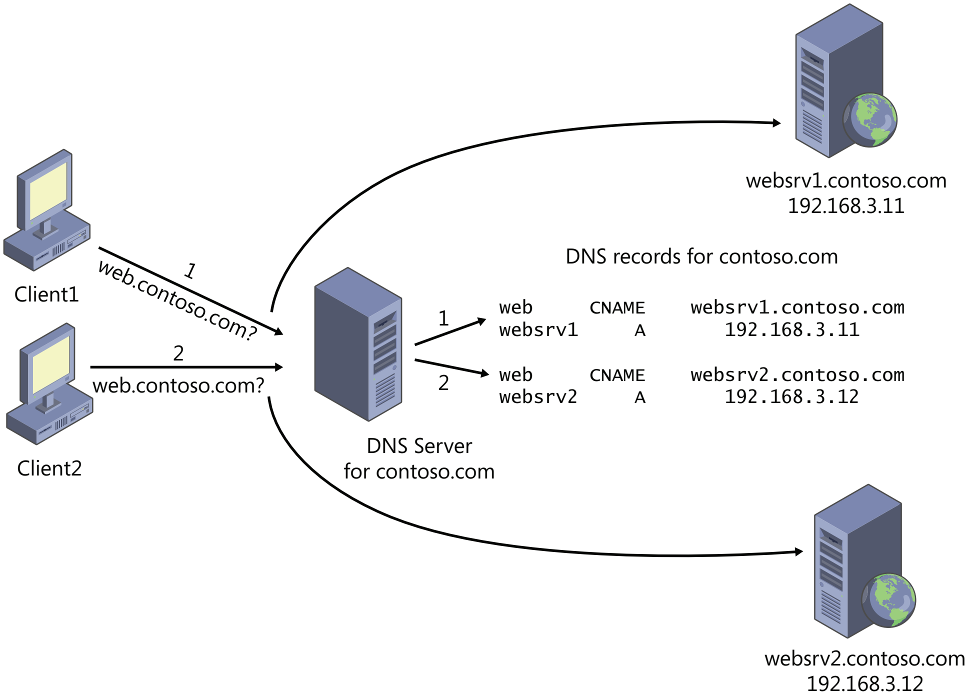

```
web       CNAME    websrv1.contoso.com
websrv1       A       192.168.3.11

web       CNAME    websrv2.contoso.com
websrv2       A       192.168.3.12
```

FIGURE 2-17 Round-robin uses DNS to distribute the client load between two or more servers.

The purpose of DNS round-robin is to load balance client requests among servers. Its main advantage is that it is very easy to configure. Round-robin DNS is enabled by default in most DNS servers, so to configure this simple sort of load balancing, you need only to create the appropriate DNS records on the DNS server.

However, there are serious limitations to round-robin as a load-balancing mechanism. The biggest drawback is that if one of the target servers goes down, the DNS server does not respond to this event, and it will keep directing clients to the inactive server until a network administrator removes the DNS record from the DNS server. Another drawback is that every record is given equal weight, regardless of whether one target server is more powerful than another or a given server is already busy. A final drawback is that round-robin does not always function as expected. Because DNS clients cache query responses from servers, a DNS client by default will keep connecting to the same target server as long as the cached response stays active.

Network Load Balancing

An installable feature of Windows Server 2008 R2, NLB transparently distributes client requests among servers in an NLB cluster by using virtual IP addresses and a shared name. From the perspective of the client, the NLB cluster appears to be a single server. NLB is a fully distributed solution in that it does not use a centralized dispatcher.

In a common scenario, NLB is used to create a *web farm*—a group of computers working to support a website or set of websites. However, NLB can also be used to create a terminal server farm, a VPN server farm, or a Forefront Threat Management Gateway (TMG) firewall cluster. Figure 2-18 shows a basic configuration of an NLB web farm located behind an NLB firewall cluster.

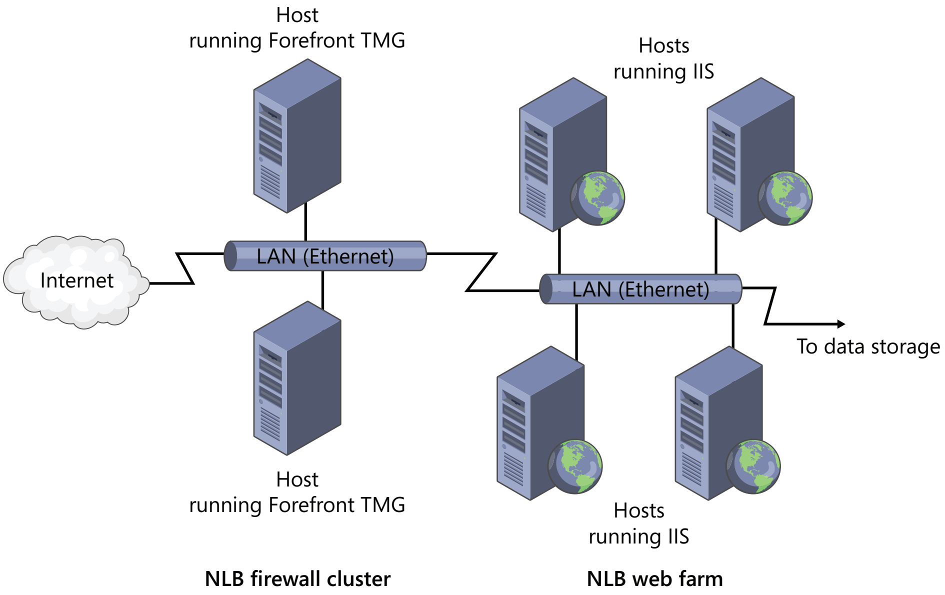

FIGURE 2-18 Basic diagram for two connected NLB clusters.

As a load-balancing mechanism, NLB provides significant advantages over round-robin DNS. First of all, in contrast to round-robin DNS, NLB automatically detects servers that have been disconnected from the NLB cluster and then redistributes client requests to the remaining live hosts. This feature prevents clients from sending requests to the failed servers. Another difference between NLB and round-robin DNS is that in NLB, you have the option to specify a load percentage that each host will handle. Clients are then distributed among hosts so that each server receives its percentage of incoming requests.

Beyond load balancing, NLB also supports scalability. As the demand for a network service such as a website grows, more servers can be added to the farm with only a minimal increase in administrative overhead.

Failover Clustering

A failover cluster is a group of two or more computers used to prevent downtime for selected applications and services. The clustered servers (called nodes) are connected by physical cables to each other and to shared disk storage. If one of the cluster nodes fails, another node begins to take over service for the lost node in a process known as failover. As a result of failover, users connecting to the server experience minimal disruption in service.

Servers in a failover cluster can function in a variety of roles, including the roles of file server, print server, mail server, and database server, and they can provide high availability for a variety of other services and applications.

In most cases, the failover cluster includes a shared storage unit that is physically connected to all the servers in the cluster, although any given volume in the storage is accessed by only one server at a time. In some cases, a separate shared volume called a *disk witness* is used (also called a *witness disk* or a *quorum disk*), which contains a copy of the cluster configuration database. A disk witness often improves the fault tolerance of the cluster.

Figure 2-19 illustrates the process of failover in a basic, two-node failover cluster.

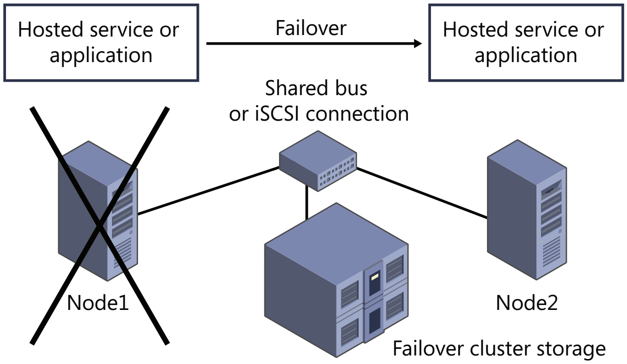
Hosted service or application
Failover
Hosted service or application
Shared bus or iSCSI connection
Node1
Node2
Failover cluster storage

FIGURE 2-19 In a failover cluster, when one server fails, another takes over and uses the same storage.

In a failover cluster, storage volumes or LUNs that are exposed to the nodes in a cluster must not be exposed to other servers, including servers in another cluster. Figure 2-20 illustrates this concept by showing two two-node failover clusters dividing up storage on a SAN.

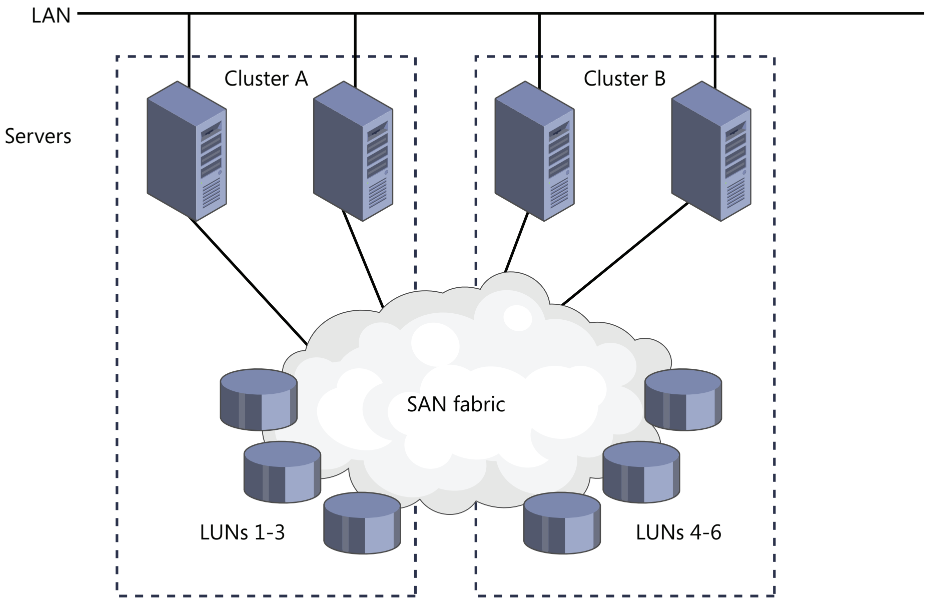

FIGURE 2-20 Each failover cluster must isolate storage from other servers.

Configuring an NLB Cluster

Creating an NLB cluster is a relatively simple process. To begin, install Windows Server 2008 R2 on two servers and then, on both servers, configure the service or application (such as IIS) that you want to provide to clients. Be sure to create identical configurations because you want the client experience to be identical, regardless of the server to which users are connected.

The next step in configuring an NLB cluster is to install the Network Load Balancing feature on all servers that you want to join the NLB cluster. For this step, simply open Server Manager and click Add Features. In the Add Features Wizard, select Network Load Balancing, click Next, and then follow the prompts to install.

The final step in creating an NLB cluster is to use Network Load Balancing Manager to configure the cluster. This procedure is outlined in the following section.

To create an NLB cluster, perform the following steps:

1. Launch Network Load Balancing Manager from Administrative Tools. (You can also open Network Load Balancing Manager by typing **Nlbmgr.exe** from a command prompt.)

   Network Load Balancing Manager is shown in Figure 2-21.

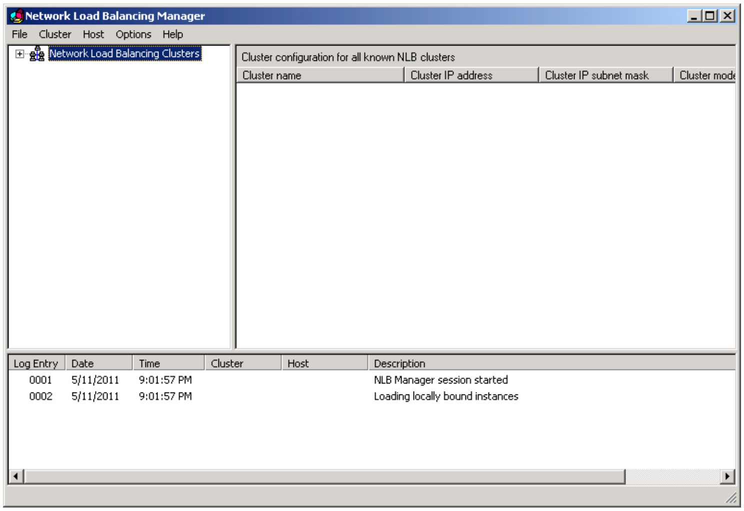
Network Load Balancing Manager
File
Cluster
Host
Options
Help
Network Load Balancing Clusters
Cluster configuration for all known NLB clusters
Cluster name
Cluster IP address
Cluster IP subnet mask
Cluster mode

Log Entry	Date	Time	Cluster	Host	Description
0001	5/11/2011	9:01:57 PM			NLB Manager session started
0002	5/11/2011	9:01:57 PM			Loading locally bound instances

FIGURE 2-21 Network Load Balancing Manager.

2. In the Network Load Balancing Manager console tree, right-click Network Load Balancing Clusters and then click New Cluster. This step opens the New Cluster: Connect page, which is shown in Figure 2-22.

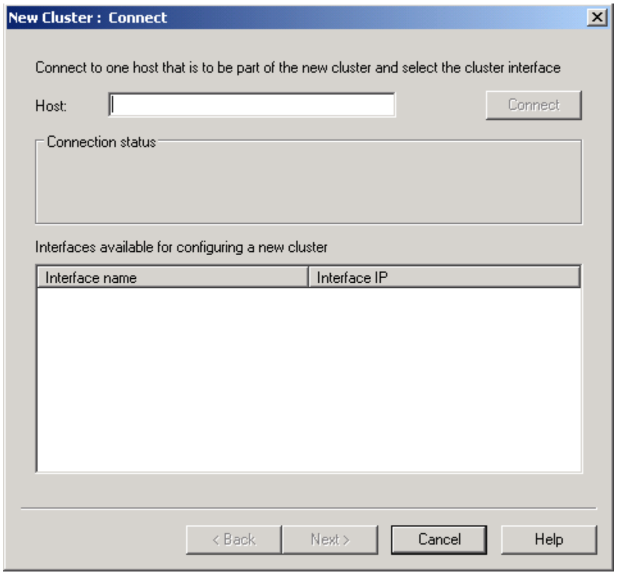

Connect to one host that is to be part of the new cluster and select the cluster interface

Interfaces available for configuring a new cluster

FIGURE 2-22 Connecting to the first host in a new NLB cluster.

3. Connect to the host that is to be part of the new cluster. In Host, enter the name of the host and then click Connect.

4. Select the interface you want to use with the cluster and then click Next. (The interface hosts the virtual IP address and receives the client traffic to load balance.)

5. On the New Cluster: Host Parameters page, shown in Figure 2-23, select a value in the Priority (Unique host identifier) drop-down list. This parameter specifies a unique ID for each host. The host with the lowest numerical priority among the current members of the cluster handles all the cluster's network traffic not covered by a port rule. You can override these priorities or provide load balancing for specific ranges of ports by specifying rules on the Port rules tab of the Network Load Balancing Properties dialog box.

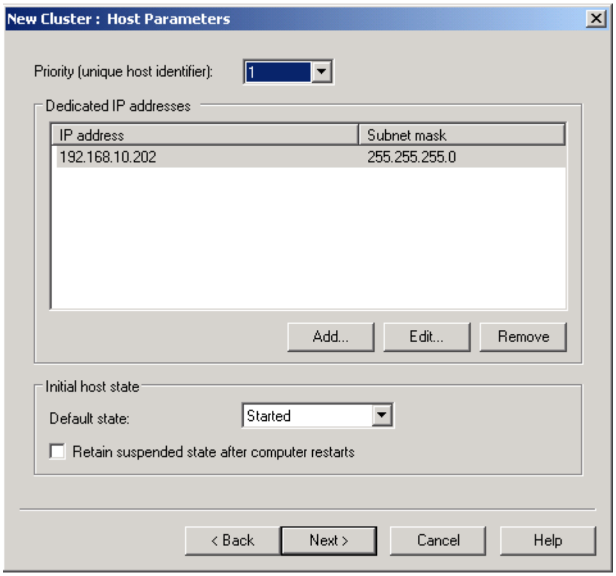
New Cluster : Host Parameters

Priority (unique host identifier): 1

Dedicated IP addresses

IP address	Subnet mask
192.168.10.202	255.255.255.0

Initial host state

Default state: Started

□ Retain suspended state after computer restarts

FIGURE 2-23 Setting a priority for an NLB cluster host.

6. On the New Cluster: Host Parameters page, verify that the dedicated IP address from the chosen interface is visible in the list. If not, click Add to add the address and then click Next to continue.

7. On the New Cluster: Cluster IP Addresses page, click Add to enter the cluster IP address shared by every host in the cluster. NLB adds this IP address to the TCP/IP stack on the selected interface of all hosts chosen to be part of the cluster. Click Next to continue.

NOTE **USE ONLY STATIC ADDRESSES**

NLB doesn't support Dynamic Host Configuration Protocol (DHCP). NLB disables DHCP on each interface it configures, so the IP addresses must be static.

8. On the New Cluster: Cluster Parameters page, shown in Figure 2-24, in the Cluster IP Configuration area, verify appropriate values for IP address and subnet mask and then type a full Internet name (Fully Qualified Domain Name) for the cluster.

   For IPv6 addresses, a subnet mask is not needed, and a full Internet name is not needed when using NLB with Terminal Services.

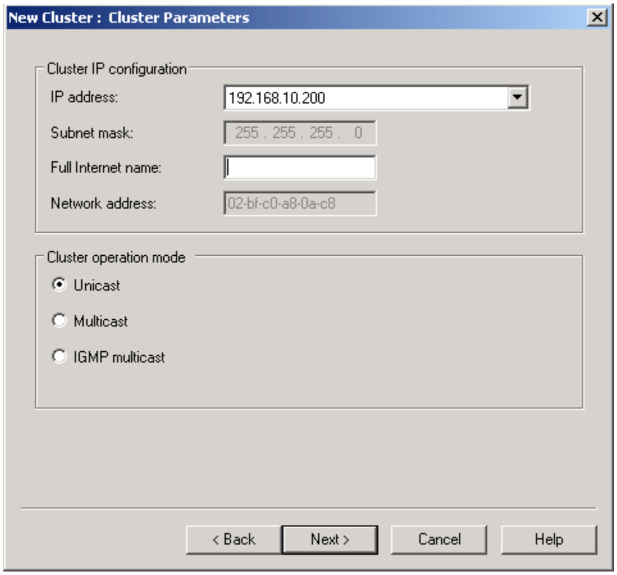

New Cluster : Cluster Parameters

Cluster IP configuration

IP address: 192.168.10.200

Subnet mask: 255 . 255 . 255 . 0

Full Internet name:

Network address: 02-bf-c0-a8-0a-c8

Cluster operation mode

● Unicast

○ Multicast

○ IGMP multicast

FIGURE 2-24 Adding a full Internet name for the NLB cluster.

9. On the New Cluster: Cluster Parameters page, in the Cluster Operation Mode area, click Unicast to specify that a unicast media access control (MAC) address should be used for cluster operations. In unicast mode, the MAC address of the cluster is assigned to the network adapter of the computer, and the built-in MAC address of the network adapter is not used. It is recommended that you accept the unicast default settings. Click Next to continue.

10. On the New Cluster: Port Rules page, shown in Figure 2-25, click Edit to open the Add/Edit Port Rule dialog box, shown in Figure 2-26. Port rules define which incoming TCP/IP requests are balanced among the hosts in the NLB cluster. Configure the NLB port rules as follows:

    - In the Port Range area, specify a range corresponding to the service you want to provide in the NLB cluster. For example, for web services, type **80** in the From and To fields so that the new rule applies only to HTTP traffic. For HTTPS traffic, type **443**. For Remote Desktop Services, type **3389**.
    - In the Protocols area, select TCP or UDP, as needed, to specify the transport-layer protocol the port rule should cover. Only the network traffic for the specified

protocol is affected by the rule. Traffic not affected by the port rule is handled by the default host.

- In the Filtering Mode area, select Multiple Host if you want multiple hosts in the cluster to handle network traffic for the port rule. Choose Single Host if you want a single host to handle the network traffic for the port rule. If you choose Single Host, the port rule will direct matching traffic to the active cluster host with the lowest handling priority.
- In Affinity (which applies only for the Multiple host-filtering mode), select None if you want multiple connections from the same client IP address to be handled by different cluster hosts (no client affinity). Leave the Single option if you want NLB to direct multiple requests from the same client IP address to the same cluster host. Select Network if you want NLB to direct multiple requests from the local subnet to the same cluster host.

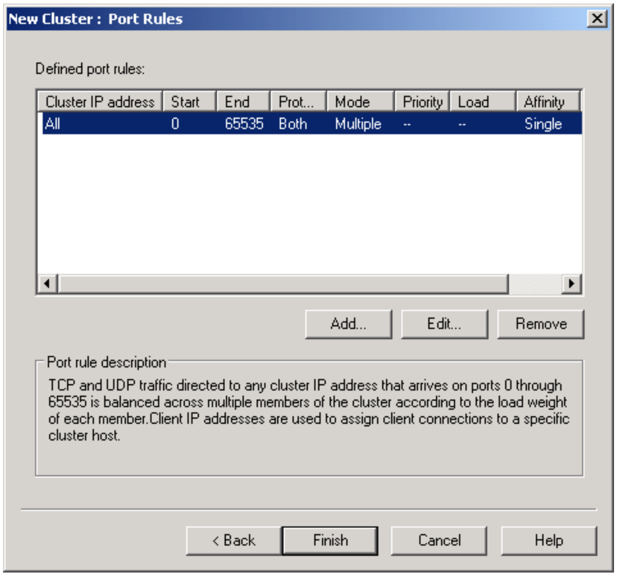

Defined port rules:

Cluster IP address	Start	End	Prot...	Mode	Priority	Load	Affinity
All	0	65535	Both	Multiple	--	--	Single

Port rule description

TCP and UDP traffic directed to any cluster IP address that arrives on ports 0 through 65535 is balanced across multiple members of the cluster according to the load weight of each member.Client IP addresses are used to assign client connections to a specific cluster host.

FIGURE 2-25 NLB port rules.

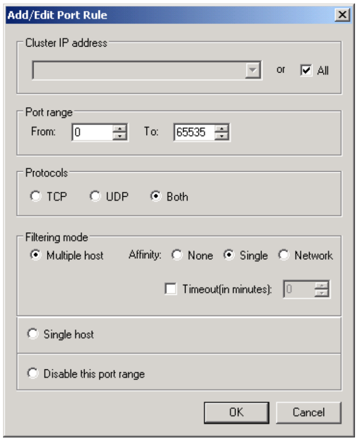

Add/Edit Port Rule

Cluster IP address

or ☑ All

Port range

From: 0 To: 65535

Protocols

○ TCP ○ UDP ● Both

Filtering mode

● Multiple host Affinity: ○ None ● Single ○ Network

□ Timeout(in minutes): 0

○ Single host

○ Disable this port range

FIGURE 2-26 Configuring NLB port rules.

EXAM TIP

Understand all the configuration options for NLB port rules for the 70-643 exam.

11. After you add the port rules, click Finish to create the cluster.

To add more hosts to the cluster, right-click the new cluster and then click Add Host To Cluster. Configure the host parameters (including host priority and dedicated IP addresses) for the additional hosts by following the same instructions you used to configure the initial host. Because you are adding hosts to an already-configured cluster, all the cluster-wide parameters remain the same.

Creating a Failover Cluster

Creating a failover cluster is a multistep process. The first step is to configure the physical hardware and shared storage for the cluster. Then, you must install the Failover Clustering feature and, optionally, run the Failover Cluster Validation tool, which ensures that the hardware and software prerequisites for the cluster are met. After the configuration has been validated by the tool, create the cluster by running the Create Cluster Wizard. Next, configure the cluster quorum settings by using the Configure Cluster Quorum Wizard. Finally, to configure the behavior of the cluster and to define the availability of selected services, install the desired application or service on cluster nodes and then run the High Availability Wizard.

Each of these steps is described in the following sections.

Preparing Failover Cluster Hardware

Failover clusters have fairly elaborate hardware requirements. To configure the hardware, review the following list of requirements for the servers, network adapters, cabling, controllers, and storage:

- **Servers** Use a set of matching computers that consist of the same or similar components (recommended).
- **Network adapters and cabling** The network hardware, like other components in the failover cluster solution, must be compatible with Windows Server 2008 R2. If you use iSCSI, each network adapter must be dedicated to either network communication or iSCSI, not both.

  In the network infrastructure that connects your cluster nodes, avoid having single points of failure. There are multiple ways of accomplishing this. You can connect your cluster nodes by multiple, distinct networks. Alternatively, you can connect your cluster nodes with one network constructed with teamed network adapters, redundant switches, redundant routers, or similar hardware that removes single points of failure.
- **Device controllers or appropriate adapters for the storage** If you are using serial attached SCSI or FC in all clustered servers, the mass-storage device controllers dedicated to the cluster storage should be identical. They should also use the same firmware version. If you are using iSCSI, each clustered server must have one or more network adapters or HBAs that are dedicated to the cluster storage. The network you use for iSCSI cannot be used for network communication. In all clustered servers, the network adapters you use to connect to the iSCSI storage target should be identical. It is also recommended that you use Gigabit Ethernet or higher. (Note also that for iSCSI, you cannot use teamed network adapters.)
- **Shared storage compatible with Windows Server 2008 R2** To configure storage for a failover cluster, you must first prepare each server node for shared storage. For example, if your servers have access to an iSCSI SAN, you must configure an iSCSI target and disks that will be shared by all nodes, and you must then use iSCSI Initiator to provision the disks on each of these nodes.

  For a two-node failover cluster, the storage should contain at least two volumes (LUNs). The first volume functions as the disk witness, witness disk, or quorum witness, a volume that holds a copy of the cluster configuration database. Disk witnesses, known only as quorum disks in Windows Server 2003, are used in many, but not all, cluster configurations.

  The second volume contains the files being shared to users. Storage requirements include the following:

  - To use the native disk support included in failover clustering, use basic disks, not dynamic disks.
  - It is recommended that you format the storage partitions with NTFS. (For the disk witness, NTFS is required.)

  After you have provisioned the disks on each node, use Disk Management on the server on which you will configure the failover cluster to bring these new disks online, initialize them, and create the desired volumes on them. (You should bring the disks online *on the one server node only* before adding them to the cluster.)

NOTE **USE COMPATIBLE HARDWARE**

When deploying a storage area network (SAN) with a failover cluster, be sure to confirm with manufacturers and vendors that the storage, including all drivers, firmware, and software used for the storage, are compatible with failover clusters in Windows Server 2008 R2.

After you have met the hardware requirements and configured the storage on the cluster server, you can then install the Failover Cluster feature.

Installing the Failover Clustering Feature

Before creating a failover cluster, you have to install the Failover Clustering feature on all nodes in the cluster.

To install the Failover Clustering feature, begin by clicking Add Features in Server Manager. In the Add Features Wizard, select the Failover Clustering check box. Click Next and then follow the prompts to install the feature.

After the feature is installed on all nodes, you are ready to validate the hardware and software configuration.

Validating the Cluster Configuration

Before you create a new cluster, you can use the Validate A Configuration Wizard to ensure that your nodes meet the hardware and software prerequisites for a failover cluster.

To run the Validate A Configuration Wizard, first open Failover Cluster Manager in the Administrative Tools program group. In Failover Cluster Manager, click Validate A Configuration in the Management area or in the Actions pane, as shown in Figure 2-27.

After the wizard completes, make any configuration changes, if necessary, and then rerun the test until the configuration is successfully validated. After the cluster prerequisites have been validated, use the Create Cluster Wizard to create the cluster.

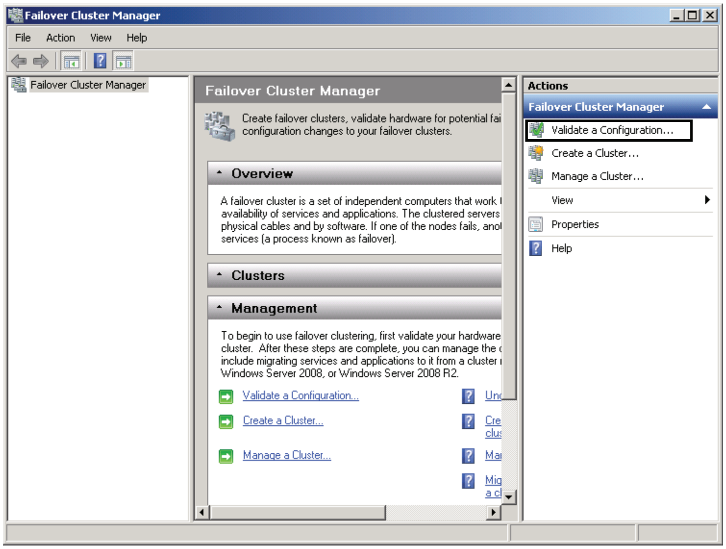

Create failover clusters, validate hardware for potential fai
configuration changes to your failover clusters.

A failover cluster is a set of independent computers that work t
availability of services and applications. The clustered servers
physical cables and by software. If one of the nodes fails, anot
services (a process known as failover).

To begin to use failover clustering, first validate your hardware
cluster. After these steps are complete, you can manage the c
include migrating services and applications to it from a cluster r
Windows Server 2008, or Windows Server 2008 R2.

FIGURE 2-27 Validating failover server prerequisites.

Running the Create Cluster Wizard

The next step in creating a cluster is to run the Create Cluster Wizard, which installs the software foundation for the cluster, converts the attached storage into cluster disks, and creates a computer account in Active Directory for the cluster. To launch this tool, in Failover Cluster Manager, click Create A Cluster in the Actions pane, as shown in Figure 2-28.

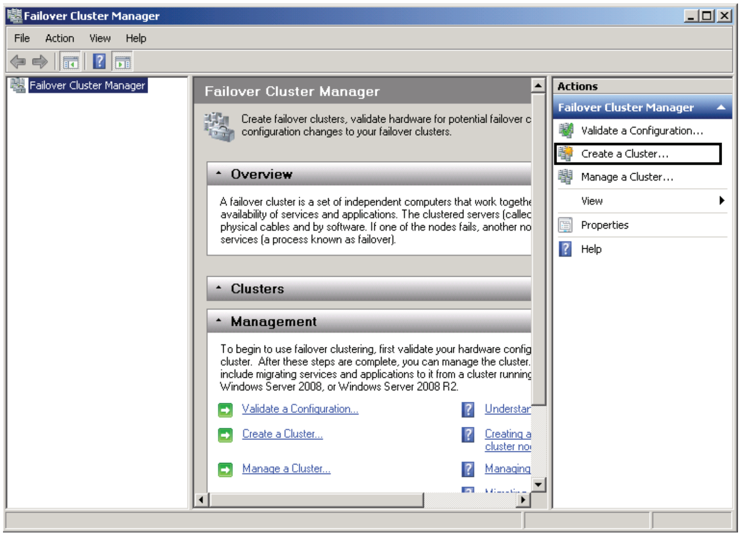

Create failover clusters, validate hardware for potential failover c
configuration changes to your failover clusters.

A failover cluster is a set of independent computers that work togethe
availability of services and applications. The clustered servers (callec
physical cables and by software. If one of the nodes fails, another no
services (a process known as failover).

To begin to use failover clustering, first validate your hardware config
cluster. After these steps are complete, you can manage the cluster.
include migrating services and applications to it from a cluster running
Windows Server 2008, or Windows Server 2008 R2.

FIGURE 2-28 Creating a failover cluster.

During the wizard, you add nodes to the cluster and give the cluster a name. If you have already configured storage for the cluster, it will be detected automatically and added to the cluster. At this point, configure the quorum settings.

Configuring Quorum Settings

The *quorum configuration* in a failover cluster determines the number of node failures the cluster can sustain before the cluster stops running. In Windows Server 2008 R2, you can choose from among the following four quorum configurations:

- **Node Majority quorum configuration** This quorum configuration is recommended for clusters with an odd number of nodes. In node majority, the failover cluster runs as long as a majority of the nodes are running.
- **Node and Disk Majority quorum configuration** This quorum configuration is recommended for clusters with an even number of nodes. In node and disk majority, the failover cluster uses a witness disk as a tiebreaker node, and the failover cluster then runs as long as a majority of these nodes are online and available.
- **Node And File Share Majority quorum configuration** This quorum configuration is recommended for clusters that have an even number of nodes and that lack access to a witness disk. In this configuration, a witness file share is used as a tiebreaker node, and the failover cluster then runs as long as a majority of these nodes are online and available.
- **No Majority: Disk Only quorum configuration** In this quorum configuration, which is recommended only for testing and not for production environments, the failover cluster remains available as long as a single node and its storage remain online.

To define the quorum configuration, right-click the *cluster* node in the Failover Cluster Manager console tree, point to More Actions in the shortcut menu, and then click Configure Cluster Quorum Settings, as shown in Figure 2-29. This step launches the Configure Cluster Quorum Wizard, the main page of which is shown in Figure 2-30.

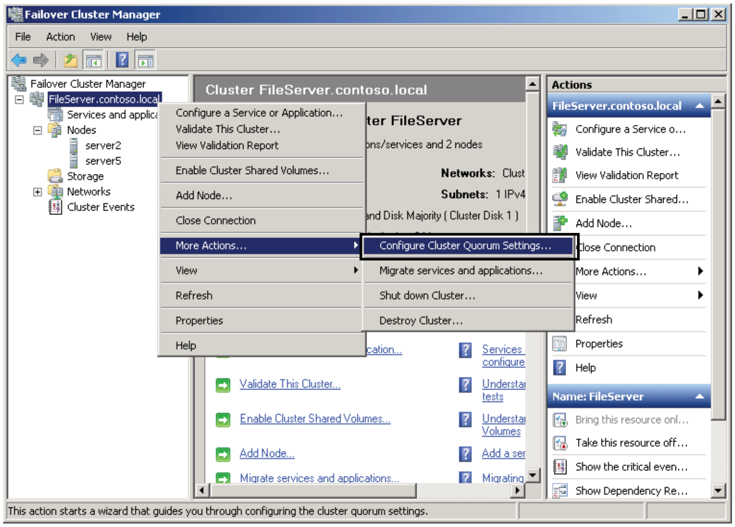

This action starts a wizard that guides you through configuring the cluster quorum settings.

FIGURE 2-29 Configuring quorum settings.

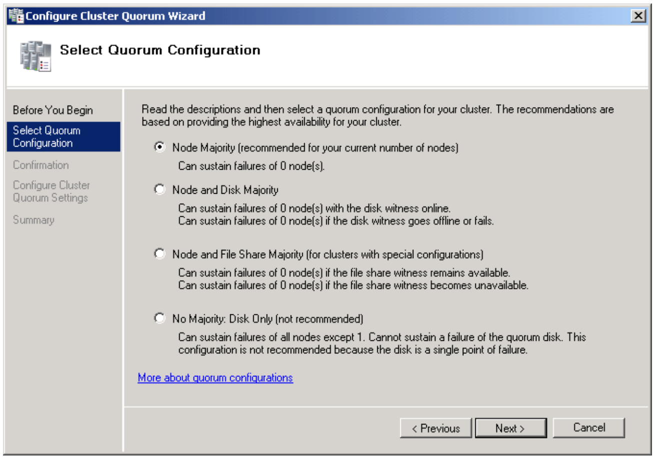
Configure Cluster Quorum Wizard
Before You Begin
Select Quorum Configuration
Confirmation
Configure Cluster Quorum Settings
Summary
< Previous
Next >
Cancel

Select Quorum Configuration

Read the descriptions and then select a quorum configuration for your cluster. The recommendations are based on providing the highest availability for your cluster.

● Node Majority (recommended for your current number of nodes)
Can sustain failures of 0 node(s).

○ Node and Disk Majority
Can sustain failures of 0 node(s) with the disk witness online.
Can sustain failures of 0 node(s) if the disk witness goes offline or fails.

○ Node and File Share Majority (for clusters with special configurations)
Can sustain failures of 0 node(s) if the file share witness remains available.
Can sustain failures of 0 node(s) if the file share witness becomes unavailable.

○ No Majority: Disk Only (not recommended)
Can sustain failures of all nodes except 1. Cannot sustain a failure of the quorum disk. This configuration is not recommended because the disk is a single point of failure.

More about quorum configurations

FIGURE 2-30 Selecting a quorum configuration.

If you select Node And Disk Majority, you must specify the volume you want to designate as the witness disk, as shown in Figure 2-31.

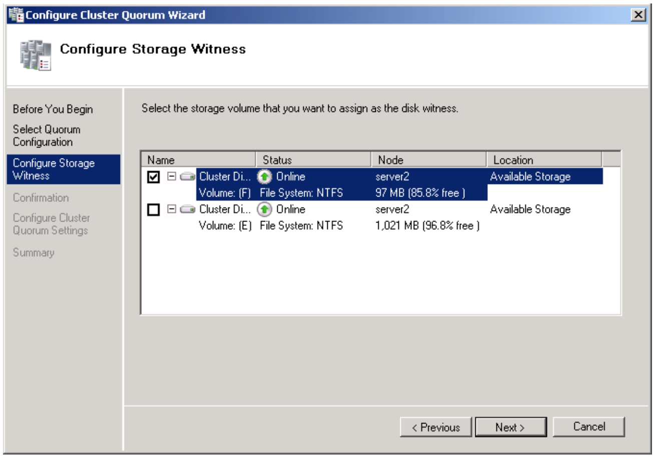
Configure Cluster Quorum Wizard
Configure Storage Witness
Before You Begin
Select Quorum Configuration
Configure Storage Witness
Confirmation
Configure Cluster Quorum Settings
Summary
< Previous
Next >
Cancel

Select the storage volume that you want to assign as the disk witness.

Name	Status	Node	Location
☑ Cluster Di...	Online	server2	Available Storage
Volume: (F)	File System: NTFS	97 MB (85.8% free)	
☐ Cluster Di...	Online	server2	Available Storage
Volume: (E)	File System: NTFS	1,021 MB (96.8% free)	

FIGURE 2-31 Selecting a disk for use as the quorum witness.

If you select Node And File Share Majority on the Select Quorum Configuration page, you use the wizard to specify a remotely hosted file share.

After you have defined the quorum configuration for the cluster, you are ready to configure the service or application for which you want to provide failover service.

Quick Check

1. What is a witness disk?
2. What is the quorum configuration of a failover cluster?

Quick Check Answers

1. A witness disk is a shared volume used in many failover clusters that contains a copy of the cluster configuration database.
2. The quorum configuration is what determines the number of node failures a failover cluster can sustain before the cluster stops running.

EXAM TIP

On the 70-643 exam, you might see basic questions about quorum configurations, witness disks, or witness file shares.

Configuring a Service for the Failover Cluster

The first step in configuring a service or application for failover service is to install the service or application on all the server nodes you want to host that service in the cluster. For example, if you want to configure high availability for a file server in a two-node cluster, you must add the File Services role on both server nodes in the cluster.

After you have added the desired service to the chosen server nodes in the cluster, you can run the the High Availability Wizard to configure failover for that service. To start the High Availability Wizard, in Failover Cluster Manager, right-click the *Services And Applications* node in the console tree and then select Configure A Service Or Application in the Action pane or the Configure area, as shown in Figure 2-32.

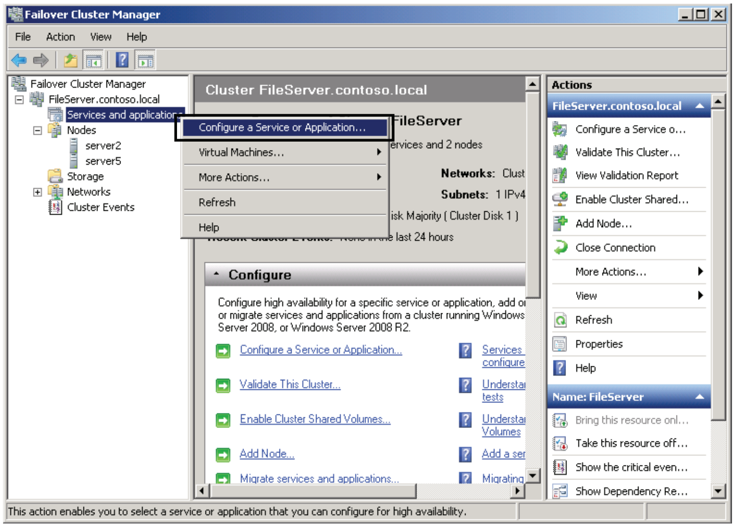

Configure high availability for a specific service or application, add o or migrate services and applications from a cluster running Windows Server 2008, or Windows Server 2008 R2.

This action enables you to select a service or application that you can configure for high availability.

FIGURE 2-32 Configuring a service for failover.

To complete the High Availability Wizard, perform the following steps:

1. On the Before You Begin page, review the text and then click Next.
2. On the Select Service Or Application page (shown in Figure 2-33), select the service or application for which you want to provide failover service (high availability) and then click Next.

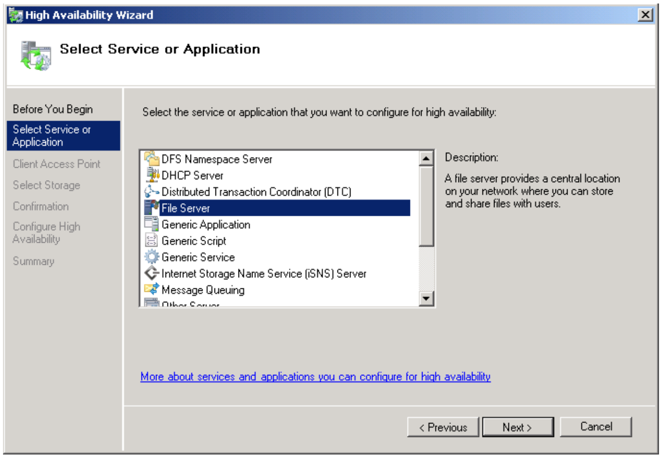
High Availability Wizard
Select Service or Application
Before You Begin
Select Service or Application
Client Access Point
Select Storage
Confirmation
Configure High Availability
Summary
DFS Namespace Server
DHCP Server
Distributed Transaction Coordinator (DTC)
File Server
Generic Application
Generic Script
Generic Service
Internet Storage Name Service (iSNS) Server
Message Queuing
Other Server
More about services and applications you can configure for high availability
< Previous
Next >
Cancel

Select the service or application that you want to configure for high availability:

Description:
A file server provides a central location on your network where you can store and share files with users.

FIGURE 2-33 Selecting a service for high availability.

3. Follow the instructions in the wizard to specify required details about the chosen service. For example, for the File Server service, you would need to specify the following:
   - A name for the clustered file server
   - Any IP address information that is not automatically supplied by your DHCP settings—for example, a static IPv4 address for this clustered file server
   - The storage volume or volumes the clustered file server should use
4. After the wizard runs and the Summary page appears, to view a report of the tasks the wizard performed, click View Report.
5. To close the wizard, click Finish.

EXAM TIP

Be sure to understand the wizards used in failover clustering for the 70-643 exam. Understand what they are used for and how to start them.

After you run the wizard, the service you have specified appears in the Failover Cluster Manager console tree beneath the *Services And Applications* node, as shown in Figure 2-34.

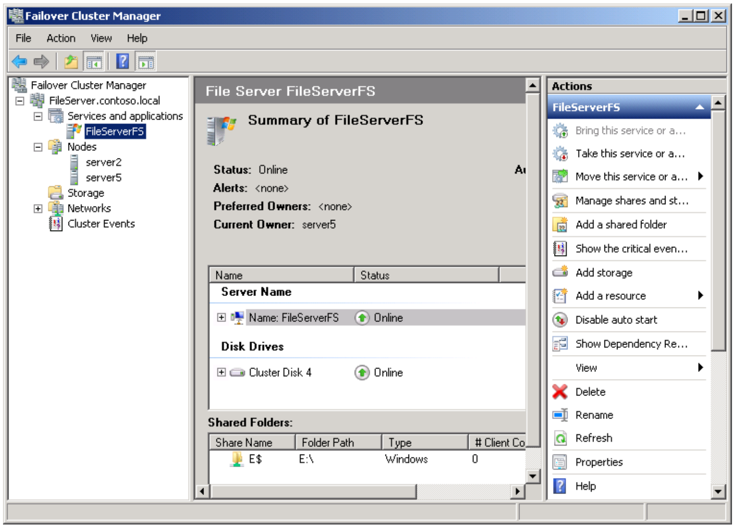

Summary of FileServerFS

Status: Online
Alerts: <none>
Preferred Owners: <none>
Current Owner: server5

Name	Status
Server Name	
Name: FileServerFS	Online
Disk Drives	
Cluster Disk 4	Online

Shared Folders:

Share Name	Folder Path	Type	# Client Co
E$	E:\	Windows	0

FIGURE 2-34 A service configured for high availability.

You can then configure preferred owner and failback behavior by right-clicking the service and then selecting Properties. The General tab, shown in Figure 2-35, enables you to designate one or more preferred owners of the service, in a specified order. These settings are used only if you enable failback on the Failover tab.

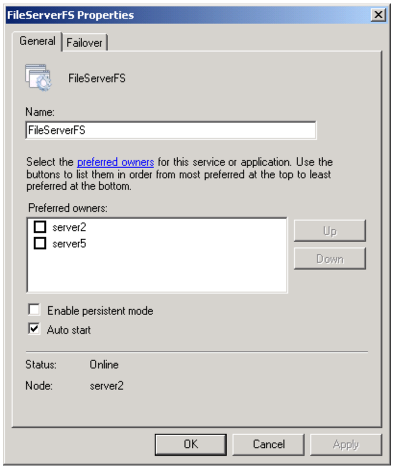

FileServerFS

Name: FileServerFS

Select the preferred owners for this service or application. Use the buttons to list them in order from most preferred at the top to least preferred at the bottom.

Preferred owners:

□ server2
□ server5

□ Enable persistent mode
☑ Auto start

Status: Online
Node: server2

FIGURE 2-35 Configuring preferred owners for a service.

Selecting the Failover tab reveals the configuration options shown in Figure 2-36. In the Failover (top) area of this tab, you can specify the maximum number of failures that you want to allow the failover cluster to sustain within a given time period. When the maximum number is exceeded, the service is left in the failed state. By default, a service is allowed to fail over only once every six hours.

In the Failback (lower) area of this tab, you can choose whether you want a failed service to be moved (fail back) automatically to the most preferred owner or node available, according to the settings on the General tab. If you choose to allow failback, you must specify the hours during which failback is allowed according to a 24-hour clock. For example, if you want to allow failback only during business hours—9 A.M. to 5 P.M.—choose the Allow Failback option and then specify failback between 9 and 17 hours.

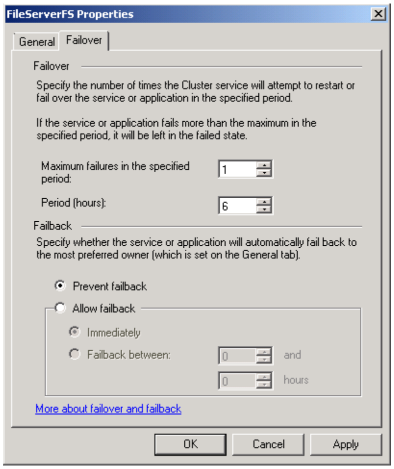

Failover

Specify the number of times the Cluster service will attempt to restart or fail over the service or application in the specified period.

If the service or application fails more than the maximum in the specified period, it will be left in the failed state.

Maximum failures in the specified period: 1

Period (hours): 6

Failback

Specify whether the service or application will automatically fail back to the most preferred owner (which is set on the General tab).

● Prevent failback
○ Allow failback
● Immediately
○ Failback between: 0 and 0 hours

FIGURE 2-36 Configuring failover properties for a service.

EXAM TIP

You need to understand the failover and failback settings for the 70-643 exam.

Testing the Failover Cluster

After you complete the wizard, test the failover cluster in Failover Cluster Management. In the console tree, make sure the *Services And Applications* node is expanded and then select the service you have just added with the High Availability Wizard. Right-click the clustered service, click Move This Service Or Application To Another Node, and then click the available node, as shown in Figure 2-37.

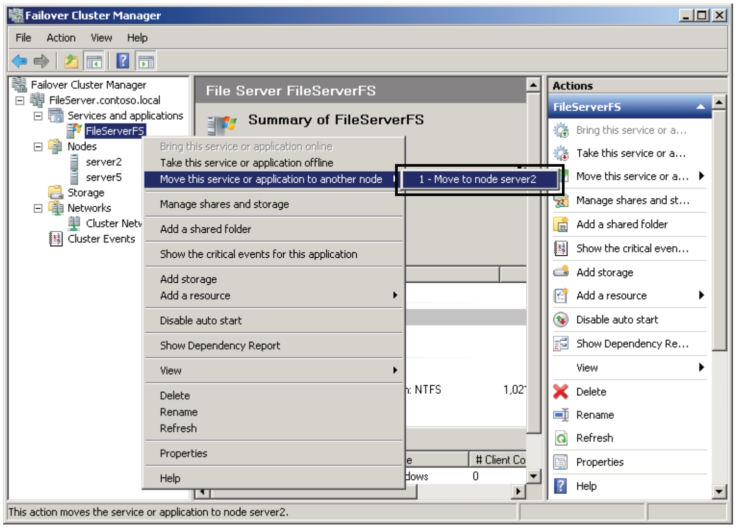

This action moves the service or application to node server2.

FIGURE 2-37 Testing a failover cluster by moving a service to another node.

You can observe the status changes in the center pane of the snap-in as the clustered service instance is moved, as shown in Figure 2-38. If the service moves successfully, the failover is functional.

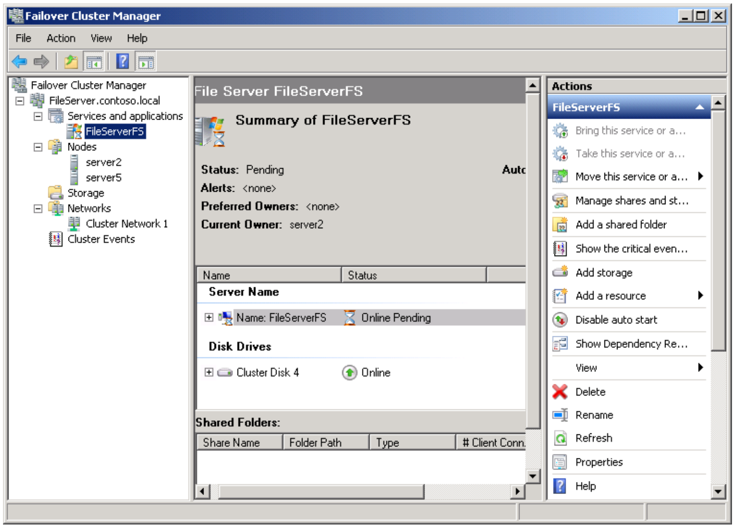

Status: Pending
Alerts: <none>
Preferred Owners: <none>
Current Owner: server2

Name	Status
Server Name	
Name: FileServerFS	Online Pending
Disk Drives	
Cluster Disk 4	Online

Shared Folders:

Share Name	Folder Path	Type	# Client Conn.

FIGURE 2-38 A service moving to another node.

Understanding Cluster Shared Volumes in Failover Clusters

In Windows Server 2008 and earlier, only one node could access a LUN or physical disk at any given time. If one application running on a LUN failed and needed to move to another node in the failover cluster, every single application on that LUN would also need to be failed over to the new node and potentially experience some downtime. To avoid this problem, applications were typically hosted on unique LUNs as a way to isolate failures. This strategy created another problem, however: a large number of LUNs that complicated setup and administration.

Windows Server 2008 R2 introduces *Cluster Shared Volumes* (CSVs), a new storage type intended for use with virtual machines. CSVs enable you to store on a single LUN the VHDs for multiple virtual machines. You can run the virtual machines on any node in the failover cluster, and when an application stored on a CSV fails, the application can be configured to fail over to a virtual machine on any other node in the cluster.

Aside from their use in failover clusters, CSVs also support live migration of virtual machines in Hyper-V. With live migration, a virtual machine is moved from one physical computer to another with almost no downtime. (Like CSVs, live migration is a new feature in Windows Server 2008 R2.)

NOTE **WHERE ARE CSVs STORED?**

On a Hyper-V host computer acting as a physical node in a failover cluster, CSVs appear as subfolders in the \ClusterStorage folder on the system drive. Example pathnames are C:\ClusterStorage\Volume1, C:\ClusterStorage\Volume2, and so on.

To enable Cluster Shared Volumes by using Failover Cluster Manager, perform the following steps:

1. In Failover Cluster Manager, select the failover cluster for which you want to enable CSVs.
2. Right-click the failover cluster and then click Enable Cluster Shared Volumes, as shown in Figure 2-39**.** The Enable Cluster Shared Volumes dialog box opens. Read and accept the terms and restrictions and click OK.

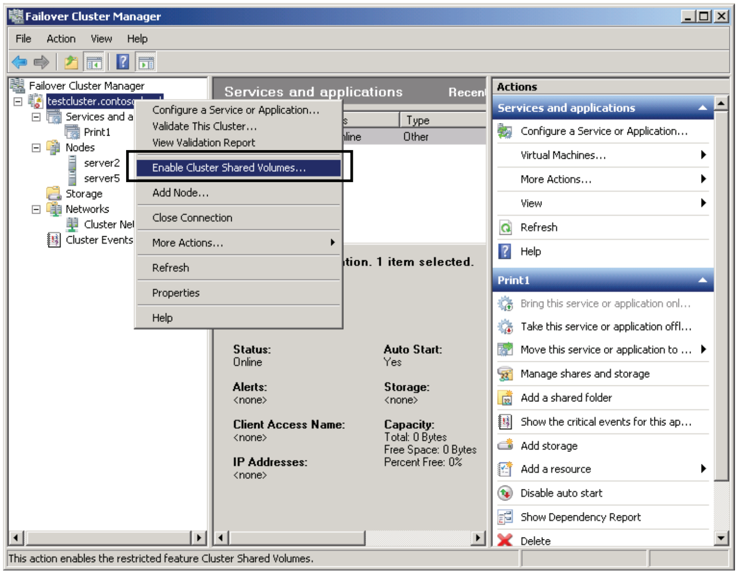

tion. 1 item selected.

Status: Online

Auto Start: Yes

Alerts: <none>

Storage: <none>

Client Access Name: <none>

Capacity: Total: 0 Bytes
Free Space: 0 Bytes
Percent Free: 0%

IP Addresses: <none>

This action enables the restricted feature Cluster Shared Volumes.

FIGURE 2-39 Enabling Cluster Shared Volumes.

To add a Cluster Shared Volume, perform the following steps:

1. In the Failover Cluster Manager snap-in, right-click the *Cluster Shared Volumes* node and then click Add Storage from the shortcut menu.
2. In Add Storage, select from the list of available disks and click OK. The disk or disks you selected appear in the Results pane for Cluster Shared Volumes.

MORE INFO **UNDERSTANDING CSVs**

For more information about the use of CSVs in a failover cluster, visit *http://technet.microsoft.com/en-us/library/ff182346(WS.10).aspx* or search on TechNet for "Using Cluster Shared Volumes in a Failover Cluster in Windows Server 2008 R2."

PRACTICE Explore Network Load Balancing and Failover Clustering

In this practice, you perform an exercise to configure a Network Load Balancing cluster. Then, you complete a virtual lab assignment in which you configure various services for failover in a failover cluster.

EXERCISE 1 Creating a Network Load Balancing Cluster

In this exercise, you begin by taking a snapshot of both the Server2 and Core1 virtual machines in Hyper-V. Then, you configure a Network Load Balancing cluster by using these two computers. Finally, you revert the virtual machines to their previous states.

1. Log on to Contoso.local from Server2 as a domain administrator. In the Server2 On Localhost window that contains the Server2 virtual machine in Hyper-V, select Snapshot from the Action menu.
2. In the Snapshot Name dialog box, type **Before Chapter 2 Lesson 2 Exercise 1** and then click Yes.
3. Repeat steps 1 and 2 on Core1.
4. On Server2, change the IPv4 address configuration of the Local Area Connection to the following static configuration:
   - IP address: 192.168.10.202
   - Subnet mask: 255.255.255.0
   - Default gateway: 192.168.10.1
   - DNS server: 192.168.10.1

   You can complete this configuration step at an elevated command prompt by typing **netsh interface ipv4 set address "local area connection" static 192.168.10.202 255.255.255.0 192.168.10.1**, pressing Enter, and then typing **netsh interface ipv4 set dns "local area connection" static 192.168.10.1** and pressing Enter again.
5. On Core1, type **sconfig** at the command prompt to start the System Configuration utility and then select option 8 to open the network settings.
6. Use the System Configuration utility to change the IPv4 address configuration of the Core1 network adapter (NIC Index# = 0) to the following static configuration:
   - IP address: 192.168.10.204
   - Subnet mask: 255.255.255.0
   - Default gateway: 192.168.10.1
   - Preferred DNS server: 192.168.10.1 (do not set an alternate DNS server)
7. Shut down both Server2 and Core1. (To shut down Core1, select 4 in the System Configuration utility to return to the main menu and then choose option 12.)

   Leave the virtual machine windows open even after the virtual machines are turned off.

8. After Server2 shuts down, in the Server2 On Localhost – Virtual Machine Connection window, select Settings from the File menu.
9. In the Settings For Server2 window, select Legacy Network Adapter from the list of hardware.
10. In the Legacy Network Adapter configuration area on the top-right part of the Settings For Server2 window, select Enable Spoofing Of MAC Addresses and then click OK.
11. Perform steps 8 to 10 on Core1, selecting Network Adapter instead of Legacy Network Adapter.

    This step is currently required to configure an NLB cluster with virtual machines in Windows Server 2008 R2.
12. Start both Server2 and Core1 and log on to Contoso.local from both computers as a domain administrator.
13. In Server Manager on Server2, right-click the *Features* node in the console tree and then click Add Features.
14. Use the Add Features Wizard to add the Network Load Balancing feature. If prompted to do so, restart the computer and log on to Contoso.local again as a domain administrator to complete the installation. (You do not need to restart if you are not prompted to do so.)
15. On Core1, type the following and then press Enter:

    dism /online /enable-feature /featurename:NetworkLoadBalancingHeadlessServer

EXAM TIP

For the exam, remember that you can use the Dism command to add a feature.

16. After you receive a message on Core1 indicating that the operation has completed successfully, switch to Server2 and open Network Load Balancing Manager from the Administrative Tools menu.
17. In the console tree, right-click the *Network Load Balancing Clusters* node and select New Cluster.
18. On the New Cluster: Connect page, type Server2.contoso.local in the Host text box and then click Connect.
19. Verify that Local Area Connection is selected beneath Interface Name with an Interface IP of 192.168.10.202 and then click Next.
20. On the New Cluster: Host Parameters page, review the default settings and then click Next.
21. On the New Cluster: Cluster IP Addresses page, click Add.

22. In the Add IP Addresses dialog box, specify an IPv4 address of 192.168.10.200 and a subnet mask of 255.255.255.0. You do not need to add or generate an IPv6 address. Click OK.

23. On the New Cluster: Cluster IP Addresses page, click Next.

24. On the New Cluster: Cluster Parameters page, type **testNLB.contoso.local** in the Full Internet Name text box. Review the other default settings and click Next.

25. On the New Cluster: Port Rules page, review the default settings and then click Finish.

    The operation takes a minute to complete. After it completes, the *Server2* node appears surrounded by a green box in the Network Load Balancing Manager console tree.

26. In the Network Load Balancing Manager console tree, right-click the *testNLB.contoso.local* node and select Add Host To Cluster from the shortcut menu.

27. On the Add Host To Cluster: Connect page, type Core1.contoso.local in the Host text box and then click Connect.

28. Verify that Local Area Connection is selected beneath Interface Name with an interface IP of 192.168.10.204 and then click Next.

29. On the Add Host To Cluster: Host Parameters page, review the default settings and then click Next.

30. On the Add Host To Cluster: Port Rules page, review the default settings and then click Finish.

    The operation takes a minute to complete. After it completes, the *Core1* node appears surrounded by a green box in the Network Load Balancing Manager console tree. The green boxes indicate that the NLB hosts have successfully converged.

31. At a command prompt on Server2, type **ping 192.168.10.200**.

    The NLB cluster responds to the ping attempt. You can also ping the NLB cluster successfully from Server1 and Core1.

32. In the Server2 On Localhost – Virtual Machine Connection window, select Revert from the Action menu.

    The Server2 virtual machine is reverted to its state before Exercise 1.

33. In the Core1 On Localhost– Virtual Machine Connection window, select Core1 from the Action menu.

    The Core1 virtual machine is reverted to its state before Exercise 1.

EXERCISE 2 Configuring a Failover Cluster

Visit *http://msevents.microsoft.com* and search for event ID 1032380228. Perform the virtual lab titled, “TechNet Virtual Lab: Creating a Highly Available Infrastructure.”

Lesson Summary

- You can configure groups of servers in Windows Server 2008 R2 to provide load balancing, scalability, or high availability for a particular service or application. These server groups are often called clusters and can be used for very different purposes. Typically, clusters are transparent and appear as a single server to clients.
- Round-robin DNS is a basic method of balancing requests for a single server between two or more servers. Round-robin is easy to configure but has significant limitations, such as the lack of awareness of server status.
- Network Load Balancing (NLB) is an installable feature of Windows Server 2008 and Windows Server 2008 R2. Like round-robin, NLB transparently distributes client requests for a single server between two or more servers. However, NLB overcomes the limitations of round-robin DNS by providing advanced features, such as the ability to redirect requests away from a downed or busy server automatically. NLB is often used to create web farms, which are NLB clusters that answer requests for a website or set of websites.
- Failover clustering is an installable feature of Windows Server 2008 R2 Enterprise Edition. A failover cluster is a group of computers that prevent downtime for selected applications and services. Servers (or nodes) in a failover cluster are connected to each other and to shared storage. Failover clusters have fairly elaborate hardware requirements, and you should be sure to review these requirements before making purchasing decisions.

Lesson Review

The following questions are intended to reinforce key information presented in this lesson. The questions are also available on the companion CD if you prefer to review them in electronic form.

> ***NOTE*** **ANSWERS**
>
> **Answers to these questions and explanations of why each answer choice is correct or incorrect are located in the "Answers" section at the end of the book.**

1. You work as a network administrator for Tailspintoys.com. Your job responsibilities include supporting company servers. The Tailspintoys.com network hosts a web server that runs on a single server named Websrv1. Recently, traffic to the website has been increasing, and the performance of the web server has been deteriorating. Traffic to the website is expected to continue to increase over the next five to eight years. You want a solution that can solve the performance problems of the web server and meet the increasing workload requirements for the website for the next five to eight years. What should you do?

   - **A.** Migrate the website to a more powerful server.
   - **B.** Use NLB to create a web farm to support the website.
   - **C.** Use failover clustering to support the website with multiple servers in a cluster.
   - **D.** Add a second web server and then use DNS round-robin to distribute web requests between the two servers. Add more servers as necessary.

2. You are configuring a failover cluster for a database server. You are assigning four nodes to the cluster. All nodes have access to a SAN, and adequate storage is available. Which of the following options should you choose for your quorum configuration?

   - **A.** Node Majority
   - **B.** Node And Disk Majority
   - **C.** Node And File Share Majority
   - **D.** No Majority: Disk Only

3. You work for a large company that uses an iSCSI-based storage area network (SAN) to provide storage for its 20 servers. You want to migrate four of your existing application servers to a failover cluster on Windows Server 2008 R2. Besides providing failover to the four applications servers, you also want to minimize the number of hardware components in the cluster, minimize the number of drives appearing in Disk Management on each physical cluster node, and prevent a single application server failure from triggering a failover on other application servers. Which of the following options will best help you achieve your goals? (Each answer presents part of the solution. Choose two.)

   - **A.** Create and configure cluster shared volumes for the failover cluster.
   - **B.** Create and configure pass-through disks for each application server.
   - **C.** Run each application server as a virtual machine in Hyper-V.
   - **D.** Run each application server on a separate LUN provisioned from the SAN.

Chapter Review

To further practice and reinforce the skills you learned in this chapter, you can perform the following tasks:

- Review the chapter summary.
- Review the list of key terms introduced in this chapter.
- Complete the case scenarios. These scenarios set up a real-world situation involving the topics of this chapter and ask you to create solutions.
- Complete the suggested practices.
- Take a practice test.

Chapter Summary

- Servers require block-based access to data to run operating systems and applications. Usually, direct-attached storage is used for this purpose. This type of storage includes all internally installed hard disks and externally attached storage.
- Windows Server 2008 and Windows Server 2008 R2 include the Virtual Disk Service (VDS) API, which exposes compatible storage subsystems to Windows Server 2008 administration tools such as Storage Manager for SANs.
- You can use Disk Management in Windows Server 2008 and Windows Server 2008 R2 to create simple volumes, spanned volumes, striped volumes, mirrored volumes, and RAID-5 volumes. You can also choose to extend or shrink existing volumes.
- Network Load Balancing (NLB) balances a workload among multiple servers. Clients connect to an NLB cluster by specifying a virtual computer name and virtual IP address. An available server in the NLB cluster then answers the request.
- Failover clustering is a solution that minimizes server downtime. In a failover cluster, cluster servers or nodes share the same storage. When one server fails, another server takes over for the failed server.

Key Terms

Do you know what these key terms mean? You can check your answers by looking up the terms in the glossary at the end of the book.

- block-based
- cluster
- Cluster Shared Volumes
- iSCSI initiator
- iSCSI target
- parity information

- partition style
- quorum configuration
- round-robin DNS
- SAN fabric
- web farm
- witness disk

Case Scenarios

In the following case scenarios, you apply what you've learned about storage and clusters in this chapter. You can find answers to these questions in the "Answers" section at the end of this book.

Case Scenario 1: Designing Storage

You are an IT support specialist for Woodgrove Bank. Your manager informs you that the bank has decided to create a SAN for shared storage among its servers, and you have been asked to research SAN technology options. Migration of chosen servers to SAN storage will occur in approximately one year.

The primary goal for the future SAN is to provide both flexible storage and extremely low latency for database servers. Other goals are to take advantage of the existing networking expertise of the IT staff as much as possible and to facilitate as much administration of the SAN as possible through the Windows Server 2008 R2 interface. No one currently employed on the IT staff has any expertise working with SANs.

1. Given the storage needs of the organization, which connection technology should you choose for the SAN?
2. Which element should you seek in vendor solutions that will enable you to meet the administrative goals of the SAN?

Case Scenario 2: Designing High Availability

You are a server administrator for Trey Research. Recently, Trey Research purchased a line-of-business application named App1 that is to be used heavily by all 500 employees throughout the day. App1 is a web-based application that connects to a back-end database.

You and other members of the IT staff are currently designing the servers to host App1 and its database. In general, the design team foresees two servers or clusters, one to host IIS and App1 and the second to host the database. All servers must run Windows Server 2008 R2. The goals for the server design are to minimize downtime and to provide the best possible performance for both the application and the database. In addition, the solution must use a single database that is always internally consistent. All tables must always be visible to App1.

Within the design team, you have been asked to research cluster solutions for the web application server and database server.

1. Which clustering technology built into Windows Server 2008 R2 is most suitable for the web application server, and why?
2. Which clustering technology built into Windows Server 2008 R2 is most suitable for the database server, and why?

Suggested Practices

To help you successfully master the exam objectives presented in this chapter, complete the following tasks.

Configure Storage

Perform the first practice on three reserved virtual machines running Windows Server 2008 R2. If you have access to a system with three extra disks, virtual or physical, you should perform the second practice.

- **Practice 1** Visit the Microsoft Download Center at *http://download.microsoft.com* and search for "Microsoft iSCSI Software Target." Download the latest version of this software and install it on a server running Windows Server 2008 R2. Use the software to create an iSCSI target that specifies iSCSI initiators running on two other servers. Create two VHDs and attach them to the iSCSI target. Use the iSCSI Initiator tool on the other servers to connect to the target and provision the new disks.
- **Practice 2** On a Windows Server 2008 system, create a RAID-5 volume. Save data to the volume. Bring one of the disks offline and then attempt to access the data.

Configure High Availability

Perform both practices. The second practice requires you to first complete Practice 1 under "Configure Storage."

- **Practice 1** Watch the following sequence of four videos available on TechNet. These videos demonstrate how to configure a failover cluster that uses cluster shared volumes in Hyper-V Server 2008 R2.
  - "Hyper-V Server 2008 R2: Bare Metal to Live Migration," available at *http://technet.microsoft.com/en-us/edge/Ff955827*.
  - "Hyper-V R2: Building a Hyper-V R2 Cluster," available at *http://technet.microsoft.com/en-us/edge/video/ff955826*.
  - "Hyper-V R2: Making Highly Available VMs," available at *http://technet.microsoft.com/en-us/edge/video/ff955812*.
  - "Hyper-V R2: Introducing Cluster Shared Volumes," available at *http://technet.microsoft.com/en-us/edge/Video/ff955813*.

- **Practice 2** Add the Failover Clustering feature and the File Services role to two virtual machines running Windows Server 2008 R2 that have provisioned disks from an iSCSI target on a third server. Create a cluster with the two servers by using the New Cluster Wizard, and then use the High Availability Wizard to configure failover service for a file server.

Take a Practice Test

The practice tests on this book's companion CD offer many options. For example, you can test yourself on just one exam objective, or you can test yourself on all the 70-643 certification exam content. You can set up the test so that it closely simulates the experience of taking a certification exam, or you can set it up in study mode so that you can look at the correct answers and explanations after you answer each question.

MORE INFO **PRACTICE TESTS**

For details about all the practice test options available, see the "How to Use the Practice Tests" section in this book's introduction.

CHAPTER 6

Managing Web Server Security

From a systems administration standpoint, one of the main goals for managing web servers is to maintain a high standard of security. Security is an important concern in all areas of IT, but it's especially important for information and applications that are readily accessible to large numbers of users. In this chapter, you learn how to configure security for a Windows Server 2008 R2 web server.

Lesson 1, "Configuring IIS Security," focuses on securing access to Internet Information Services 7 (IIS 7) and the content it contains. You learn how to configure permissions for remote management and how to increase the security of the server by disabling or removing unneeded features and options. In Lesson 2, "Controlling Access to Web Services," you learn about ways in which you use authentication and authorization. You also learn how to increase security through server certificates and IP address restrictions.

Exam objectives in this chapter:

- Configure Web applications.
- Manage Web sites.
- Manage the Web Server (IIS) role.
- Configure SSL security.
- Configure Web site authentication and permissions.

Lessons in this chapter:

Before You Begin

To complete the lessons in this chapter, you should have:

- Installed the Web Server (IIS) server role on Server2.contoso.local by using the default installation options for this server role. If you have created additional websites or web applications in previous exercises, you may leave them configured on this server.
- The ability to create and manage websites and web applications. These topics are covered in Chapter 5, "Installing and Configuring Web Applications."

REAL WORLD

Anil Desai

The primary goal for systems administrators who are responsible for managing access to Web Services is to minimize the potential for unauthorized access to and misuse of applications or data. One of the primary ways to secure a server is by reducing its attack surface. If certain web applications do not require a particular technology (for example, support for the Microsoft .NET Framework), you can reduce potential unauthorized access to the system by disabling that feature.

Another major strategy related to web server security is *defense in depth*. This technique involves a multilayered security approach. Security options include authentication, authorization, file system permissions, and other settings that provide multiple barriers to access. These security mechanisms work together to ensure that only authorized users have access to the system. In addition, if one layer of security is incorrectly configured or is compromised, other security settings can help restrict or prevent unauthorized access.

Security settings can often be difficult and complicated to manage, and this complexity reduces security because many systems administrators find it challenging to set up the appropriate permissions. IIS simplifies security management through a hierarchical arrangement that helps organize settings and content. For example, you can apply security-related settings at the server level, for specific websites, for specific web applications, or directly on virtual directories, physical files, and folders.

In general, applying permissions at higher levels in the hierarchy simplifies administration. Figure 6-1 shows how objects, such as the web server, websites, web applications, and other items, are arranged into nested parent–child relationships. In general, settings placed on higher-level objects (such as a website) apply automatically to all the lower-level objects (such as multiple web applications). Administrators can then override settings for specific web applications by using

whatever method is dictated by business or technical requirements. The result of this configuration strategy is a high level of security with minimal administrative effort.

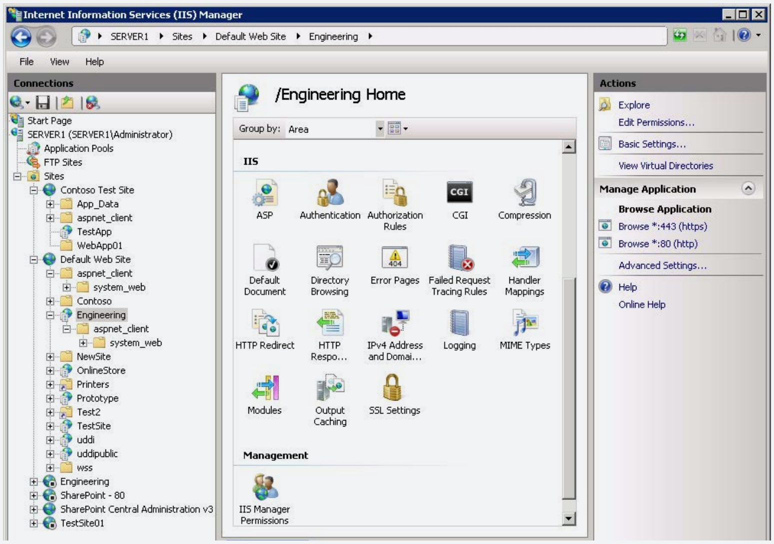

FIGURE 6-1 Viewing a hierarchy of objects in IIS Manager.

Lesson 1: Configuring IIS Security

In this lesson, you learn how to configure and manage security for the Web Server (IIS) server role and its associated components. You first learn how to determine the permissions administrators have on web servers. You learn ways to extend IIS administration capabilities to other users and web developers in your organization through remote management and delegation settings. You then learn how to increase security by configuring request handlers and their associated settings to minimize risks related to the execution of unwanted or malicious code or content.

After this lesson, you will be able to:

- Describe the security architecture of IIS, including built-in accounts.
- Enable remote management features for IIS web servers.
- Configure IIS Manager users, permissions, and feature delegation for distributed administration.
- Manage request handlers and handler mappings to reduce the attack surface of the web server.

Estimated lesson time: 75 minutes

Understanding IIS 7 Security Accounts

When you add the Web Server (IIS) role to a computer running Windows Server 2008 or Windows Server 2008 R2, as discussed in Chapter 5, the process makes numerous changes and additions to the configuration of the server. In earlier versions of IIS, each installation used service accounts that were based on the name of the server. Because the accounts and their security identifiers (SIDs) were different, copying web content and settings between web servers required multiple steps.

In IIS 7, a built-in, internal account named IUSR and a local security group called IIS_IUSRS are used on each computer running Windows Server 2008 and Windows Server 2008 R2 web server. Passwords for the accounts are managed internally, so administrators do not need to keep track of them.

Managing File System Permissions

To implement security, web server administrators must be able to define which content should be protected. They must also be able to specify which users or groups of users have access to protected content. Permissions settings for web content are managed through NTFS file system permissions. These permissions can be administered directly by using Windows Explorer or by right-clicking a specific object in the IIS Manager hierarchy and then clicking Edit Permissions. As shown in Figure 6-2, the permissions settings display which users or groups of users have access to the content and which permissions they have. IIS uses these permissions to determine whether credentials are required when attempting to complete a request from a web client.

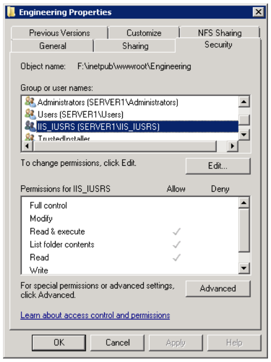

Object name: F:\inetpub\wwwroot\Engineering

Group or user names:

Administrators (SERVER1\Administrators)
Users (SERVER1\Users)
IIS_IUSRS (SERVER1\IIS_IUSRS)
TrustedInstaller

To change permissions, click Edit.

Permissions for IIS_IUSRS	Allow	Deny
Full control		
Modify		
Read & execute	✓	
List folder contents	✓	
Read	✓	
Write		

For special permissions or advanced settings, click Advanced.

Learn about access control and permissions

FIGURE 6-2 Viewing permissions for a folder within the Engineering website.

Configuring IIS Administration Features

When you add the Web Server (IIS) role to a computer running Windows Server 2008 or Windows Server 2008 R2, the default configuration enables only local administration of the server. This enhances security because users of other computers are unable to use IIS Manager to make changes to the server's configuration. Although this is appropriate for small, simple installations, often systems administrators benefit from the ability to use IIS Manager to configure the server remotely.

In many environments, multiple systems administrators manage websites and web applications. In large deployments, it is common for several administrators to be responsible for the same web server. For example, a single IIS server might host several important web applications, each of which is administered by a different individual or group. In hosting situations—when an organization provides IIS server access to subscribers—you must enable subscribers to control certain web content and features. In this case, subscribers act as remote administrators for certain portions of the servers. Remote administration is helpful for both multiple administrators and management performance from multiple locations.

To enable remote administrators to manage IIS, you must first enable remote management on the server. You can then define and configure IIS Manager users. *Feature delegation* enables you to specify which actions remote administrators can perform.

Enabling Remote Management

To enable remote management functionality, you first add the Management Service role service to the local server. You can do this by using Server Manager. Right-click the Web Server (IIS) role in the Roles folder and then select Add Role Services. Add Management Service, which is located in the Management Tools section of the available role services.

The Management Service role service is associated with the Web Management Service (WMSVC), which works by using a standard HTTP or HTTPS connection. Communications are configured to transmit over port 8172 by default. Assuming that traffic is allowed on this port through any firewalls or network security devices, this enables remote administrators to manage their IIS servers over a local network connection or over the Internet.

After you have added the Management Service role service to the Web Server (IIS) role, you can use IIS Manager to enable remote management. To do this, open IIS Manager and select the server object in the left pane. Double-click Management Service from the Management section in the Features view. (See Figure 6-3.)

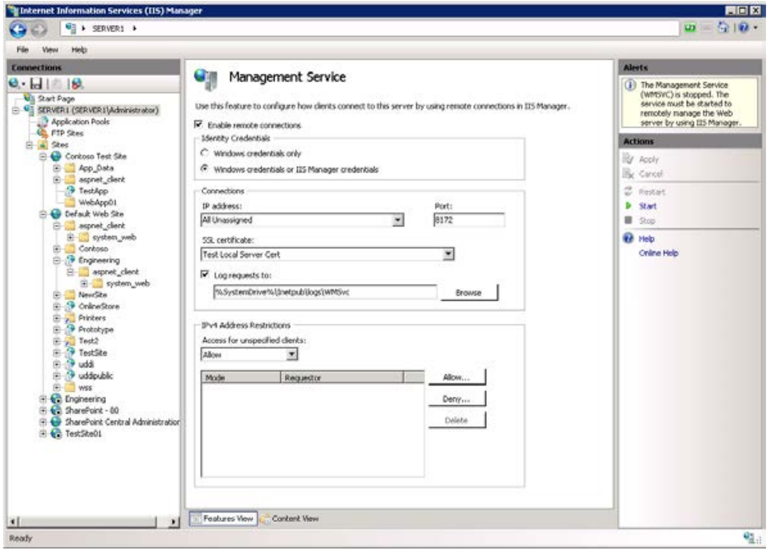

Internet Information Services (IIS) Manager

Management Service

Use this feature to configure how clients connect to this server by using remote connections in IIS Manager.

☑ Enable remote connections

Identity Credentials

○ Windows credentials only

● Windows credentials or IIS Manager credentials

Connections

IP address: All Unassigned

Port: 8172

SSL certificate: Test Local Server Cert

☑ Log requests to: %SystemDrive%\Inetpub\logs\WMSvc

IPv4 Address Restrictions

Access for unspecified clients: Allow

Mode	Requestor

Alerts

The Management Service (WMSVC) is stopped. The service must be started to remotely manage the Web server by using IIS Manager.

Ready

FIGURE 6-3 Configuring Management Service using IIS Manager.

Initially, the Enable Remote Connections option will be cleared. To enable authorized users to connect to IIS over the network, select Enable Remote Connections. The Identity Credentials section enables you to specify whether you will allow authentication by using Windows credentials only (the default setting) or by using IIS Manager credentials as well.

The Connections portion of the settings enables you to specify the IP addresses and ports on which the management service will respond. The default setting is for the service to respond to all available IP addresses on port 8172. If your web server is configured with multiple network connections or IP addresses, you can increase security by restricting remote access connections to a specific address. The SSL Certificate section enables you to select one of the SSL certificates that has been configured on the local server. You can also configure the path into which remote management requests are logged. The default is *%SystemDrive%* \Inetpub\Logs\WMSvc.

Finally, the IPv4 Address Restrictions section enables you to increase security by restricting which computers can connect to IIS remotely. As shown in Figure 6-4, you can configure rules based on a specific IPv4 address or on an address range (which is defined by a combination of an IP address and subnet mask). The Access For Unspecified Clients drop-down list in the Management Service feature defines whether IP addresses without entries are allowed or denied. You can then create Allow or Deny entries to define which IP addresses can connect. These options are most useful when you have control over the groups of computers that will be used for administering web services.

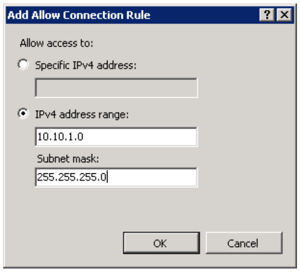

Allow access to:

○ Specific IPv4 address:

● IPv4 address range: 10.10.1.0

Subnet mask: 255.255.255.0

FIGURE 6-4 Configuring IPv4 address restrictions for Management Service in IIS Manager.

Because WMSVC is stopped by default, you must click the *Start* command in the Actions pane to start allowing remote connections. You must stop the service to make changes to the configuration.

Understanding IIS Manager Users

To connect to a Windows Server 2008 or Windows Server 2008 R2 web server by using IIS Manager, users must have the necessary permissions. Users who are logged on to a computer running Windows Server 2008 or Windows Server 2008 R2 with administrator credentials have the necessary permissions to complete all the available tasks on the server. For other types of users, such as remote systems administrators, you must decide how you want to manage permissions.

By default, the Web Server (IIS) role enables permissions to be assigned using Windows authentication only. This means that all administrators who attempt to manage IIS must have Windows-based credentials and permissions. Windows authentication is most appropriate for environments in which all the web server administrators belong to the same domain. Users who are logged on to the domain will not have to supply credentials manually when they connect to a server using IIS Manager, assuming they have the necessary permissions. Windows authentication is also useful when you plan to create either local or domain accounts for all the administrators who need access to IIS Manager.

In some cases, it might be impractical to create local or domain accounts for each of the potential IIS administrators. For example, web service hosting companies can have hundreds of users who require the ability to manage their servers. In these environments, each user can generally modify specific settings for his or her own website. These users should not have access to other users' websites and often will be restricted to changing only certain settings. To support these scenarios, you must enable the Windows Credentials Or *IIS Manager*

Credentials option. When this option is enabled by using the Management Service described in the previous section, you can create username and password combinations solely for the purpose of managing IIS. These credentials can then be given to other users and administrators so they can connect to the web server without requiring individual Windows accounts for each of the users.

Creating IIS Manager Users

The IIS Manager utility enables you to define which users can connect to and administer websites and web services. To configure these settings:

1. Open IIS Manager and select a server in the left pane.
2. Click IIS Manager Users in the Management section of the features view. By default, the IIS installation will not contain any locally defined users.
3. To create a new user, first click Open Feature in the Actions pane and then click the *Add User* command in the Actions pane. You will be prompted to provide a username and to type and confirm a password. (See Figure 6-5.) These settings are defined locally in IIS, so it is not necessary to use a fully qualified username that is compatible with your domain design.

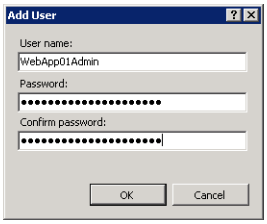

User name: WebApp01Admin

FIGURE 6-5 Adding an IIS Manager user.

In addition to configuring permission through IIS Manager users, you can use group membership settings to determine which users can connect remotely. Users who have permission to log on to the local computer and to use IIS Manager will be able to do so from a remote computer.

Defining IIS Management Permissions

So far, you have learned how to enable remote management and how to specify which users can use IIS Manager to administer a web server. Next, you must decide which permissions remote administrators will have after they connect. In some cases, you might want to enable a remote administrator to have full administrative access to the web server. In other cases, you will want to restrict access to only specific websites or web applications. You can configure IIS Manager Permissions at the website and application levels. However, you cannot configure permissions directly at the server level. This helps ensure that users are given permissions to modify the settings for only the specific websites and web applications to which they need access.

To manage permissions, select a website or web application and then click IIS Manager Permissions in the Management section of the Features View. By default, new IIS Manager users are not given permissions to connect to a specific website or web application. To enable a new user to connect at the selected level, first click Open Feature in the Actions pane and then click the *Allow User* command in the Actions pane. You are given the opportunity to specify a Windows user or an IIS Manager user (if IIS Manager credentials are accepted), as shown in Figure 6-6. If you are using the Windows option, you can select an existing user or group that is defined either in the domain (if the server is a member of a domain) or locally.

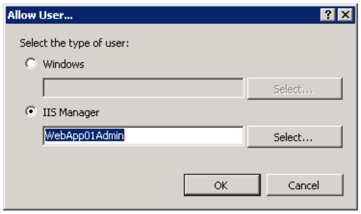

Allow User...

Select the type of user:

○ Windows

● IIS Manager

WebApp01Admin

FIGURE 6-6 Allowing a user to administer a website.

When users connect to IIS remotely, they can access only those websites and web applications on which they have been allowed. By default, permissions from higher-level objects are inherited automatically by lower-level objects.

To simplify administration of many users, two commands are available when managing permissions for a website. *Show All Users* provides a list of all the users available on the IIS installation. *Show Only Site Users* restricts the display to only the users who have access to the site.

Configuring Feature Delegation

The ability to define users and permissions enables you to manage administration based on site content structure. However, it is also important to determine which features users can view and configure. For example, you might want a web server administrator to connect to the Default Web Site, but you do not want her to be able to change authentication settings. Delegation is the process by which an administrator can determine which features of IIS a user can view and change.

Default settings for feature delegation are defined initially at the server level in IIS. To access these settings by using IIS Manager, select the *server* object in the left pane and then double-click Feature Delegation in the Management section of Features View, as shown in Figure 6-7.

The list of items available for delegation includes all the features that have been added through the Web Server (IIS) server role and enabled role services. To change the setting for a feature, select it from the list and use the commands in the Set Feature Delegation section of the Actions pane. Most features have options of Read Only or Read/Write. In addition, some items have a Configuration Read/Write or Configuration Read Only setting. These settings enable web developers to specify settings in their configuration files or to manage them based on database settings. The Not Delegated setting means that the feature has not been enabled for delegation at lower levels and is not available for configuration. You can also use the Delegation option in the Group By drop-down list to determine quickly how all the settings have been configured, as shown in Figure 6-8.

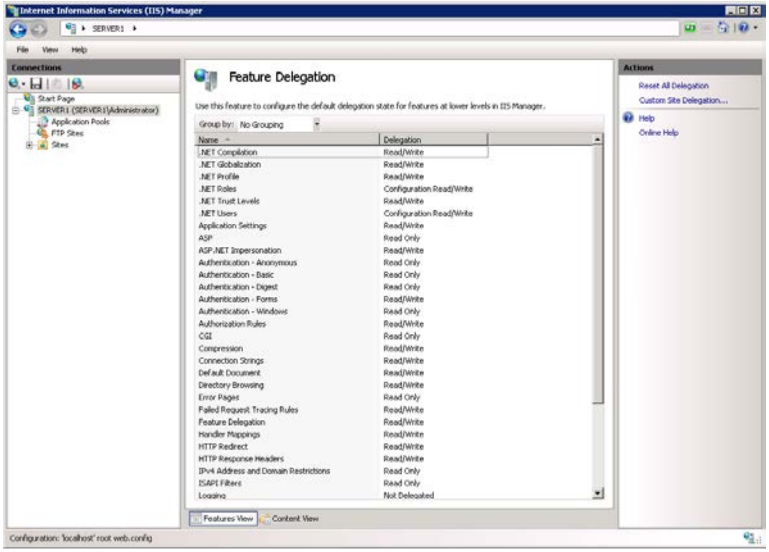

Feature Delegation

Use this feature to configure the default delegation state for features at lower levels in IIS Manager.

Group by: No Grouping

Name	Delegation
.NET Compilation	Read/Write
.NET Globalization	Read/Write
.NET Profile	Read/Write
.NET Roles	Configuration Read/Write
.NET Trust Levels	Read/Write
.NET Users	Configuration Read/Write
Application Settings	Read/Write
ASP	Read Only
ASP.NET Impersonation	Read/Write
Authentication - Anonymous	Read Only
Authentication - Basic	Read Only
Authentication - Digest	Read Only
Authentication - Forms	Read/Write
Authentication - Windows	Read Only
Authorization Rules	Read/Write
CGI	Read Only
Compression	Read/Write
Connection Strings	Read/Write
Default Document	Read/Write
Directory Browsing	Read/Write
Error Pages	Read Only
Failed Request Tracing Rules	Read/Write
Feature Delegation	Read/Write
Handler Mappings	Read/Write
HTTP Redirect	Read/Write
HTTP Response Headers	Read/Write
IPv4 Address and Domain Restrictions	Read Only
ISAPI Filters	Read Only
Logging	Not Delegated

Configuration: 'localhost' root web.config

FIGURE 6-7 Viewing Feature Delegation settings for an IIS web server.

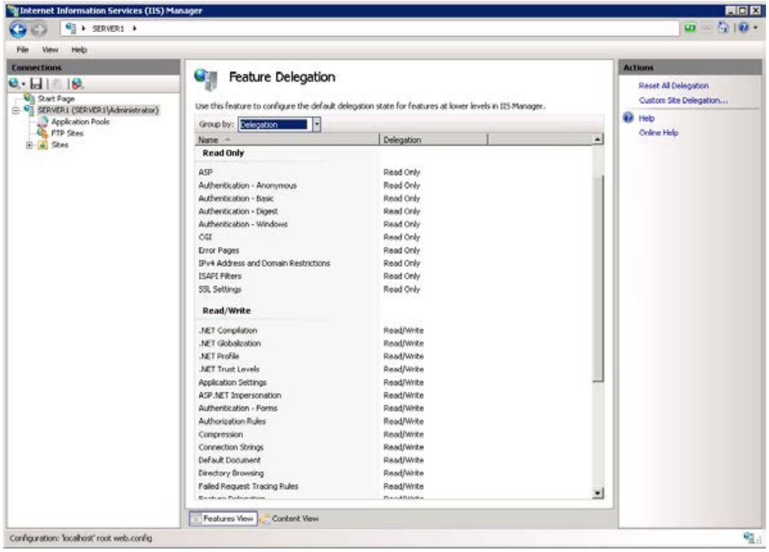

Feature Delegation

Use this feature to configure the default delegation state for features at lower levels in IIS Manager.

Group by: Delegation

Name	Delegation
Read Only	
ASP	Read Only
Authentication - Anonymous	Read Only
Authentication - Basic	Read Only
Authentication - Digest	Read Only
Authentication - Windows	Read Only
CGI	Read Only
Error Pages	Read Only
IPv4 Address and Domain Restrictions	Read Only
ISAPI Filters	Read Only
SSL Settings	Read Only
Read/Write	
.NET Compilation	Read/Write
.NET Globalization	Read/Write
.NET Profile	Read/Write
.NET Trust Levels	Read/Write
Application Settings	Read/Write
ASP.NET Impersonation	Read/Write
Authentication - Forms	Read/Write
Authorization Rules	Read/Write
Compression	Read/Write
Connection Strings	Read/Write
Default Document	Read/Write
Directory Browsing	Read/Write
Failed Request Tracing Rules	Read/Write

Configuration: 'localhost' root web.config

FIGURE 6-8 Viewing Feature Delegation configuration grouped by the delegation setting.

The settings you define at the server level automatically apply to all child websites and applications by default. In some cases, you will want to restrict feature delegation at the site level. To do this, click the *Custom Site Delegation* command in the Actions pane. This brings up the Custom Site Delegation screen, as shown in Figure 6-9, and enables you to select specific sites to which you want delegation settings to apply.

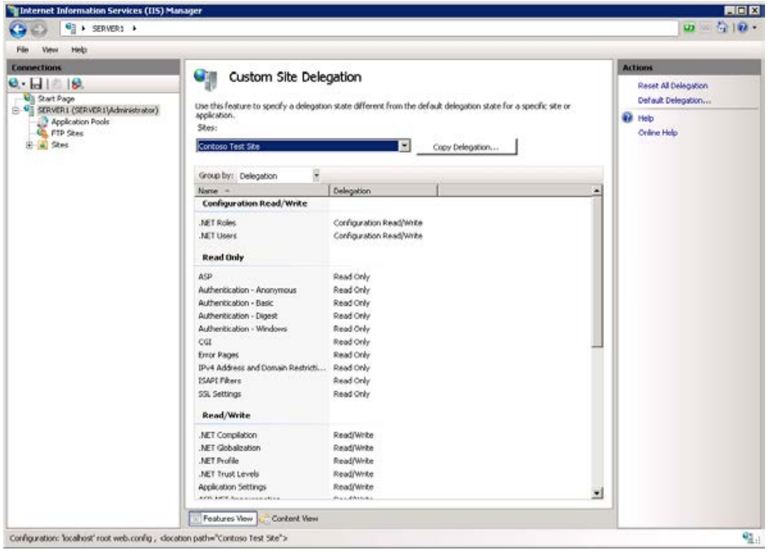

Use this feature to specify a delegation state different from the default delegation state for a specific site or application.

Sites: Contoso Test Site

Group by: Delegation

Name	Delegation
Configuration Read/Write	
.NET Roles	Configuration Read/Write
.NET Users	Configuration Read/Write
Read Only	
ASP	Read Only
Authentication - Anonymous	Read Only
Authentication - Basic	Read Only
Authentication - Digest	Read Only
Authentication - Windows	Read Only
CGI	Read Only
Error Pages	Read Only
IPv4 Address and Domain Restricti...	Read Only
ISAPI Filters	Read Only
SSL Settings	Read Only
Read/Write	
.NET Compilation	Read/Write
.NET Globalization	Read/Write
.NET Profile	Read/Write
.NET Trust Levels	Read/Write
Application Settings	Read/Write

Configuration: 'localhost' root web.config , <location path="Contoso Test Site">

FIGURE 6-9 Specifying Custom Site Delegation settings.

The *Copy Delegation* button enables you to copy the currently selected settings to one or more websites on the server. You can also use the *Reset To Inherited* and *Reset All Delegation* commands in the Actions pane to change groups of settings quickly to earlier values. (*Reset To Inherited* appears in the Actions pane only when you select a particular item or group of items for configuration within a site.) You use feature delegation settings to determine which parts of the system configuration will be available when remote users connect to the server by using IIS Manager.

NOTE **IMPLEMENTING REMOTE MANAGEMENT SECURITY**

When implementing remote management security, keep in mind the specific administration requirements. Some settings, such as IIS Manager Users and Feature Delegation, can be configured only at the level of the web server. That makes these settings applicable to all the lower-level objects. Alternatively, IIS Manager Permissions can be configured for specific websites and web applications to enable you to implement granular security for those users who should have access to only limited portions of the web server.

Connecting to a Remote Server by Using IIS Manager

After you have enabled remote management and configured the appropriate permissions and settings, remote users will be able to connect to the server by using the IIS Manager console. To verify the configuration from either the local computer or from a remote computer that has the IIS Manager console installed, use Start Page in IIS Manager or the File menu to connect to IIS. As shown in Figure 6-10, remote users will be able to connect to the server at one of several levels. The available commands include:

- *Connect To A Server*
- *Connect To A Site*
- *Connect To An Application*

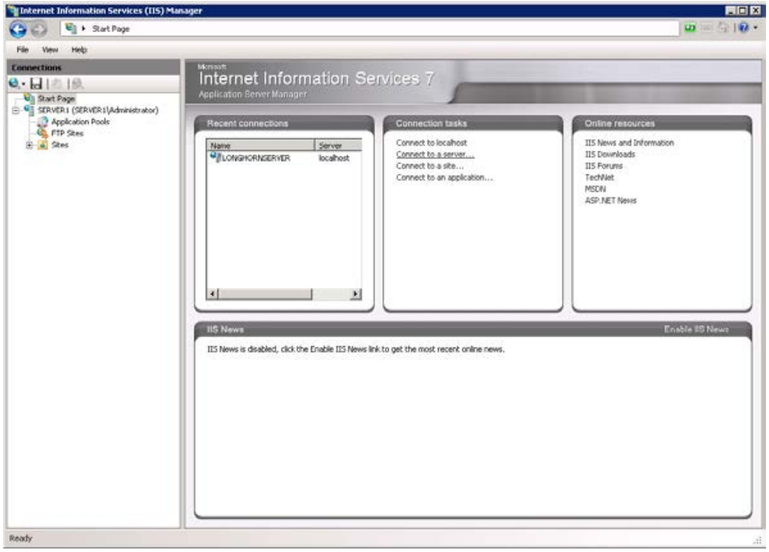

IIS News is disabled, click the Enable IIS News link to get the most recent online news.

FIGURE 6-10 Connecting to a remote installation of IIS.

> ***MORE INFO*** **DOWNLOADING THE IIS MANAGER CONSOLE**
>
> **Users of Windows Server 2003, Windows XP, Vista, and Windows 7 can download a copy of the IIS Manager console to install on their own computers. To find the download, visit *http://www.iis.net/download/IISManager*. After remote users install the program, they can connect to installations of Windows Server 2008 and Windows Server 2008 R2 that include the Web Server (IIS) server role and for which remote management is enabled.**

Figure 6-11 shows the options available for connecting directly to a web application. Remote administrators will be prompted to provide credentials (including a username and password) to make the connection. If the connection is successful, remote administrators see a new object in the left pane of the IIS Manager. These administrators also can name or rename these connections to keep track of multiple connections.

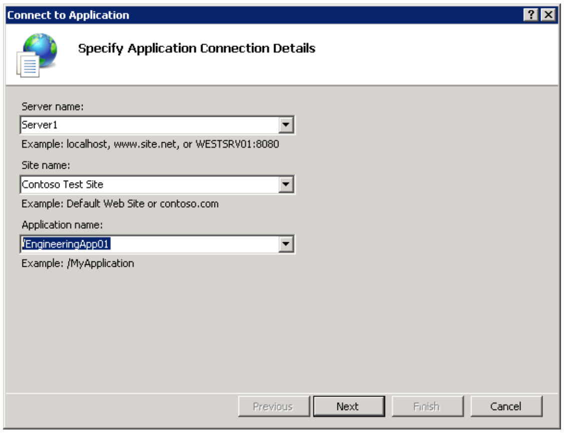

Specify Application Connection Details

Server name:
Server1
Example: localhost, www.site.net, or WESTSRV01:8080

Site name:
Contoso Test Site
Example: Default Web Site or contoso.com

Application name:
/EngineeringApp01
Example: /MyApplication

FIGURE 6-11 Creating a connection to a web application.

The specific items available for management are based on feature delegation settings. Although the same icons might appear, remote administrators will be unable to make or save configuration changes for particular items. For most settings, they can access the configuration page that shows the details, but the controls themselves will be disabled. Figure 6-12 shows an example.

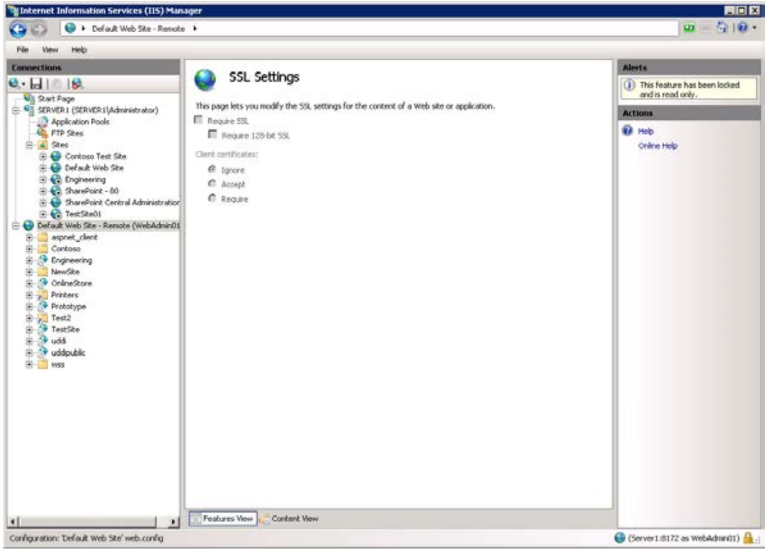

SSL Settings

This page lets you modify the SSL settings for the content of a Web site or application.

□ Require SSL

□ Require 128-bit SSL

Client certificates:

● Ignore

○ Accept

○ Require

Alerts

This feature has been locked and is read only.

Configuration: 'Default Web Site' web.config

(Server1:8172 as WebAdmin01)

FIGURE 6-12 Viewing SSL options that are disabled due to feature delegation settings.

✓ Quick Check

1. **Which steps must you take to enable users to manage IIS remotely by using the IIS Manager console?**
2. **What are the two types of identity credentials that the Management Service can authenticate?**

Quick Check Answers

1. **To enable remote management, you must add the Management Service role service and enable Web Management Service (WMSVC).**
2. **The Management Service can authenticate users by using Windows authentication or IIS Manager credentials.**

Managing Request Handlers

To provide support for various web application technologies, the architecture of IIS allows for enabling and disabling request handlers. *Request handlers* are programs that can process web requests and generate responses that are then returned to clients. Web servers and web applications can be configured with their own sets of request handlers, based on the types of content that must be supported. For example, a web application might be configured to support static content (such as HTML) as well as ASP.NET webpages.

The primary benefit is that web developers can choose the technologies most useful for their tasks. However, there is a drawback from a security standpoint. When IIS is configured with multiple request handlers, the security *attack surface* is increased. A vulnerability in any of the enabled request handlers can result in unauthorized access or related issues. Therefore, it is recommended that systems administrators enable only those request handlers they plan to use. In this section, you learn how to enable and disable request handlers.

REAL WORLD

Anil Desai

Web developers and systems administrators tend to grant far too many permissions on their web servers. Their motivation is simple: it's just easier to provide complete access for all features and settings. That way, it's unlikely that you'll miss some strange requirement. Often, systems administrators don't understand the complexities of web application security, and web developers don't appreciate the importance of minimizing the attack surface of production web servers. The result is less than ideal security and increased risk of unauthorized access. So what's the solution?

The most important aspect of determining ideal security settings is communication. Server administrators should ask web application developers for a list of specific requirements for applications running in production. A preproduction checklist that includes details about intended users, required IIS handlers, authentication requirements, and code access security requirements is a good start. Web developers should understand the importance of minimizing exposure of services and of reducing execution permissions for their applications. To ensure that these goals are being met, both teams can develop tests that validate the configuration from functional and security standpoints.

Overall, web developers and web server administrators tend to have different technical backgrounds and areas of expertise. This is a positive difference as long as both groups understand the benefits of implementing production server security.

Understanding Handler Mappings

When the web server receives a request, IIS uses the definition of handler mappings to determine which request handler to use. A *handler mapping* includes the following information:

- **Verb** HTTP requests include verbs that define the type of request being made. The two most common verbs are GET, which obtains information from the web server, and POST, which can also include information sent from the client browser to the web server.
- **Request extension** Web servers commonly return a wide array of content types. The most common types of information are standard HTML pages and images such as .jpg and .gif files. IIS can use the file extension information from the HTTP request to determine which type of content must be processed. For example, the default file extension for ASP.NET webpages is .aspx. Requests for .aspx pages are mapped automatically to the ASP.NET request handler. Most web development platforms have their own conventions for extensions. It is also possible to create new extensions and provide the appropriate mappings for them.
- **Handler information** The handler mapping includes details related to the specific request handler IIS should call based on the verb and request extension. This information can be provided in different ways, including as a full path to an executable or as the name of a program designed to handle the request.

In addition to specific handler mappings based on these settings, IIS provides the ability to return content by using a default handler. The StaticFile handler mapping is configured to respond to requests that do not map to an existing file. The specific response is based on the settings for the web application. If a default document is specified for the web application or virtual directory, that document is returned if a file is not specified in the URL. For example, a request to *http://Server1.contoso.local/TestSite* results automatically in the return of the default.htm document (if one exists).

If a default document does not exist or the feature is disabled, the StaticFile handler checks whether directory browsing is enabled. If it is, a listing of the contents of the folder is returned to the requester. Finally, if neither of these methods is able to complete the request, the user receives an error stating that the request is forbidden. The complete error message is HTTP Error 403.14, The Web Server Is Configured To Not List The Contents Of This Directory. (See Figure 6-13.)

NOTE **LOCAL VS. REMOTE ERROR MESSAGES**

For security purposes, IIS is configured to provide one type of error message to web users who access the server from the local computer and another type of error message to users who access it remotely. This is done to maintain security; potentially sensitive information is not exposed to remote web browser users, but useful troubleshooting information is still provided to systems administrators and web developers.

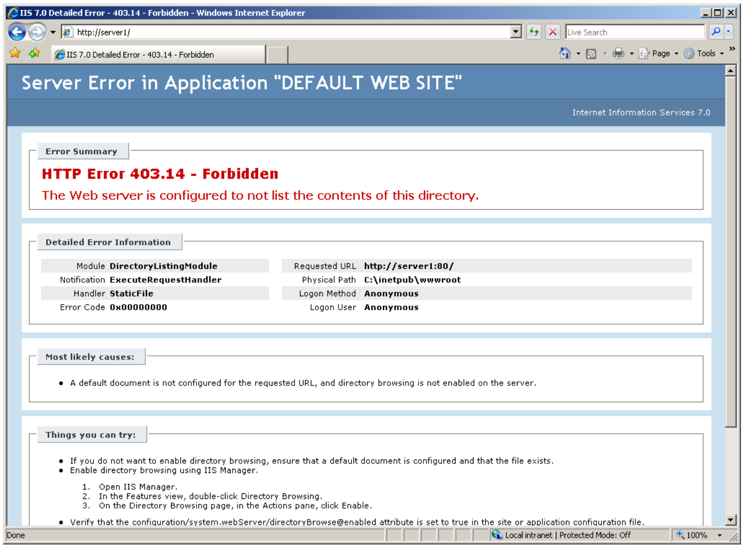
IIS 7.0 Detailed Error - 403.14 - Forbidden - Windows Internet Explorer
http://server1/
Live Search
IIS 7.0 Detailed Error - 403.14 - Forbidden
Page
Tools

Server Error in Application "DEFAULT WEB SITE"

Internet Information Services 7.0

Error Summary

HTTP Error 403.14 - Forbidden

The Web server is configured to not list the contents of this directory.

Detailed Error Information

Module	**DirectoryListingModule**	Requested URL	**http://server1:80/**
Notification	**ExecuteRequestHandler**	Physical Path	**C:\inetpub\wwwroot**
Handler	**StaticFile**	Logon Method	**Anonymous**
Error Code	**0x00000000**	Logon User	**Anonymous**

Most likely causes:

- A default document is not configured for the requested URL, and directory browsing is not enabled on the server.

Things you can try:

- If you do not want to enable directory browsing, ensure that a default document is configured and that the file exists.
- Enable directory browsing using IIS Manager.
  1. Open IIS Manager.
  2. In the Features view, double-click Directory Browsing.
  3. On the Directory Browsing page, in the Actions pane, click Enable.
- Verify that the configuration/system.webServer/directoryBrowse@enabled attribute is set to true in the site or application configuration file.

FIGURE 6-13 A detailed Request Not Found error page.

Configuring Handler Mappings

When you add the Web Server (IIS) role to Windows Server 2008 or Windows Server 2008 R2, a default set of handler mappings is defined for the web server and for the default website. New websites and web applications are also configured with a default set of handler mappings. In addition, when you add role services to the Web Server (IIS) role, additional handler mappings might be added automatically to the configuration.

You can use IIS Manager to configure handler mappings. After you have connected to an installation of IIS, you must choose the level at which you want to configure mappings. You can configure mappings at the following levels:

- Web Server
- Web Sites
- Web Applications
- Virtual Directories
- Web Folders

Child items in the hierarchy automatically inherit handler mappings. For example, a child item automatically inherits the default handler mappings for a new web application from the configuration of the parent website. Settings made at lower levels override the settings from higher levels. This enables a specific web application to support a certain type of file content, such as ASP.NET pages, whereas other applications and the parent website might support only static content.

To view the handler mappings that are configured at a specific level, click the relevant object in the left pane of IIS Manager. Then, double-click Handler Mappings from the Features View in the center pane. Figure 6-14 shows the handler mappings that are defined for a website.

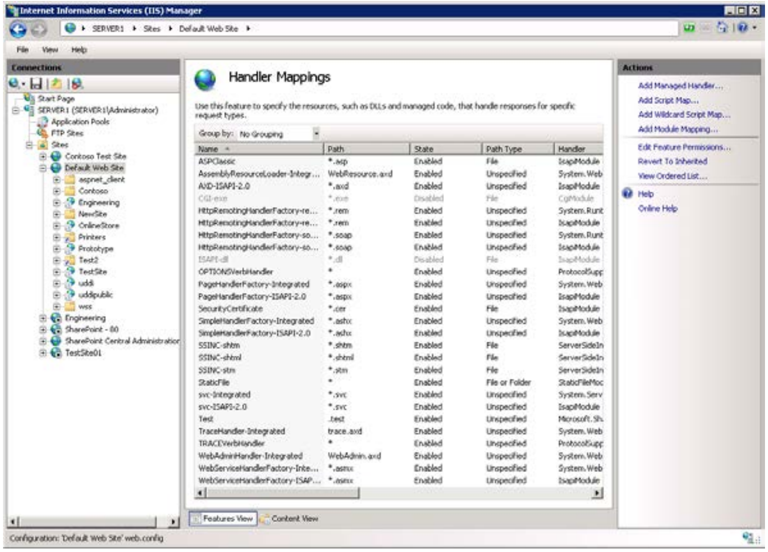

Handler Mappings

Use this feature to specify the resources, such as DLLs and managed code, that handle responses for specific request types.

Group by: No Grouping

Name	Path	State	Path Type	Handler
ASPClassic	*.asp	Enabled	File	IsapiModule
AssemblyResourceLoader-Integr...	WebResource.axd	Enabled	Unspecified	System.Web
AXD-ISAPI-2.0	*.axd	Enabled	Unspecified	IsapiModule
CGI-exe	*.exe	Disabled	File	CgiModule
HttpRemotingHandlerFactory-re...	*.rem	Enabled	Unspecified	System.Runt
HttpRemotingHandlerFactory-re...	*.rem	Enabled	Unspecified	IsapiModule
HttpRemotingHandlerFactory-so...	*.soap	Enabled	Unspecified	System.Runt
HttpRemotingHandlerFactory-so...	*.soap	Enabled	Unspecified	IsapiModule
ISAPI-dll	*.dll	Disabled	File	IsapiModule
OPTIONSVerbHandler	*	Enabled	Unspecified	ProtocolSupp
PageHandlerFactory-Integrated	*.aspx	Enabled	Unspecified	System.Web
PageHandlerFactory-ISAPI-2.0	*.aspx	Enabled	Unspecified	IsapiModule
SecurityCertificate	*.cer	Enabled	File	IsapiModule
SimpleHandlerFactory-Integrated	*.ashx	Enabled	Unspecified	System.Web
SimpleHandlerFactory-ISAPI-2.0	*.ashx	Enabled	Unspecified	IsapiModule
SSINC-shtm	*.shtm	Enabled	File	ServerSideIn
SSINC-shtml	*.shtml	Enabled	File	ServerSideIn
SSINC-stm	*.stm	Enabled	File	ServerSideIn
StaticFile	*	Enabled	File or Folder	StaticFileMoc
svc-Integrated	*.svc	Enabled	Unspecified	System.Serv
svc-ISAPI-2.0	*.svc	Enabled	Unspecified	IsapiModule
Test	.test	Enabled	Unspecified	Microsoft.Sh.
TraceHandler-Integrated	trace.axd	Enabled	Unspecified	System.Web
TRACEVerbHandler	*	Enabled	Unspecified	ProtocolSupp
WebAdminHandler-Integrated	WebAdmin.axd	Enabled	Unspecified	System.Web
WebServiceHandlerFactory-Inte...	*.asmx	Enabled	Unspecified	System.Web
WebServiceHandlerFactory-ISAP...	*.asmx	Enabled	Unspecified	IsapiModule

Configuration: 'Default Web Site' web.config

FIGURE 6-14 Viewing handler mappings for a website.

The display includes information about all the handler mappings defined at the selected level. The name specifies information about the request handler itself. Examples include StaticFile and ASPClassic. Built-in handler mappings have default names, but administrators can provide names for new mappings when they are created. The Path column shows the specific request extensions for which the handler will be used.

The State column specifies whether the handler is enabled or disabled. If the handler is disabled, requests that match the mapping will not be processed. The Handler column specifies details about the program to be called. Finally, the Entry Type specifies whether the handler mapping is inherited from a parent object or is Local (defined directly for this object).

You can use the Group By drop-down list to view handler mappings based on different criteria. These view options make it easy to determine the security attack surface for each component of the web server.

Removing Handler Mappings

To secure your web content, it is a good idea to remove any request handlers that you know will not be required when running in production. To remove a handler mapping, click it and then select the *Remove* command from the Actions pane. After a handler is removed, requests for the types of content that it handled will not be processed. For example, Figure 6-15 shows the result returned to a local web browser when the StaticFile request handler for the web application has been removed. In this case, the request file (default.htm) is present in the web application folder. However, because no request handler is available for the .htm file extension, the request cannot be processed. To the requester, it appears that the file does not exist.

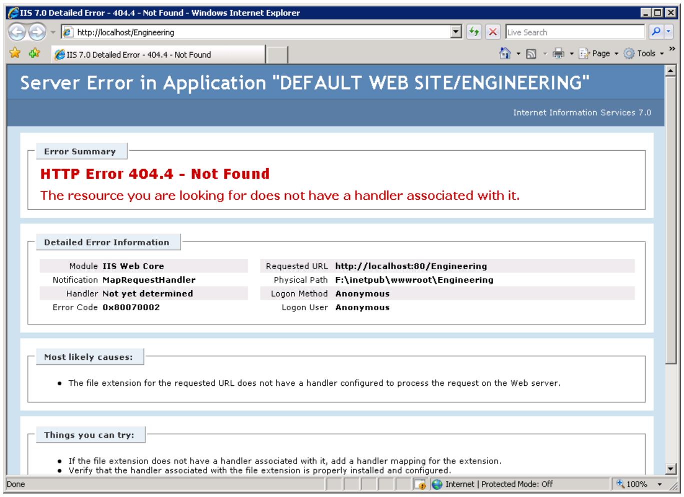
IIS 7.0 Detailed Error - 404.4 - Not Found - Windows Internet Explorer
http://localhost/Engineering
Live Search
IIS 7.0 Detailed Error - 404.4 - Not Found
Page
Tools
Done
Internet | Protected Mode: Off
100%

Server Error in Application "DEFAULT WEB SITE/ENGINEERING"

Internet Information Services 7.0

Error Summary

HTTP Error 404.4 - Not Found

The resource you are looking for does not have a handler associated with it.

Detailed Error Information

Module	**IIS Web Core**	Requested URL	**http://localhost:80/Engineering**
Notification	**MapRequestHandler**	Physical Path	**F:\inetpub\wwwroot\Engineering**
Handler	**Not yet determined**	Logon Method	**Anonymous**
Error Code	**0x80070002**	Logon User	**Anonymous**

Most likely causes:

- The file extension for the requested URL does not have a handler configured to process the request on the Web server.

Things you can try:

- If the file extension does not have a handler associated with it, add a handler mapping for the extension.
- Verify that the handler associated with the file extension is properly installed and configured.

FIGURE 6-15 A detailed request handler error page.

Managing Handler Inheritance

The inheritance feature of handler mapping settings can significantly simplify the administration of servers that host many websites and web applications. In general, configure handler mappings at the highest applicable level. For example, if you are sure that none of the web applications in a specific website must respond to the .soap file extension, you can remove this handler mapping at the level of the website. As mentioned earlier, to increase security, minimize the numbers and types of handlers that are enabled.

By default, it is possible for lower-level objects on the web server to override handler mapping settings from parent objects. In some cases, you might want to prevent some types of requests from being processed on the entire server, regardless of settings for websites and web applications. You do this by locking the configuration of the request handler. To lock the configuration, click the web server object in IIS Manager and then double-click Handler Mappings. Select the handler mapping you wish to lock and then click the *Lock* command in the Actions pane.

It is also possible to restore the handler mappings settings to their default values. To do this, click the *Revert To Parent* command in the Actions pane in IIS Manager. Performing this action restores mappings from the parent object, but it also results in the loss of any locally defined handler mappings.

Adding Handler Mappings

The architecture of IIS enables systems administrators to add new handler mappings based on specific needs. For example, if you want to provide support for a type of file that has a .mypage extension, you can add a handler for this path type. In addition, web developers can create their own programs to manage new types of requests.

To add a handler mapping, select the appropriate object and then double-click Handler Mappings in the Features View in IIS Manager. The Actions pane contains several options for adding new types of request handlers. They are:

- **Add Managed Handler** A managed handler processes requests based on a .NET-based code library. The Type setting enables you to choose from the existing .NET code modules registered on the local server, as shown in Figure 6-16. These types of options all belong to the *System.Web* namespace.

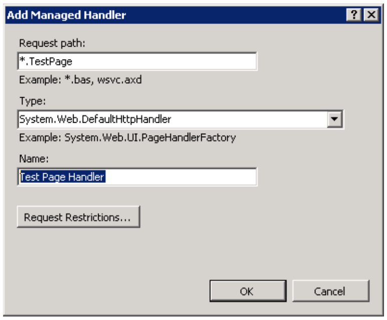

Request path: *.TestPage
Example: *.bas, wsvc.axd

Type: System.Web.DefaultHttpHandler
Example: System.Web.UI.PageHandlerFactory

Name: Test Page Handler

FIGURE 6-16 Adding a managed handler for a website.

- **Add Script Map** Scripting mappings send request processing to a dynamic link library (DLL) or executable (.exe) file type. These types of programs are designed to process request information and generate a response for IIS to send back to the end user.

- **Add Wildcard Script Map** Wildcard script mappings specify a default handler for types of documents that are not managed by other handlers. The Executable path option points to either a .dll or an .exe file designed to handle requests.
- **Add Module Mapping** *Modules* are programs designed to integrate with the IIS request processing pipeline. They can provide a wide range of functions and are included with the default and optional role services that are part of the Web Server (IIS) role. Examples include *FastCGIModule*, for processing scripts based on the Common Gateway Interface (CGI) specification, and *StaticCompressionModule*, which compresses static HTML content to reduce bandwidth usage. In addition to specifying the module to be used for processing, administrators can define an optional executable or .dll file to be used when processing requests, as shown in Figure 6-17.

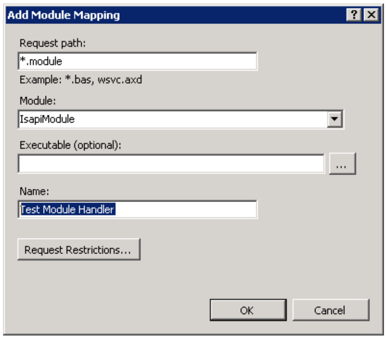

Request path: *.module

Example: *.bas, wsvc.axd

Module: IsapiModule

Executable (optional):

Name: Test Module Handler

FIGURE 6-17 Adding a module mapping to a web application.

When you add a new request handler, you are prompted to provide information about the request path. You can use wildcards, or you can specify a list of specific files. Examples include *.mypage (for responding to a request for any file with this extension) and Config.mypage (for responding to requests for this specific file name). You use the Name setting to help other developers and administrators identify the purpose of the handler mapping.

Configuring Request Restrictions

Besides specifying the paths and file names to which specific request handlers will be mapped, you can further secure IIS through request restrictions. To see the available options, click Request Restrictions in the dialog box when you are adding a mapping. Three tabs organize the request restrictions options: Mapping, Verbs, and Access.

You can use the Mapping tab to specify additional details related to whether files, folders, or both will be included in the mapping. The default setting is for the handler to handle requests automatically for both files and folders. You can choose either files or folders to limit whether the handler will respond to default documents or explicit file requests.

You can use the Verbs tab, shown in Figure 6-18, to specify the HTTP request verbs to which the handler will respond. Although the most common types of verbs are GET and POST, some applications might use other verbs (such as HEAD) to request other details from the web server. By default, all verb types are sent to the request handler. If you want to use different handlers for different verbs, or if you want the handler mapping to apply only to specific types of requests, you can specify this by using the One Of The Following Verbs option.

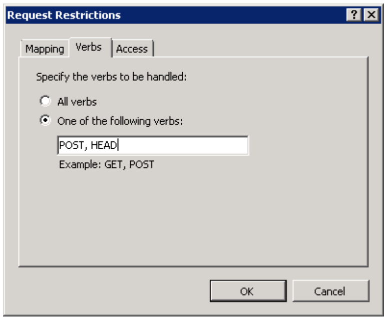

Specify the verbs to be handled:

○ All verbs

● One of the following verbs:

POST, HEAD

Example: GET, POST

FIGURE 6-18 Viewing Verb Request Restrictions options for a handler mapping.

Finally, the Access tab specifies the access permissions that will be granted to the request handler. To improve security, minimize the types of access the handler will have. The default setting is *Script*, which is acceptable for most types of executable handlers. Other options include *None*, *Read*, *Write*, and *Execute*.

Configuring Feature Permissions

Feature permissions specify which types of actions a request handler can take. You can configure these options by double-clicking Handler Mappings and clicking Edit Feature Permissions in the Actions pane, as shown in Figure 6-19.

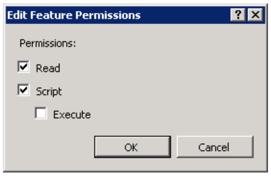

Permissions:

☑ Read

☑ Script

□ Execute

FIGURE 6-19 Configuring Feature Permissions for a request handler.

The three permission options are:

- ***Read*** Enables the handler to read files stored within the file system.
- ***Script*** Enables the handler to perform basic scripting-related tasks on the server.
- ***Execute*** Enables the handler to run executable program code (such as .dll or .exe) files on the computer when processing a request. For *Execute* to be enabled, *Script* permissions must also be assigned.

By default, the *Read* and *Script* feature permissions are enabled for new handler mappings.

Understanding Request Filtering

In Windows Server 2008 R2, you can use the Request Filtering item in Features View of IIS Manager to restrict the kinds of HTTP requests your web server will process. Request Filtering is a security feature that helps you limit the attack surface of your web server.

> ***NOTE*** **REQUEST FILTERING IN WINDOWS SERVER 2008**
>
> **In the RTM version of Windows Server 2008, the Request Filtering item is available in IIS Manager only if you install the Administration Pack for IIS 7.0, which is available on the Microsoft website. However, the full capabilities of this administration feature are still available. You can filter requests in Windows Server 2008 by running Appcmd.exe commands in a command-line window, by editing configuration files directly, or by writing WMI scripts. For more information about Request Filtering in IIS 7.0, visit *http://www.iis.net/ConfigReference/system.webServer/security/requestFiltering*.**

The Request Filtering item in IIS 7.5 provides the following tabs to help you restrict HTTP requests.

- File Name Extensions

  With this tab, you can allow or deny access to web services according to a list of file name extensions specified in HTTP requests.

- Rules

  You can use the Rules tab to configure rules for allowing or denying web access according to various parameters, such as headers and deny strings.

- Hidden Segments

  This tab enables you to define a list of URL segments for which web access will be denied access and excluded from directory listings. (A segment is any part of a URL path between two slash [/] marks.)

- URL

  This tab enables you to allow specific URLs or to deny access to specific sequences within an URL. For example, if you specify "admin/config.xml" as a URL sequence, requests to *http://contoso.com/application/admin/config.xml* will be denied.

- HTTP Verbs

  Use this tab to create a list of HTTP verbs (such as GET, POST, or HEAD) whose access will be specifically allowed or denied.

- Headers

  Use this tab to create a list of HTTP headers whose access will be denied if a specified size limit is surpassed.

- Query Strings

  With this tab, you can create a list of query strings for which web access will be explicitly allowed or denied. An example of a query string is “%3b”, which represents HTTP URL encoding for the apostrophe character and is used in some SQL injection attacks.

PRACTICE Manage IIS Security Settings

This practice walks you through the steps required to manage security for a computer running Windows Server 2008 R2 that has the Web Server (IIS) role installed. Specifically, you learn how to enable remote administration and the effects of configuring handler mappings to increase security. The steps assume that you have already installed the Web Server (IIS) role by using the default options on Server2.contoso.local, and that you are familiar with the process of adding role services.

EXERCISE 1 Configuring and Managing Remote Administration

In this exercise, you use the *IIS Management Service* features to enable a user to connect to the computer. First, you must install the IIS Management Service role service, and then you create a new user based on IIS Manager credentials and configure permissions to access the Default Web Site. Finally, you connect to IIS, using the new user account to verify that the permissions and feature delegation settings are in effect. The final steps can be performed locally on Server2, or you can use another computer, running either Windows 7 or Windows Server 2008 R2, that has the IIS 7 Manager console installed. The steps assume that you perform the tasks locally on Server2.

1. Log on to Server2 as a user who has Administrator permissions.
2. Using Server Manager, add the Management Service role service to the web Server (IIS) server role. You can begin this process by right-clicking the *Web Server (IIS)* node and then selecting Add Role Services. On the Select Role Services page of the Add Role Services Wizard, you can find the Management Service role service within the Management Tools group.
3. After you have added the Management Service role service, open IIS Manager and connect to the local server (Server 2).
4. Click the server object in the left pane and then double-click the Management Service icon in Features View.
5. On the Management Service page, you should see a message stating that the Management Service (WMSVC) is stopped. This is necessary to make configuration changes. Select the Enable Remote Connections option.
6. In the Identity Credentials section, choose Windows Credentials Or IIS Manager Credentials. This enables you to create IIS Manager users later. Leave all other settings at their default values. Note that Management Service responds on port 8172 by default.

7. Start Management Service by clicking Start in the Actions pane. Note that you are unable to modify settings while the service is running.
8. If the Management Service message box appears, click Yes to save the settings before starting the service.
9. Return to Features View by clicking Back in the top toolbar.
10. Double-click IIS Manager Users to view a list of users who have been allowed to access the system. Note that, by default, there are no users in the list.
11. Click Add User in the Actions pane to create a new IIS Manager user. Use the WebAdmin01 username and the 1w3b!admin password. (Always use strong passwords.) Click OK to create the new user and verify that it appears in the list of IIS Manager Users.
12. In the left pane of IIS Manager, expand the *Sites* container and then click the *Default Web Site* object. Next, double-click IIS Manager Permissions in the Management section of Features View.
13. Click Allow User in the Actions pane. For the type of user, select IIS Manager and then type **WebAdmin01** in the text box.

    You can also click Select to select from all the users who have been defined on the server.
14. Click OK.
15. In IIS Manager, click the *Server2* object and then double-click Feature Delegation in the Management section of Features View. In the Group By drop-down list, select Delegation. Note which features in the list are set to Read Only. In later steps, you attempt to change SSL Settings to verify that feature delegation is working.
16. In IIS Manager, click Start Page in the left pane. In the center pane, click the Connect To A Site link.
17. For Server Name, type **server2.contoso.local**. For Site Name, type **Default Web Site**. Click Next.
18. For Username, type **WebAdmin01** and type **1w3b!admin** for Password. Click Next.
19. For the name of the connection, change the name to **Default Web Site – Test** to specify that this is a test connection. Click Finish.

    When the connection is complete, you see a new item called Default Web Site – Test in the left pane of IIS Manager. You can click this connection to administer the site, just as you would with the default local connection. However, note that the new connection shows only the contents of Default Web Site. You will have only the permissions that have been assigned to the WebAdmin01 user.
20. To verify the feature delegation settings, double-click SSL Settings in the IIS section of Features View.

    Note the message in the Actions pane stating that the feature has been locked and set to read only. Verify that you are unable to make changes to these settings.

21. Remove the new connection in IIS Manager by right-clicking it and selecting Remove Connection.
22. When you are finished, close IIS Manager. You can click Yes to save the changes.

EXERCISE 2 Managing Handler Mappings

In this exercise, you learn how to configure and manage handler mappings for a web application. Initially, you verify that content is being presented correctly to web users. You then disable a request handler mapping and verify that the content is no longer accessible. Finally, you revert the handler mappings to their inherited settings to restore access to the content.

1. If you have not done so already, log on to Server2 as a user who has Administrator permissions.
2. Using Windows Explorer, navigate to the *%SystemDrive%*\Inetpub\Wwwroot folder.
3. From the Organize menu in Windows Explorer, select Folder And Search Options.
4. On the View tab of the Folder Options dialog box, clear the Hide Extensions For Known File Types check box and then click OK.
5. Make a copy of the Iisstart.htm file and name it **Iisstart.test**.
6. When you are finished, close Windows Explorer.
7. Open IIS Manager and connect to the local server.
8. In the left pane of IIS Manager, expand the *Sites* node and select Default Web Site. In the Actions pane, click *Browse *:80 (http)*.

   This launches Internet Explorer and connects to the default content for the site. Note that the default document (in this case, Iisstart.htm) is displayed and that the page contains a .png image type.
9. In Internet Explorer, modify the URL to request the iisstart.test page. An example of the full URL would be *http://Server2/iisstart.test*.

   Although the file exists, you receive an HTTP Error 404.3. The error states that no handler is available to process the request.
10. When you are finished, close Internet Explorer.
11. In IIS Manager, with the Default Web Site still selected, double-click Handler Mappings in Features View.

    You see a list of all the default handlers that have been registered on the system.
12. Click the Add Module Mapping link to create a new mapping. For Request Path, type ***.test**. For Module, select StaticFileModule. For Name, type **Test Page Handler**. Leave the other settings at their default values and then click OK to create the mappings.

    This enables the web server to process files that have the .test extension.
13. Open Internet Explorer and navigate to the Iisstart.test page by using the same URL you used in step 9.

This time, you see a blank page, and no error message appears. This indicates that the new handler mapping you created is functioning properly. (By default, the HTML in the file with the .test extension cannot be read without coding a new custom handler. The default.png image does not appear in iisstart.test for this reason.)

14. Close Internet Explorer.
15. In IIS Manager, return to the Handler Mappings section for Default Web Site and then click Revert To Parent in the Actions pane. Click Yes to confirm the changes.

    This restores the default handler mappings and removes the Test Handler Mapping you created in a previous step.
16. When you are finished, close IIS Manager.

Lesson Summary

- When implementing IIS security, consider the overall goals of implementing defense-in-depth best practices and reducing the server's attack surface.
- IIS 7 uses consistent built-in user and group accounts for managing security.
- You can enable remote management of IIS by adding the Management Service role service.
- You can manage remote management capabilities by creating users, assigning permissions, and configuring feature delegation.
- Request handler mappings determine which types of content IIS will allow for a particular component in the hierarchy.

Lesson Review

You can use the following questions to test your knowledge of the information in Lesson 1, "Configuring IIS Security." The questions are also available on the companion CD if you prefer to review them in electronic form.

> ***NOTE*** **ANSWERS**
>
> **Answers to these questions and explanations of why each answer choice is correct or incorrect are located in the "Answers" section at the end of the book.**

1. You are a systems administrator responsible for securing a Windows Server 2008 R2 web server. You have created a new website called Contoso Intranet that will contain seven web applications. One of the application developers has told you that her web application requires a new request handler that is processed by using a .NET library her team created. How can you meet these requirements while also maximizing security for the server?

   - **A.** Add a new managed handler to the Contoso Intranet website.
   - **B.** Add a new managed handler for the specific web application that requires it.
   - **C.** Add a new module mapping to the Contoso Intranet website.
   - **D.** Add a new module mapping for the specific web application that requires it.

2. You are a systems administrator responsible for managing a Windows Server 2008 R2 web server. Recently, your organization set up a new IIS website that users outside your organization will access. Consultants should be able to connect to this website by using IIS Manager. Your organization's security policy prevents you from creating domain accounts or local user accounts for these users. You attempt to use the IIS Manager Permissions feature for the website. However, when you click Allow User, you are able to select only Windows users. How can you resolve this problem?
   - **A.** Verify that Management Service has been started.
   - **B.** Reconfigure the file system permissions for the root folder of the website.
   - **C.** Reconfigure Management Service to enable Windows and IIS Manager credentials.
   - **D.** Verify the authentication settings for the website.

Lesson 2: Controlling Access to Web Services

Web servers are deployed in many configurations. Some servers provide content intended for the general public through the Internet. Others contain web application content intended only for a limited set of users. Web server administration must be able to define which users can connect to a web service and, after users have proven their identities, rules must be in place for determining which content is available to them.

In this lesson, you learn about how you can configure authentication and authorization for protecting web content in IIS. Due to the many security standards and approaches for web services, it is important to understand how to select the most appropriate one for a given scenario. You also learn how you can use features such as IP Address And Domain Restrictions and .NET trust levels to secure your web services further.

After this lesson, you will be able to:

- Describe the authentication options available for IIS web services.
- Configure authentication options for a web server, website, or web application.
- Implement and manage authorization rules to limit access to specific web content.
- Configure server certificates and enable Secure Sockets Layer (SSL) functionality for an IIS server.
- Create and manage IP Address And Domain Restrictions settings to limit access to an IIS web server.
- Configure .NET trust levels based on the needs of specific web applications.

Estimated lesson time: 75 minutes

Managing IIS Authentication

Authentication refers to the process by which a user or computer proves its identity for security purposes. The most familiar method is through a logon or username and an associated password. When working with web servers such as IIS, authentication settings and options determine how users will provide their credentials to access content stored on the web server. IIS provides numerous methods for securing content. By default, content stored in new websites, web applications, and virtual directories will allow access to anonymous users. This means that users are not required to provide any authentication information to retrieve the data. In this section, you learn about the authentication modes supported by IIS and how you can configure them.

Understanding Anonymous Authentication

For many types of web servers, users should be able to access at least a default page or some content without being required to provide authentication information. When you enable the Web Server (IIS) role by using default options, anonymous authentication is enabled for the Default Web Site and its associated web content. Anonymous authentication is designed to provide access to content that should be available to all users who can connect to the web server. An example is the default IIS webpage for Default Web Site. When IIS receives a request for content, it automatically uses a specific identity to attempt to complete the request. By default, anonymous authentication uses the IUSR built-in account. (See Figure 6-20.) As long as this user account has permission to access the content (based on NTFS permissions), the request is processed automatically.

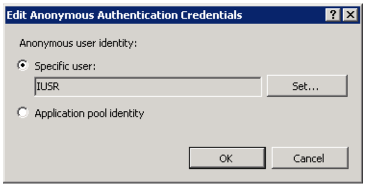

Edit Anonymous Authentication Credentials

Anonymous user identity:

● Specific user:

IUSR

○ Application pool identity

FIGURE 6-20 Editing settings for anonymous authentication credentials.

It is also possible to use the Set option to provide a username and password for a different account. This is useful when you plan to use different NTFS permissions for different web content. Finally, there is an option to use Application Pool Identity. This setting instructs IIS to use the same credentials that are applied to the application pool used by the website or web application.

If all the content on the web server should be available to all users, then no further authentication configuration is required. More commonly, however, you want to restrict access to at least some content on the server. For example, an intranet server might include a web application or virtual directory that is intended for only members of the Human Resources department. To restrict access to content, you can use NTFS permissions. If the credentials configured for the anonymous authentication option are insufficient to access the content, it will not be returned to the user automatically. Generally, you enable one of the other available authentication methods so that authorized users can access the content.

NOTE **SIMPLIFYING CONTENT PROTECTION**

On all web servers, some content exists that should not be accessible to any users. Examples include contents of system folders (such as the Windows system folder) and application source code stored within web content folders. You can use *Deny* NTFS permissions to ensure that users cannot use anonymous credentials to access this content. If you are using multiple accounts for anonymous authentication of different content, it is best to create a group that contains these accounts. You can then deny permission to the group to simplify administration.

Understanding Forms Authentication

A common security approach web developers use is to use standard HTTP forms to transmit logon information. Forms authentication uses an HTTP 302 (Login/Redirect) response to redirect users to a logon page. Generally, the logon page provides users with locations to enter a logon name and password. When this information is submitted back to the logon page, it is validated. Assuming that the credentials are accepted, users are redirected to the content they originally requested. By default, form submissions send data in an unencrypted format. To secure the transmission of logon information, enable encryption through SSL or TLS.

Forms authentication is the most common approach used on the Internet because it does not have any specific web browser requirements. Web developers typically build their own logon pages. Logons are often validated against user account information stored in a relational database (for Internet sites) or against an Active Directory Domain Services (AD DS) domain.

The default settings for forms authentication are designed for use by ASP.NET web applications, and the Forms Authentication option appears only when you have added the ASP.NET role service to the Web Server (IIS) role. You can edit the settings of forms authentication to manage several settings. (See Figure 6-21.) The primary setting is the Login URL. This specifies the name of the webpage to which users are sent when they attempt to access protected content.

After the user has provided authentication information, cookies are sent from the web browser to the web server during each request. This enables the client to prove that it has authenticated with the web server, which is necessary because HTTP is a stateless protocol. The Cookie Settings section enables you to configure how cookies will be used by the site. The Mode options include:

- Do Not Use Cookies
- Use Cookies
- Auto Detect
- Use Device Profile

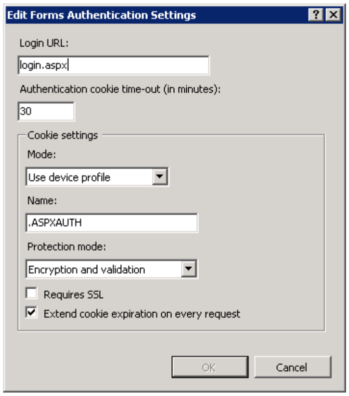

Login URL: login.aspx

Authentication cookie time-out (in minutes): 30

Cookie settings

Mode: Use device profile

Name: .ASPXAUTH

Protection mode: Encryption and validation

☐ Requires SSL

☑ Extend cookie expiration on every request

FIGURE 6-21 Configuring settings for forms authentication.

The most appropriate option will be based on web browser requirements (for example, whether your website requires users to enable support for cookies) and the requirements of the web application or web content.

Understanding Challenge-Based Authentication

Users who access secure websites on the Internet are familiar with the process of providing a username and password to access secured content or to perform actions such as placing online orders. IIS supports three methods of presenting a security challenge to users who are attempting to access web content that has been secured using file system permissions. Each of these methods relies on sending an HTTP 401 Challenge—a standard method that prompts users to provide logon information. These three authentication methods are:

- ***Basic authentication*** Presents an authentication challenge to web users through a standard method that is supported by all web browsers. The main drawback to this method is that information users provide is encoded but not encrypted. This means that, if the information is intercepted, the logon and password details can be obtained easily. To transfer basic authentication information securely, either ensure that your network connections are secure (for example, in a data center environment) or enable encryption using SSL or TLS.
- ***Digest authentication*** Relies on the HTTP 1.1 protocol to provide a secure method of transmitting logon credentials. It does this by using a Windows domain controller to authenticate the user. A potential drawback is that it requires clients' web browsers to support HTTP 1.1. Current versions of most popular browsers support this method, so it is possible to use digest authentication for both Internet and intranet environments.
- ***Windows authentication*** Provides a secure and easy-to-administer authentication option. It relies on the use of either the NTLM or Kerberos authentication protocol to validate users' credentials against a Windows domain or local security database. This method is designed primarily for use in intranet environments, where clients and web servers are members of the same domain. To simplify administration, administrators can use AD DS domain accounts to control access to content.

One important consideration about these challenge-based authentication methods is their interaction with anonymous authentication. If you want to require users to provide logon information before accessing web content, you must disable anonymous authentication. If anonymous authentication remains enabled, content that is not protected by using file system permissions is automatically made available to users without requiring authentication. Another requirement to note is that you cannot enable both forms authentication and challenge-based authentication for the same content.

Understanding ASP.NET Impersonation

Like form authentication, *ASP.NET impersonation* appears as an authentication option when you add the ASP.NET role service. Impersonation is a security method by which an IIS web request is processed by using the security information provided by a specific user account or by the user who is accessing the site. When ASP.NET impersonation is disabled (the default setting), the security context for processing requests is based on the account the web application uses. When you enable impersonation, you can specify a user account for determining the security context. (See Figure 6-22.) To provide the username and password information, click Set.

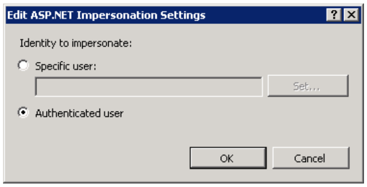
Edit ASP.NET Impersonation Settings

Identity to impersonate:

○ Specific user:

● Authenticated user

FIGURE 6-22 Configuring ASP.NET impersonation settings.

Another option is to configure ASP.NET impersonation to the Authenticated User option. This setting specifies that the security permissions of a user who has been authenticated (using one of the other authentication options) will be used to provide access to content. This setting is useful when you want to use file system permissions that use specific users and groups to decide which content should be protected. When used in this way, it is most appropriate for environments that support relatively small numbers of users, such as department-level intranet web servers.

Understanding Client Certificate Authentication

In addition to the other available authentication options, IIS provides support for using client certificates to validate the identity of a web user. This method requires users to have security certificates installed on their computers. When a request is made for protected content, IIS automatically validates the identity of the client by querying the certificate information. There are three main modes by which client certificates can be used:

- **One-To-One mappings** In this configuration, the web server must contain a copy of the client certificate used by every computer that will access restricted content. The server compares its copy of the certificate with the one presented by the client to validate requests.
- **Many-To-One mappings** It is often impractical to manage certificates for all possible web users on the server. Although this method is slightly less secure, many-to-one mappings are based on the web server performing authentication by using certain information found in the client certificate. A common example is validating the organization information in the certificate to ensure that the user is coming from a trusted company.
- **Active Directory mappings** Active Directory Certificate Services can simplify the creation and management of client certificates. To enable this method, organizations must first set up their own certificate-based infrastructure.

Because of the certificate requirements for client certificate authentication, this method is most often used in environments in which systems administrators have control over end users' computers. It is impractical to require certificates for publicly accessible Internet websites and applications.

Understanding Authentication Requirements

Handlers and modules manage IIS authentication. The specific authentication options available for a web server are based on the Web Server (IIS) role services installed. The list of available role services includes:

- Basic Authentication
- Windows Authentication
- Digest Authentication
- Client Certificate Mapping Authentication
- IIS Client Certificate Mapping Authentication
- ASP.NET (for forms authentication and ASP.NET impersonation)

To add or remove a security-related role service, open Server Manager, expand the Roles section, right-click Web Server (IIS), and then select either Add Role Services or Remove Role Services. (See Figure 6-23.) Because role services affect the available authentication options for the entire web server, determine the requirements of all the web applications and web content on your server.

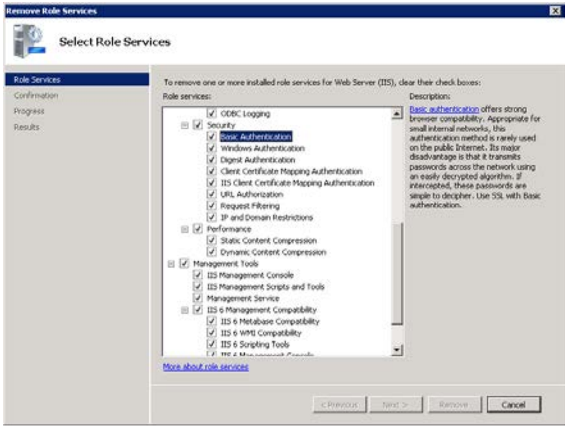

To remove one or more installed role services for Web Server (IIS), clear their check boxes:

Role services:

- ☑ ODBC Logging
- ☑ Security
  - ☑ Basic Authentication
  - ☑ Windows Authentication
  - ☑ Digest Authentication
  - ☑ Client Certificate Mapping Authentication
  - ☑ IIS Client Certificate Mapping Authentication
  - ☑ URL Authorization
  - ☑ Request Filtering
  - ☑ IP and Domain Restrictions
- ☑ Performance
  - ☑ Static Content Compression
  - ☑ Dynamic Content Compression
- ☑ Management Tools
  - ☑ IIS Management Console
  - ☑ IIS Management Scripts and Tools
  - ☑ Management Service
  - ☑ IIS 6 Management Compatibility
    - ☑ IIS 6 Metabase Compatibility
    - ☑ IIS 6 WMI Compatibility
    - ☑ IIS 6 Scripting Tools
    - ☑ IIS 6 Management Console

Description:

Basic authentication offers strong browser compatibility. Appropriate for small internal networks, this authentication method is rarely used on the public Internet. Its major disadvantage is that it transmits passwords across the network using an easily decrypted algorithm. If intercepted, these passwords are simple to decipher. Use SSL with Basic authentication.

FIGURE 6-23 Viewing installed authentication-related role services.

In addition to role service settings, each of the authentication methods has specific module requirements, as shown in Table 6-1. For more information about managing modules, see the "Managing Request Handlers" section earlier in this chapter.

TABLE 6-1 IIS Authentication Methods and Their Requirements

AUTHENTICATION METHODS	REQUIRED MODULE(S)
Anonymous	AnonymousAuthModule
ASP.NET Impersonation	ManagedEngine
Basic	BasicAuthModule TokenCacheModule
Client Certificates	iisClientCertificateMappingModule
Client Certificates (Active Directory Mapping)	CertificateMappingAuthenticationModule
Digest	DigestAuthModule
Forms	FormsAuthenticationModule
Windows	WindowsAuthenticationModule

Configuring Authentication Settings

IIS enables you to define configuration settings by using the web object hierarchy. Authentication settings can be configured for objects at the following levels:

- Web server
- Websites
- Web applications
- Virtual directories
- Physical folders and individual files

Authentication settings that are defined at higher levels (such as for a web application) apply automatically to lower-level objects. This method makes it easier to manage settings for multiple websites, web applications, and their related content.

To configure authentication settings by using IIS Manager, select the appropriate object in the left pane and then double-click Authentication in Features View. Figure 6-24 shows the default authentication options for the Default Web Site object.

The default display shows a complete list of the available authentication options, grouped by the response type used. Each method can be enabled or disabled by selecting the item and using the *Enable* or *Disable* commands in the Actions pane. In addition, some authentication options provide additional commands for managing settings. By default, when you enable or disable an authentication option, the setting applies to all lower-level objects and content in the IIS hierarchy. You can override this behavior by explicitly enabling or disabling authentication methods at lower levels.

To verify your authentication-related settings, you should always test access to content by using a web browser. In some cases, it might be necessary to use a second computer to ensure that authentication is working properly. For example, if you are already connected to a computer running Windows Server 2008 R2 as a member of the administrators group and you want to test Windows Authentication, you should attempt to connect from another computer in the environment to help prevent automatic authentication from affecting your test results.

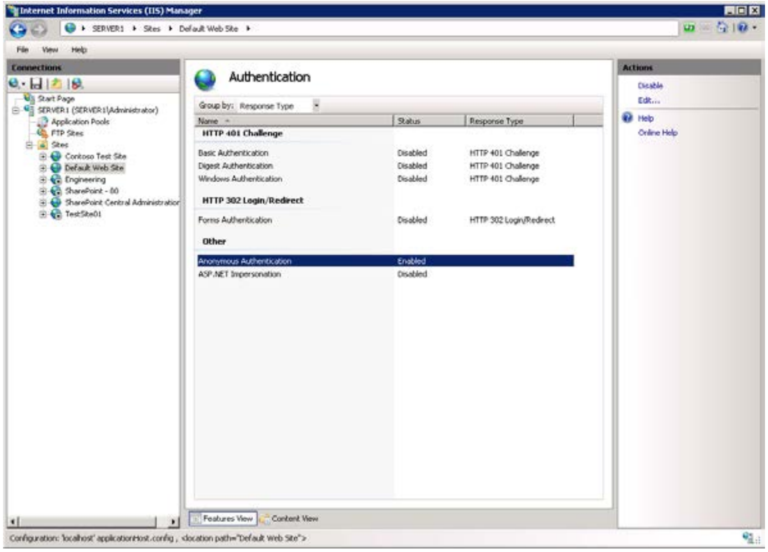

Group by: Response Type

Name	Status	Response Type
HTTP 401 Challenge		
Basic Authentication	Disabled	HTTP 401 Challenge
Digest Authentication	Disabled	HTTP 401 Challenge
Windows Authentication	Disabled	HTTP 401 Challenge
HTTP 302 Login/Redirect		
Forms Authentication	Disabled	HTTP 302 Login/Redirect
Other		
Anonymous Authentication	Enabled	
ASP.NET Impersonation	Disabled	

Configuration: 'localhost' applicationHost.config , <location path="Default Web Site">

FIGURE 6-24 Viewing authentication options for Default Web Site by using IIS Manager.

Managing URL Authorization Rules

Authorization is a method by which systems administrators can determine which resources and content are available to specific users. Authorization relies on authentication to validate the identity of a user. After the identity has been proven, authorization rules determine which actions a user or computer can perform. IIS provides methods of securing different types of content by using URL-based authorization. Because web content is generally requested by using a URL that includes a full path to the content being requested, you can configure authorization settings easily by using IIS Manager.

Creating URL Authorization Rules

To enable URL authorization, the UrlAuthorizationModule must be enabled. Authorization rules can be configured at the level of the web server for specific websites, for specific web applications, and for specific files (based on a complete URL path). *URL authorization rules* use inheritance so that lower-level objects inherit authorization settings from their parent objects (unless they are specifically overridden).

To configure authorization settings, select the appropriate object in the left pane of IIS Manager and then select Authorization Rules in Features View. Figure 6-25 shows an example of multiple rules configured for a website.

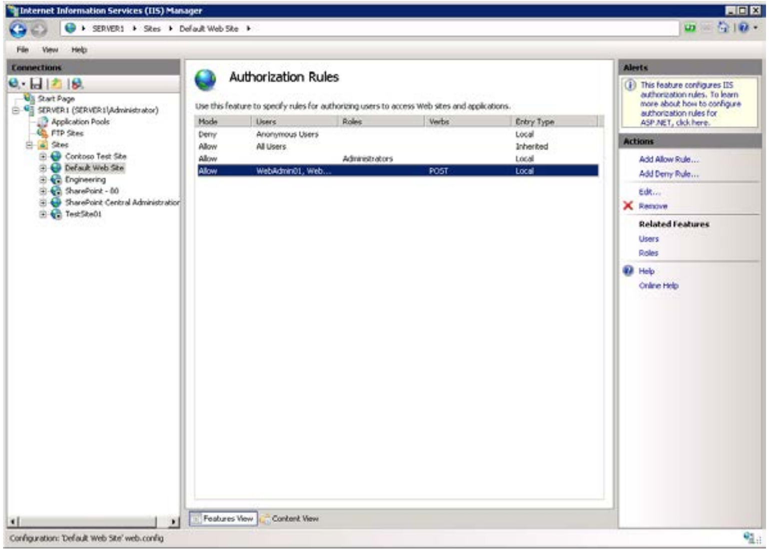

Authorization Rules

Use this feature to specify rules for authorizing users to access Web sites and applications.

Mode	Users	Roles	Verbs	Entry Type
Deny	Anonymous Users			Local
Allow	All Users			Inherited
Allow		Administrators		Local
Allow	WebAdmin01, Web...		POST	Local

This feature configures IIS authorization rules. To learn more about how to configure authorization rules for ASP.NET, click here.

Configuration: 'Default Web Site' web.config

FIGURE 6-25 Viewing authorization rules for a website.

There are two types of rules: Allow and Deny. You can create new rules by using the *Add Allow Rule* and *Add Deny Rule* commands in the Actions pane. The available options for both types of rules are the same, as shown in Figure 6-26. When creating a new rule, the main setting is determining the users to which the rule applies. The options are:

- All Users
- All Anonymous Users
- Specific Roles Or User Groups
- Specific Users

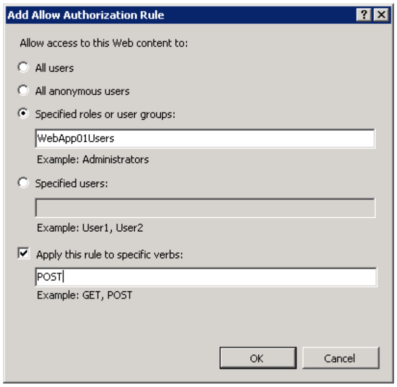

Allow access to this Web content to:

○ All users

○ All anonymous users

● Specified roles or user groups:

WebApp01Users

Example: Administrators

○ Specified users:

Example: User1, User2

☑ Apply this rule to specific verbs:

POST

Example: GET, POST

FIGURE 6-26 Creating a new Allow Rule for a web application.

When you choose to specify users or groups to which the rule applies, you can type the appropriate names in a comma-separated list. The specific users and groups are defined using .NET role providers. This is a standard feature that is available to ASP.NET web developers. Developers can create their own roles and user accounts and can define permissions within their applications. Generally, information about users and roles is stored in a relational database or relies on a directory service such as Active Directory.

In addition to user and role selections, you can further configure an authorization rule based on specific HTTP verbs. For example, if you want to apply a rule only for *POST* commands (which typically send information from a web browser to a web server), add only the POST verb to the rule.

Managing Rule Inheritance

As mentioned earlier in this section, authorization rules are inherited automatically by lower-level objects. This is useful when your website and web content is organized hierarchically based on intended users or groups. The Entry Type column shows whether a rule has been inherited from a higher level or has been defined locally. IIS Manager automatically prevents you from creating duplicate rules. You can remove rules at any level, including both *Inherited* and *Local* entry types.

Configuring Server Certificates

One of the many challenges related to security is verifying the identity of a web server, and, after you are reasonably sure that the server can be trusted, you must protect communications between the web client and the web server. On many networks, and especially on the Internet, providing secure communications for sensitive data is a key concern. Server certificates are designed to provide added security for web services. IIS provides built-in support for creating and managing server certificates and for enabling encrypted communications. In this section, you learn how to configure and enable these options.

Understanding Server Certificates

Server certificates are a method by which a web server can prove its identity to the clients attempting to access it. The general approach to providing this functionality is by a hierarchy of trust authorities. The party that issues a server certificate is known as a *certificate authority (CA)*. On the Internet, numerous third-party organizations are available for validating servers and generating certificates. Assuming that users trust these third parties, they should be able to extend the trust to validated websites. Organizations can also serve as their own CA for internal servers. This enables systems administrators to validate and approve new server deployments by using a secure mechanism.

The general process for obtaining a server certificate involves three major steps:

- **Generating a certificate request** The request is created on a web server, which produces a text file containing the information about the request in an encrypted format. The certificate request identifies the web server uniquely.

- **Submitting the certificate request to a CA** The certificate request is submitted to a CA (generally by using a secure website or email). The CA then verifies the information in the request and creates a trusted server certificate.
- **Obtaining and installing a certificate on the web server** The CA returns a certificate to the requester, usually in the form of a small text file. This file can then be imported into the web server configuration to enable secure communications.

NOTE **CLIENT CERTIFICATES VS. SERVER CERTIFICATES**

Certificate-based technology can be used with a web server by several methods. Use client-based certificates to verify access to a web server by validating clients. In this case, the client holds a certificate the server can validate. You learned about this method earlier in this lesson. Server-side certificates are installed on web server computers to prove their identity to web clients and to enable encrypted communications. Client-side certificates are generally used in intranet or extranet environments, whereas server-side certificates are common for securing all types of web servers.

Creating an Internet Certificate Request

Use IIS Manager to obtain a certificate for use on an IIS web server. To begin the process, connect to a web server running Windows Server 2008 or Windows Server 2008 R2 and double-click Server Certificates in Features View, as shown in Figure 6-27. Note that certificate requests are generated at the level of the web server and not for other objects, such as websites or web applications.

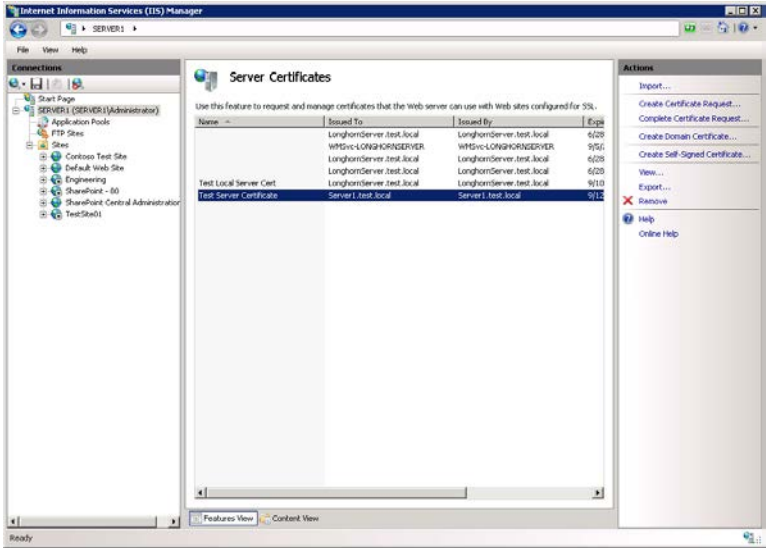

Use this feature to request and manage certificates that the Web server can use with Web sites configured for SSL.

Name	Issued To	Issued By	Expi
	LonghornServer.test.local	LonghornServer.test.local	6/28
	WMSvc-LONGHORNSERVER	WMSvc-LONGHORNSERVER	9/5/
	LonghornServer.test.local	LonghornServer.test.local	6/28
	LonghornServer.test.local	LonghornServer.test.local	6/28
Test Local Server Cert	LonghornServer.test.local	LonghornServer.test.local	9/10
Test Server Certificate	Server1.test.local	Server1.test.local	9/12

FIGURE 6-27 Viewing Server Certificate options for an IIS web server.

Depending on the configuration of the local server, some certificates might already be included in the default configuration. The Actions pane provides commands for creating new certificates.

To begin the certificate request process, click Create Certificate Request. As shown in Figure 6-28, you will be required to provide information about the requesting organization. The CA uses this information to determine whether to issue the certificate. Therefore, it is important for information to be exact. For example, the Organization field should include the complete legal name of the requesting company. The Common Name field generally defines the domain name that will be used with the certificate.

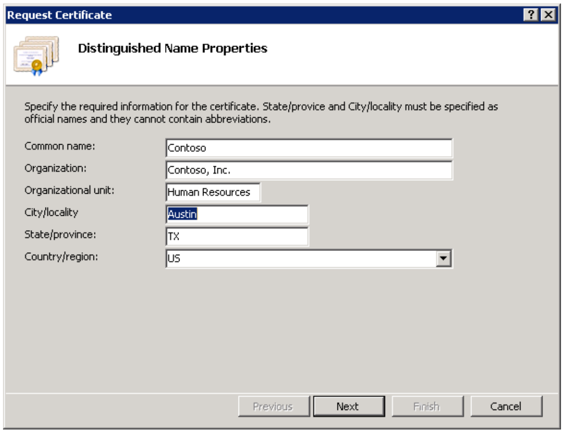
Request Certificate

Distinguished Name Properties

Specify the required information for the certificate. State/provice and City/locality must be specified as official names and they cannot contain abbreviations.

Common name: Contoso
Organization: Contoso, Inc.
Organizational unit: Human Resources
City/locality Austin
State/province: TX
Country/region: US

FIGURE 6-28 The Distinguished Name Properties page.

The second step of the certificate request process requires you to choose the cryptographic method to secure the certificate request. (See Figure 6-29.) The Cryptographic Service Provider setting should use a method accepted by the certificate authority. (The default option, Microsoft RSA SChannel Cryptographic Provider, is accepted by most third-party CAs.) The Bit Length setting indicates the strength of the encryption. Larger values take more time to process (due to computational overhead) but provide added security.

The final step of the process involves storing the certificate request to a file. Here you can provide a fully qualified path and file name into which the request will be stored. The request itself will be stored in a text file that contains encrypted information.

The next step of the process involves submitting the certificate request to a CA. Generally, the issuer's website will request that you either upload the certificate request or copy and paste the contents into a secure website. The issuer will also require additional information, such as details about your organization and payment information.

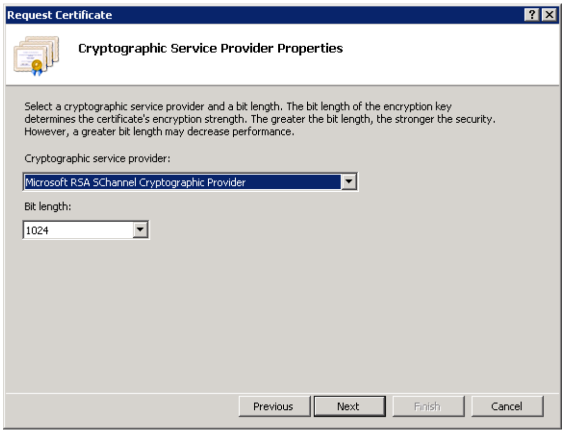

Request Certificate

Cryptographic Service Provider Properties

Select a cryptographic service provider and a bit length. The bit length of the encryption key determines the certificate's encryption strength. The greater the bit length, the stronger the security. However, a greater bit length may decrease performance.

Cryptographic service provider: Microsoft RSA SChannel Cryptographic Provider

Bit length: 1024

FIGURE 6-29 The Cryptographic Service Provider Properties page.

Completing an Internet Certificate Request

The amount of time a public third-party CA can take to process a request varies. After the request has been processed and approved, the CA sends a response by email or through its website. You can then store this response in a text file and provide it to IIS to complete the process by selecting the appropriate request in the Server Certificates feature view and then clicking the *Complete Certificate Request* command in the Actions pane. You will be asked to specify the path and file name of the response along with a friendly name for administration purposes, as shown in Figure 6-30. The convention is to use a file name with a .cer extension for the response; however, any type of standard text file will work.

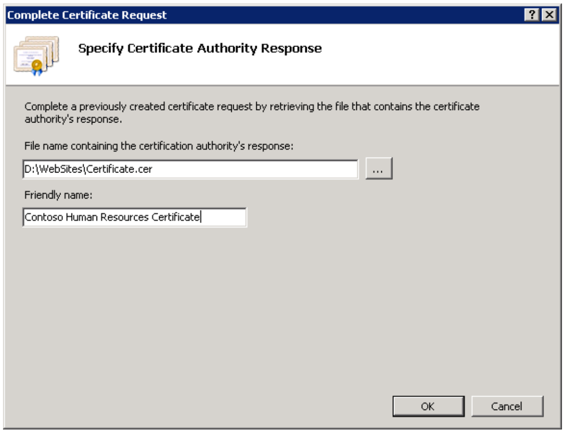

Complete Certificate Request

Specify Certificate Authority Response

Complete a previously created certificate request by retrieving the file that contains the certificate authority's response.

File name containing the certification authority's response: D:\WebSites\Certificate.cer

Friendly name: Contoso Human Resources Certificate

FIGURE 6-30 Completing the certificate request process.

Assuming that the certificate request matches the response, the certificate is imported into the configuration of IIS and is ready for use.

Creating Other Certificate Types

In addition to the standard certificate request process, you can use two other commands to create certificates. These commands are also available in the Actions pane in the properties of the Server Certificates feature. The *Create Domain Certificate* option generates a request to an internal certificate authority. This is commonly used in organizations that have their own certificate services infrastructure. Instead of sending the request to a third-party CA, the request is designed to be sent to an internal server. Figure 6-31 shows the available options. The Specify Online Certificate Authority text box accepts the path and name of an internal CA server. The Friendly Name can identify the purpose of the certificate.

> ***MORE INFO*** **ACTIVE DIRECTORY CERTIFICATE SERVICES**
>
> **Windows Server 2008 and Windows Server 2008 R2 include the Active Directory Certificate Services server role, which enables administrators to create their own certificate-based security infrastructure. The details of implementing these services are outside the scope of this book and the scope of the 70-643 exam. For more information about configuring certificate services, see *http://technet.microsoft.com/en-us/windowsserver/dd448615.aspx*.**

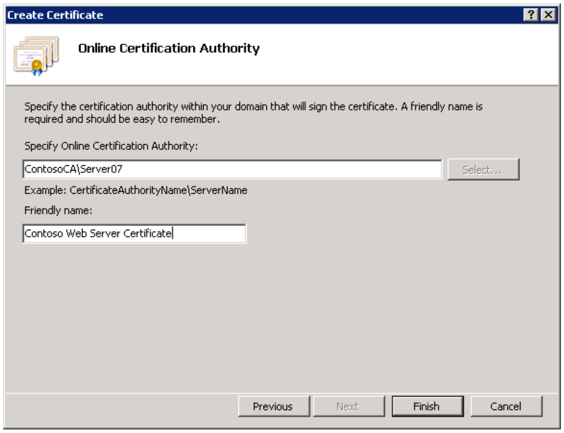

Specify the certification authority within your domain that will sign the certificate. A friendly name is required and should be easy to remember.

Specify Online Certification Authority: ContosoCA\Server07

Example: CertificateAuthorityName\ServerName

Friendly name: Contoso Web Server Certificate

FIGURE 6-31 Specifying Online Certificate Authority settings for a Domain Certificate.

Creating a Self-Signed Certificate

The certificate creation and management process can require several steps, and usually requires an added cost for obtaining a certificate from a trusted third-party CA. Although these steps are necessary to ensure security in a production environment, an easier method is preferable for development and test environments. *Self-signed certificates* can test certificate functionality by creating a local certificate. By avoiding the CA process, it is easy to create these certificates by using the *Create Self-Signed Certificate* command in the Actions pane. Figure 6-32 shows the dialog box.

Unlike other certificate types, it is not necessary to provide organizational information for the certificate because the certificate itself is created immediately on the local computer. The primary drawback of self-signed certificates is that users who access the web server using a secure connection receive a warning that the certificate has not been issued by a third party. (See Figure 6-33.) Whereas this is generally not a problem in test environments, it prevents the use of self-signed certificates for production web servers.

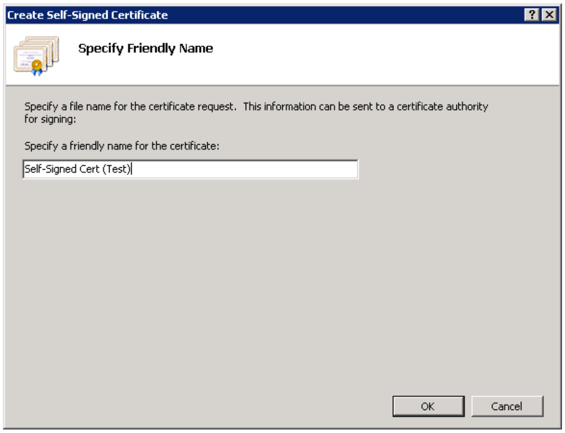

Specify a file name for the certificate request. This information can be sent to a certificate authority for signing:

Specify a friendly name for the certificate:

Self-Signed Cert (Test)

FIGURE 6-32 Creating a self-signed certificate.

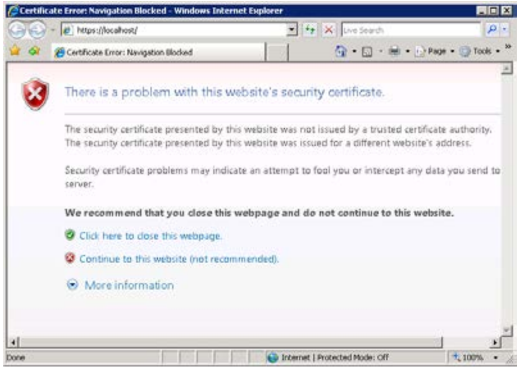
Certificate Error: Navigation Blocked - Windows Internet Explorer
https://localhost/
Live Search
Certificate Error: Navigation Blocked
Page
Tools

There is a problem with this website's security certificate.

The security certificate presented by this website was not issued by a trusted certificate authority.
The security certificate presented by this website was issued for a different website's address.

Security certificate problems may indicate an attempt to fool you or intercept any data you send to server.

We recommend that you close this webpage and do not continue to this website.

Click here to close this webpage.

Continue to this website (not recommended).

More information

FIGURE 6-33 Viewing a certificate-related error when accessing a server that is using a self-signed certificate.

Viewing Certificate Details

The contents of a server certificate include several details and properties. To view this information, double-click an item in the Server Certificates list for a web server. The Certificate dialog box, shown in Figure 6-34, provides information about the server certificate. The General tab displays details about the issuer of the certificate. For an Internet-based certificate, this is the name of the trusted third party that issued it. Additionally, certificates have a range of valid dates.

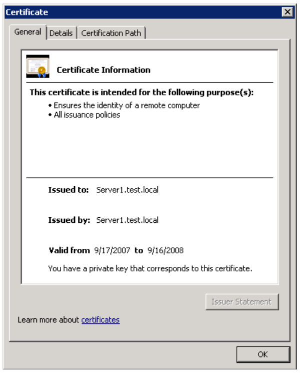

Certificate Information

This certificate is intended for the following purpose(s):

- Ensures the identity of a remote computer
- All issuance policies

Issued to: Server1.test.local

Issued by: Server1.test.local

Valid from 9/17/2007 to 9/16/2008

You have a private key that corresponds to this certificate.

Learn more about certificates

FIGURE 6-34 Viewing general information for a server certificate.

The Details tab displays additional properties of the certificate, including the encryption method. The Certification Path tab shows the entire trust hierarchy for the certificate. In environments that have multiple levels of CAs, this is useful for tracking all the trust relationships that are used. For the certificate to be considered valid, all the levels must be trusted. (See Figure 6-35.)

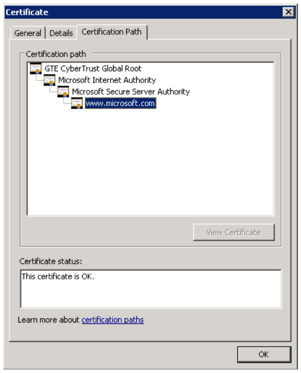

Certification path

- GTE CyberTrust Global Root
  - Microsoft Internet Authority
    - Microsoft Secure Server Authority
      - www.microsoft.com

Certificate status:

This certificate is OK.

Learn more about certification paths

FIGURE 6-35 Viewing certificate information for a public website by using Internet Explorer.

Importing and Exporting Certificates

After a certificate has been installed on a web server, you might need to export it to a file. You can use IIS Manager to do this by right-clicking the certificate and choosing the *Export* command. You can then provide an export location and file name for the file along with a password to protect the certificate from being installed by unauthorized users. (See Figure 6-36.) By default, exported certificate files use the .pfx extension. However, you can use any other extension. The contents of the exported certificate are encrypted and protected with the password you provide.

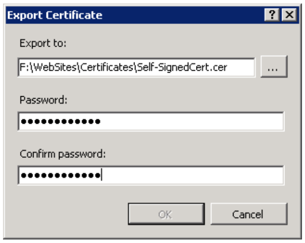
Export Certificate

Export to:

F:\WebSites\Certificates\Self-SignedCert.cer

Password:

••••••••••••

Confirm password:

••••••••••••

FIGURE 6-36 Exporting a server certificate by using IIS Manager.

To import a certificate, click the *Import* command in the Actions pane. You are prompted to provide the file system location of the exported certificate file along with the password to open it. In addition, you can choose whether you want to allow the certificate to be exported in the future.

Enabling Secure Sockets Layer

After you have added a server certificate to an IIS web server, you can enable SSL-based connections. SSL-based connections rely on certificates to validate the identity of the web server. After the identity has been proven, users can create a secure connection by using the HTTP Secure (HTTPS) protocol. By default, HTTPS connections use TCP port 443 for communications. To modify the details or to enable HTTPS for a website, you must configure the site bindings for a website. (For complete details about configuring site bindings, see Chapter 5.)

You can also require SSL-enabled connections for specific websites by using IIS Manager. To do this, select a website, a web application, or a folder and then double-click SSL Settings in the Features view. In Windows Server 2008 R2, websites have only an HTTP binding by default, so you first must add an HTTPS binding before you can configure SSL. (Adding an HTTPS binding requires you to specify a server SSL certificate.) Figure 6-37 shows the available SSL settings after you have added an HTTPS binding to a website. The check boxes enable you to specify whether SSL is required to access this content. If the option is enabled, standard HTTP connections are not enabled. Optionally, you can specify whether client certificates will be ignored, accepted, or required.

Overall, server certificates and SSL provide a standard method of protecting web-based connections and web server content. Support for server certificates and SSL is often expected for all types of web servers that contain sensitive information.

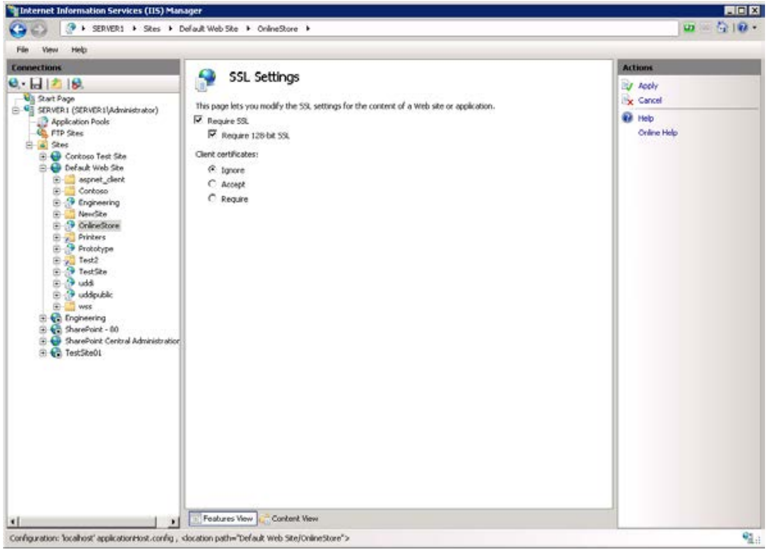

SSL Settings

This page lets you modify the SSL settings for the content of a Web site or application.

☑ Require SSL

☑ Require 128-bit SSL

Client certificates:

● Ignore

○ Accept

○ Require

Configuration: 'localhost' applicationHost.config , <location path="Default Web Site/OnlineStore">

FIGURE 6-37 Configuring SSL settings for a web application.

Configuring IP Address and Domain Restrictions

Although some web servers are configured to provide public access to all content, it's also common to need to restrict access to only specific groups of users. By default, IIS is configured to accept requests on all connections based on site binding settings such as IP address and TCP port. Systems administrators can further restrict access to websites by responding only to requests that originate from specific IP addresses or domains using IIS Manager.

The first step is to select the level at which you want to assign the restrictions. The IP Address and Domain Restrictions feature is available at the server, site, web application, virtual directory, and folder levels. In general, assign restrictions at the highest level for which the settings will apply. For example, if all the web applications in a particular site should respond to requests from a single domain only, configure the request settings at the site level. By default, IIS does not include any restrictions. To configure request settings, select the appropriate object in the left pane of IIS Manager and then double-click IP Address and Domain Restrictions in Features View. (This feature becomes available when you add the IP and Domain Restrictions component of the Security role service.) Figure 6-38 provides an example of the settings.

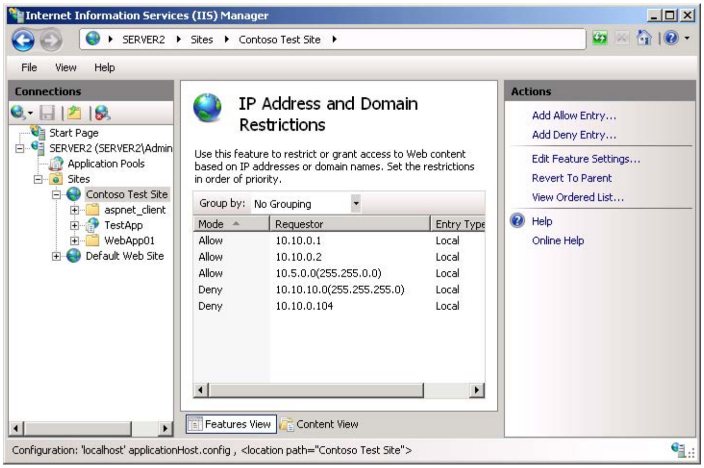

IP Address and Domain Restrictions

Use this feature to restrict or grant access to Web content based on IP addresses or domain names. Set the restrictions in order of priority.

Group by: No Grouping

Mode	Requestor	Entry Type
Allow	10.10.0.1	Local
Allow	10.10.0.2	Local
Allow	10.5.0.0(255.255.0.0)	Local
Deny	10.10.10.0(255.255.255.0)	Local
Deny	10.10.0.104	Local

Configuration: 'localhost' applicationHost.config , <location path="Contoso Test Site">

FIGURE 6-38 Configuring IPv4 Address and Domain Restrictions for a website.

Adding Allow and Deny Entries

You can add two main types of entries to the IP Address and Domain Restrictions configuration. Allow entries specify which IP addresses can access web content; Deny entries define which addresses cannot access the content. When configuring *IP address restrictions*, you can specify either a single IP address or a range of IP addresses. (See Figure 6-39.) When specifying a range, you can enter the initial IP address and the subnet mask. This determines the range of addresses that will be allowed or denied. It is possible to exclude specific addresses or ranges by using additional allow or deny rules. Overall, however, try to keep the configuration simple to make administration and management as easy as possible.

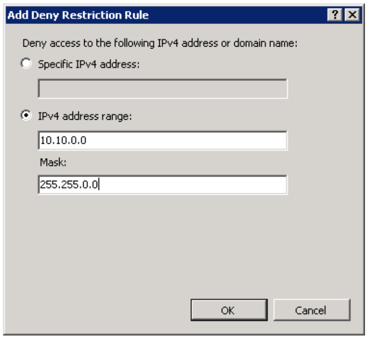

Deny access to the following IPv4 address or domain name:

○ Specific IPv4 address:

● IPv4 address range:

10.10.0.0

Mask:

255.255.0.0

FIGURE 6-39 Adding a Deny entry IP address restriction for a website.

The single address option is useful if only a few users require access to the site or if only a few other servers require access to the content. This is common in environments that support distributed server-side web applications that are not designed for direct user access. IP address ranges are more appropriate when groups of users and computers should have access to the environment. For example, if all the users in the Human Resources department are located on the same subnet, that subnet can be allowed while other subnets are denied.

When evaluating connection rules, IIS evaluates all allow and deny rules to determine whether an address has access. Deny rules take precedence over allow rules. If users are denied access to a site, they see a screen similar to the one shown in Figure 6-40.

An additional setting defines the default behavior for any IP addresses that are not explicitly added to either the Allow or Deny list. By default, IIS allows access automatically from these addresses. To change the setting, click Edit Feature Settings in the Actions pane and choose Deny for the Access For Unspecified Clients setting. (See Figure 6-41.)

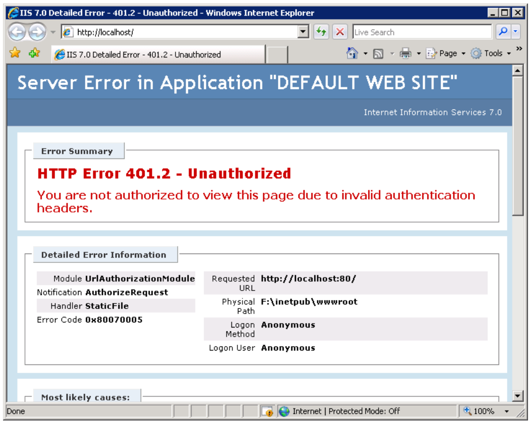
IIS 7.0 Detailed Error - 401.2 - Unauthorized - Windows Internet Explorer

http://localhost/

Server Error in Application "DEFAULT WEB SITE"

Internet Information Services 7.0

Error Summary

HTTP Error 401.2 - Unauthorized

You are not authorized to view this page due to invalid authentication headers.

Detailed Error Information

Module	UrlAuthorizationModule	Requested URL	http://localhost:80/
Notification	AuthorizeRequest	Physical Path	F:\inetpub\wwwroot
Handler	StaticFile	Logon Method	Anonymous
Error Code	0x80070005	Logon User	Anonymous

Most likely causes:

Done

Internet | Protected Mode: Off

FIGURE 6-40 An error message returned to a user, based on site restriction settings.

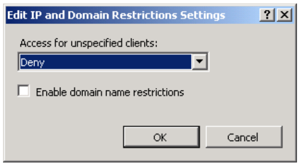
Edit IP and Domain Restrictions Settings

Access for unspecified clients:

Deny

□ Enable domain name restrictions

FIGURE 6-41 Configuring feature settings for IPv4 Address and Domain Restrictions.

Adding Domain Restrictions

Managing access to web services by using IP addresses is useful when the list of incoming clients is well known. This is typical of intranet and internal network environments in which network administrators can configure and manage IP address ranges. In other types of web server scenarios—such as public web servers or extranets—managing IP address ranges can be time consuming and impractical.

An alternative to using IP address-based restrictions is specifying allow and deny settings by using domain name restrictions. This method depends on a Domain Name System (DNS) reverse lookup operation. Whenever a user attempts to connect to IIS, the web server performs a reverse DNS lookup to resolve the requester's IP address to a domain name. IIS then uses the domain name to determine whether the user should have access. Domain-based restrictions are disabled by default because this feature can decrease server performance significantly. Every incoming request must be resolved, adding overhead to request processing. In addition, this can place significant load on the DNS server infrastructure. From a management standpoint, however, this feature can sometimes be useful (especially in low-volume scenarios).

To enable domain name restrictions, double-click the IP Address And Domain Restrictions feature for a portion of the website and then click Edit Feature Settings in the Actions pane. As shown in Figure 6-41, you can select the Enable Domain Name Restrictions check box to enable this feature. Figure 6-42 shows the confirmation warning when you enable this feature.

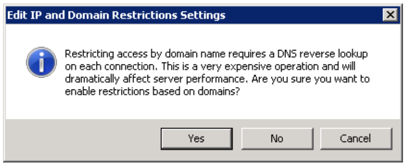

Edit IP and Domain Restrictions Settings

Restricting access by domain name requires a DNS reverse lookup on each connection. This is a very expensive operation and will dramatically affect server performance. Are you sure you want to enable restrictions based on domains?

FIGURE 6-42 Viewing a warning when enabling domain name restrictions.

After you have enabled domain name restrictions, you can use the *Add Allow Entry* and *Add Deny Entry* commands to configure the rules. As shown in Figure 6-43, the dialog boxes include a setting for Domain Name.

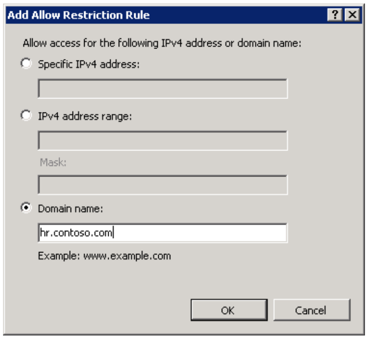

Add Allow Restriction Rule

Allow access for the following IPv4 address or domain name:

○ Specific IPv4 address:

○ IPv4 address range:

Mask:

● Domain name: hr.contoso.com

Example: www.example.com

FIGURE 6-43 Adding a domain name restriction to a website.

As mentioned earlier, the default behavior for allow and deny entries is for these restrictions to flow from parent objects to child objects. If you have made explicit changes to the settings for an object such as a web application, you can use the *Revert To Parent* command in the Actions pane to remove settings at that level. The effective settings are then based on the parent hierarchy.

Configuring .NET Trust Levels

The .NET Framework technology provides web developers with a strong set of features for implementing applications. The functionality includes web applications (based on the ASP.NET platform) and other managed code features. It is relatively simple to create .NET applications that can perform a wide array of operations on a computer. From a security standpoint, however, it is important to restrict the permissions that are granted to a .NET application. Malicious or defective code can cause problems ranging from unauthorized access to data to the accidental deletion of content.

To help systems administrators manage permissions on production servers better, IIS supports Code Access Security (CAS) policy. CAS policies can determine which operations are available to .NET-based application code. There are two main types of configuration. The full trust option provides ASP.NET application code with all permissions on the computer. For compatibility reasons, this is the default setting for applications that are based on.NET Framework 1.0, 1.1, and 2.0.

Understanding Partial Trust Levels

The other CAS policy option is partial trust, which limits the actions .NET applications can perform. These options are available to applications that are built using .NET Framework 1.1 and .NET Framework 2.0. The goal of partial trust is to enable only the permissions that are necessary for a specific web application.

Trust levels can be configured at different levels in the web server object hierarchy. These levels include:

- Web server
- Websites
- Web applications
- Virtual directories and physical folders

As with other security-related settings, trust levels that are defined at parent levels automatically apply to child objects unless they are specifically overridden. In general, define *.NET trust level* settings at the highest relevant setting. For example, if none of the web applications in a website should have full permissions, you can configure these settings at the site level. You can then manage exceptions by assigning the necessary .NET trust level settings for specific web applications or folders.

Understanding .NET Trust Levels

The .NET Framework contains many features and operations that can cause security issues on a web server. To provide a simpler method of configuring and applying trust settings, IIS includes five built-in levels that can be applied to IIS objects. The specific settings for each level are defined within various .config files. (For more information about using configuration files, see Chapter 5.) It is also possible to view and modify the settings in these files by using an XML editor or text editor. Table 6-2 lists the levels and their effects.

TABLE 6-2 .NET Trust Levels and Their Descriptions

.NET TRUST LEVEL	CONFIG FILE NAME	DESCRIPTION	RESTRICTED ACTIONS
Full (internal)	N/A	Provides full permissions to an ASP.NET application	N/A
High	Web_hightrust.config	Provides access to most actions on the server and is designed for well-trusted and well-tested web applications	■ Calling unmanaged code ■ Calling serviced components ■ Writing to the event log ■ Accessing message queuing services ■ Accessing ODBC, OLEDB, and Oracle data sources
Medium	Web_medium trust.config	Provides additional restrictions for web applications that should not need to access the file system or registry	■ Accessing files outside of the application's directory ■ Accessing the registry ■ Making network or web service calls
Low	Web_lowtrust.config	Further restricts application capabilities	■ Writing to the file system ■ Calling the Assert method (a method that is often used for testing application code)
Minimal	Web_minimal trust.config	Allows only *Execute* permissions and prevents access to other resources on the computer	Performing actions that require permissions greater than *Execute*

EXAM TIP

Expect to see questions on .NET trust levels. Familiarize yourself with the purpose of each .NET trust level and keep in mind which types of operations are considered the most risky. The levels are cumulative, from a standpoint of restrictions. For example, the Low level adds further restrictions to the Medium level and the levels above it. On the exam, be sure to understand a web application's requirements before deciding which trust level is most appropriate.

The default setting is Full (internal), which provides the best compatibility but also the greatest security risk. Whenever possible, lower the .NET trust levels to ensure that application code is being run with minimal permissions. Often, this involves interactions with web developers to determine requirements and perform complete testing at various security levels.

Configuring .NET Trust Levels

To configure .NET trust levels by using IIS Manager, select the object for which you want to assign the settings and then double-click .NET Trust Levels from Features View, as shown in Figure 6-44. To change the setting, select the appropriate level from the drop-down list and click Apply. After the trust level is set, it will apply to all ASP.NET applications running at the selected level and to any child objects unless the settings are explicitly overridden.

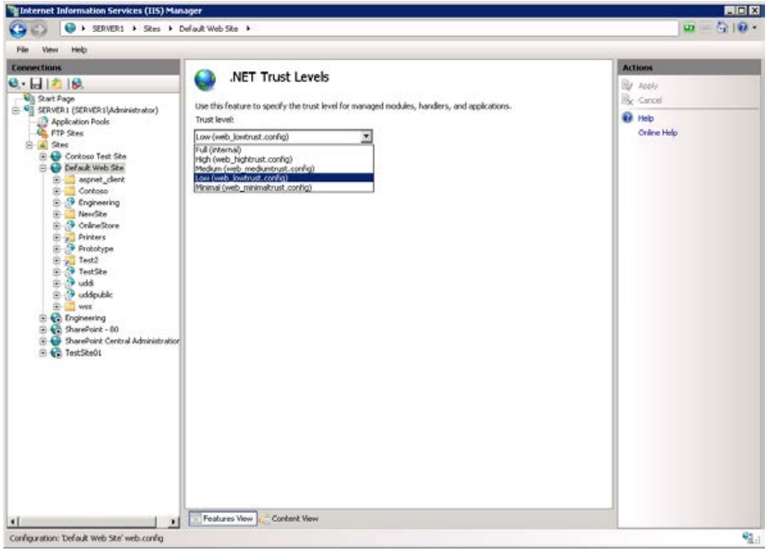

.NET Trust Levels

Use this feature to specify the trust level for managed modules, handlers, and applications.

Trust level: Low (web_lowtrust.config)

- Full (internal)
- High (web_hightrust.config)
- Medium (web_mediumtrust.config)
- Low (web_lowtrust.config)
- Minimal (web_minimaltrust.config)

Configuration: 'Default Web Site' web.config

FIGURE 6-44 Viewing .NET Trust Levels options for a website.

Quick Check

1. How can you manage which content is available to users without requiring any authentication?
2. What are the requirements for enabling SSL on an IIS web server that will be accessible from the Internet?
3. How can you restrict access to an IIS web application to only a limited set of computers?

Quick Check Answers

1. Assuming that anonymous authentication is enabled, IIS will use NTFS file system permissions settings to determine which content requires credentials to access.
2. To provide SSL security for Internet-based connections, obtain a security certificate from a trusted third-party issuer and install the certificate on the web server. You can then enable SSL through an HTTPS site binding.

3. **You can use IP address restrictions to specify which computers should have access to an IIS web server. Other options are also possible, including the use of client certificates.**

PRACTICE Secure Web Servers and Web Content

In these exercises, you apply the information you learned about ways to add security to specific web content. The steps assume that you have installed the Web Server (IIS) role using the default settings and that you are familiar with the process of adding role services.

EXERCISE 1 Managing and Testing Authentication Settings

In this exercise, you configure and verify the effects of various authentication settings.

1. Log on to Contoso.local from Server2 as a domain administrator.
2. Using Server Manager, add the following role services to the Web Server (IIS) role. You can find these role services in the Security group.
   - Basic Authentication
   - Windows Authentication
   - Digest Authentication
   - URL Authorization
   - IP and Domain Restrictions
3. When you are finished, close Server Manager.
4. Open IIS Manager, browse to the *Sites* container, and select Default Web Site in the left pane. Double-click Authentication in Features View.

   The default settings specify that only anonymous authentication is enabled.
5. Click Default Web Site again and then click Browse *:80 (http) in the Actions pane. Verify that the default IIS start page is displayed. Keep the web browser open, but return to IIS Manager.
6. Again, double-click Authentication in Features View. Select Digest Authentication and then click Enable in the Actions pane.
7. Return to Internet Explorer and type **http://server2** in the address box.

   Note that you are not prompted to provide authentication information. This is because anonymous authentication is still enabled for the site.
8. Return to IIS Manager, select anonymous authentication, and then click Disable in the Actions pane.
9. Return to Internet Explorer and refresh the page. This time, you are prompted to provide logon information to access the site. Enter your username and password and then click OK to verify that the site loads. Optionally, you can provide invalid logon

information (such as a user account that does not exist) to see that you cannot access the site. When you are finished, close Internet Explorer.

10. To restore the original authentication settings, return to IIS Manager. Disable Digest Authentication and enable anonymous authentication.
11. When you are finished, close IIS Manager and log off Server2.

EXERCISE 2 Configuring Server Certificates

In this exercise, you create a self-signed security certificate for Server2.contoso.local and then require SSL to access Default Web Site and test the settings by using Internet Explorer.

1. Log on to Server2 as a user with *Administrator* permissions on the computer.
2. Open IIS Manager and select the server object in the left pane.
3. Double-click Server Certificates in the IIS section of Features View.

   Depending on which roles and role services have been installed on the local server, some certificates might already be available on the server.
4. Click Create Self-Signed Certificate in the Actions pane.
5. For the name of the certificate, type **Test Local SSL Certificate** and then click OK.

   You should now see the new certificate in the Server Certificates view of IIS Manager.
6. To view the properties of the new certificate, right-click it and select View.

   Note details such as the issuer (which is the name of the server) and the dates for which the certificate is valid. (New certificates expire in one year.) The Certification Path tab shows only the certificate itself and not a chain of trust, which signifies that it has not been issued by a trusted certificate authority (CA). For this reason, the certificate is not suited for access by users on public networks such as the Internet.
7. Click OK when you are finished.
8. In IIS Manager, right-click the *Default Web Site* object and select Edit Bindings.
9. In the Site Bindings dialog box, Click Add.
10. In the Add Site Bindings dialog box, in the Type drop-down list, select HTTPS.
11. In the SSL Certificate list, select Test Local SSL Certificate. Click OK to save the settings and then click Close. (If, after you click OK, a message box informs you that the binding is already being used, first click Yes and then click Close.)
12. In IIS Manager, ensure that the *Default Web Site* object is selected and then double-click SSL Settings in Features View. Enable the Require SSL option and then click Apply in the Actions pane.
13. Click Back to return to Features View for Default Web Site. In the Actions pane, choose Browse *:80 (http). This launches Internet Explorer and attempts to connect to the site by using a non-SSL (HTTP) connection. You receive an error stating, “The Page You Are Trying To Access Is Secure With Secure Sockets Layer (SSL).” Close Internet Explorer.
14. In IIS Manager, click Browse *:443 (https) in the Actions pane.

This time, you receive a warning stating that there is a problem with the website's security certificate. This is because a self-signed certificate was not issued by a trusted CA.

15. To access the site anyway, click Continue To This Website. In the Security Alert message box, click OK.

    The address bar turns red, and a Certificate Error message appears. The site content is, however, accessible.

16. When you are finished, close Internet Explorer.
17. In IIS Manager, double-click the SSL Settings feature for Default Web Site and disable the Require SSL option. Click Apply in the Actions pane to save the setting.
18. When you are finished, close IIS Manager.

Lesson Summary

- Anonymous authentication provides access to site content without requiring users to provide credentials.
- Forms authentication is useful for public websites and applications that manage their own security.
- URL authorization rules can determine which users or groups have access to which website content.
- Web server administrators can use Internet server certificates to enable encrypted connections through SSL over the Internet.
- Administrators can create self-signed server certificates for testing and development purposes.
- You can use IP Address And Domain Restrictions to restrict access to web content.
- .NET trust levels restrict the permissions that managed code has on a web server.

Lesson Review

You can use the following questions to test your knowledge of the information in Lesson 2, "Controlling Access to Web Services." The questions are also available on the companion CD if you prefer to review them in electronic form.

> ***NOTE*** **ANSWERS**
>
> **Answers to these questions and explanations of why each answer choice is correct or incorrect are located in the "Answers" section at the end of the book.**

1. You are an IIS web server administrator implementing authentication settings for a new Human Resources website. According to the requirements for the website, users should be prompted for authentication information when they attempt to access the site. The site will be accessed only by users who have accounts in your organization's

Active Directory domain. You have already configured the file system permissions for the content based on the appropriate settings. You also want to maximize security of the site. Which two actions should you take to meet these requirements? (Each correct answer presents part of a complete solution. Choose two.)

- **A.** Enable Windows authentication.
- **B.** Enable basic authentication.
- **C.** Disable anonymous authentication.
- **D.** Enable anonymous authentication.

2. You are a systems administrator troubleshooting a problem with accessing a web server running Windows Server 2008 R2. Previously, another administrator created and installed a server certificate on the computer. Users report that they are able to connect to the site using HTTP, but that they receive a warning in Internet Explorer when trying to connect by HTTPS. You want to enable users to connect using both HTTP and HTTPS. You attempt to access the site by using an instance of Internet Explorer on the server itself, and you receive the same warning message for HTTPS connections. How can you resolve this issue?

   - **A.** Change the site binding for the website to enable connections on port 443.
   - **B.** Change the SSL settings for the website to enable the Require SSL option.
   - **C.** Obtain and install an Internet certificate on the web server.
   - **D.** Export and reimport the existing security certificate.
   - **E.** Reconfigure clients' firewall settings to enable traffic on port 443.

Chapter Review

To further practice and reinforce the skills you learned in this chapter, you can perform the following tasks:

- Review the chapter summary.
- Review the list of key terms introduced in this chapter.
- Complete the case scenarios. These scenarios set up real-world situations involving the topics of this chapter and ask you to create a solution.
- Complete the suggested practices.
- Take a practice test.

Chapter Summary

- Web server administrators should focus on implementing defense in depth and reducing the attack surface of IIS by using features such as request handler mappings.
- IIS allows for managing remote administration by configuring users, permissions, and feature delegation for the management service.
- Server administrators can control access to the web server by using authentication settings, URL authorization rules, server certificates, and IP Address And Domain Restrictions.

Key Terms

Do you know what these key terms mean? You can check your answers by looking up the terms in the glossary at the end of the book.

- ASP.NET impersonation
- attack surface
- certificate authority (CA)
- Client Certificate Authentication
- defense in depth
- domain restrictions (IIS)
- feature delegation (IIS)
- handler mappings (IIS)
- IIS Management Service
- IIS Manager credentials
- Internet certificate request (IIS)
- IP address restrictions (IIS)
- modules (IIS)

- .NET trust levels
- request handlers
- self-signed certificate
- server certificates
- URL authorization rules

Case Scenarios

In these case scenarios, you apply the information that you have learned about securing IIS.

Case Scenario 1: Configuring Remote Management for IIS

You are a systems administrator responsible for managing four web servers running Windows Server 2008 R2. You would like to use a single instance of IIS Manager to connect to all the servers. In addition, three other systems administrators need to manage the servers. One of these administrators is a consultant, and she does not have a Windows domain or local user account. You would like to create a username and password for her that is limited to managing IIS. You want all administrators other than you to be able to view but not change settings for the Default Document and Directory Browsing features.

1. What is the easiest method of managing settings for all the web servers by using IIS Manager?
2. How can you set up a username and password for a remote systems administrator?
3. How can you prevent the other users from modifying the Default Document and Directory Browsing features when using IIS Manager?

Case Scenario 2: Increasing Website Security

You are a systems administrator responsible for implementing and managing security for a production web server running Windows Server 2008 R2. The server is accessible from the Internet and contains eight websites. Each site contains at least one web application. A web application named Customer Database contains an ASP.NET 2.0 web application that needs to access a remote database server. Another website, named Service Desk, contains static content, most of which should be available to all users. However, a folder called Admin should be available only to specific users. Finally, you have a new requirement for an application named Contoso Central that specifies that all connections should use an encrypted connection.

1. Which .NET trust level should you configure for the Customer Database application?
2. How can you configure security for the Admin folder within the Service Desk application?
3. How can you require encryption security for connections to the Contoso Central application?

Suggested Practices

To help you successfully master the exam objectives presented in this chapter, complete the following tasks. The practices in this section enable you to apply the methods you have learned to secure IIS-based web servers, websites, and web applications.

- **Practice 1** Create a new website by using IIS Manager. The content of the website can contain copies of the Iisstart.htm file or other HTML files you have available. Place some of the files within folders and create scenarios in which you want to protect content. Apply file system permissions, authentication settings, and URL authorization rules to ensure that only certain users can access the site. For example, create a new subfolder within a web application called SecureDocuments. Place the appropriate limitations to ensure that users must provide credentials to access the content. Also, test the effects of changing handler mappings. For example, remove the StaticFile handler mapping for a website and test the effects by using Internet Explorer. You can also add your own custom handler mappings for new file types (such as files that have a .secure extension).
- **Practice 2** Add the Management Service role service to a web server running Windows Server 2008 R2. Practice using a variety of security features to support web server administrators with different levels of restrictions. Options to test include:
  - Creating IIS Manager users.
  - Assigning IIS Manager Permissions settings to control which websites and web applications administrators can access.
  - Assigning permissions to non-administrator users who have Windows accounts.
  - Creating IP address restrictions to control which computers can administer IIS.
  - Using feature delegation to control which settings can be modified by using IIS Manager.

  To test settings most efficiently, it is recommended that you use a remote computer running Windows 7 or Windows Server 2008 R2 that has IIS 7 Manager installed.
- **Practice 3** View the following webcasts and resources for more information about IIS:
  - The webcast entitled, "Secure, Simplified Web Publishing Using Internet Information Services 7.0 (Level 300)," by Robert McMurray, available on the companion CD in the Webcasts folder. Alternatively, you can find this webcast by visiting *http://msevents.microsoft.com* and searching for event ID 1032352159.
  - The webcast entitled, "Securing and Tuning Internet Information Services 7.0 (Level 300)," by Nazim Lala, available on the companion CD in the Webcasts folder. Alternatively, you can find this webcast by visiting *http://msevents.microsoft.com* and searching for event ID 1032352141.
  - The Microsoft Internet Information Services website at *http://www.microsoft.com/iis*.
  - The IIS.NET website at *http://www.iis.net*.

- IIS 7 webcasts at *http://learn.iis.net/Videos*.
- IIS 7 virtual labs at *http://technet.microsoft.com/en-us/virtuallabs/bb499672*.

Take a Practice Test

The practice tests on this book's companion CD offer many options. For example, you can test yourself on just one exam objective, or you can test yourself on all the 70-643 certification exam content. You can set up the test so that it closely simulates the experience of taking a certification exam, or you can set it up in study mode so that you can look at the correct answers and explanations after you answer each question.

MORE INFO **PRACTICE TESTS**

For details about all the practice test options available, see the "How to Use the Practice Tests" section in this book's introduction.

Index

About the Authors

J.C. MACKIN (MCITP, MCTS, MCSE, MCDST, MCT) is a writer, editor, and trainer who has been working with Microsoft networks since Windows NT 4.0. Books he has previously authored or coauthored include *MCSA/MCSE Self-Paced Training Kit (Exam 70-291): Implementing, Managing, and Maintaining a Microsoft Windows Server 2003 Network Infrastructure*, *MCITP Self-Paced Training Kit (Exam 70-443): Designing a Database Server Infrastructure Using Microsoft SQL Server 2005*, and *MCITP Self-Paced Training Kit (Exam 70-685): Windows 7, Enterprise Desktop Support Technician*. He also holds a master's degree in Telecommunication and Network Management. When not working with computers, J.C. can be found with a panoramic camera photographing villages in Italy or France.

ANIL DESAI is an independent consultant based in Austin, Texas. He specializes in evaluating, implementing, and managing information technology (IT) solutions. He has worked extensively with IT management, development, and database technology. Anil holds many certifications and is a Microsoft Most Valuable Professional (MVP) (Windows Server – Virtualization). Anil is the author of numerous technical books focusing on the Windows Server platform, virtualization, databases, and IT management best practices. He is also a frequent contributor to numerous IT publications and conferences. For more information, please see *http://AnilDesai.net*.

Windows Server 2008 Resource Kit— Your Definitive Resource!

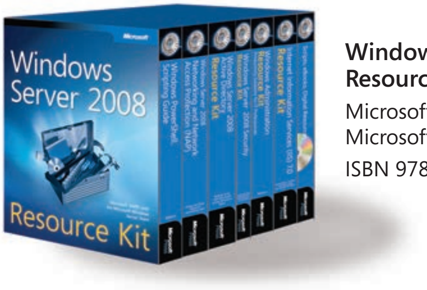

Windows Server® 2008 Resource Kit

Microsoft® MVPs with the
Microsoft Windows Server Team

ISBN 9780735623613

Your definitive reference for deployment and operations—from the experts who know the technology best. Get in-depth technical information on Active Directory®, Windows PowerShell® scripting, advanced administration, networking and network access protection, security administration, IIS, and other critical topics—plus an essential toolkit of resources on CD.

ALSO AVAILABLE AS SINGLE VOLUMES

Windows Server 2008 Security Resource Kit

Jesper M. Johansson et al. with
Microsoft Security Team

ISBN 9780735625044

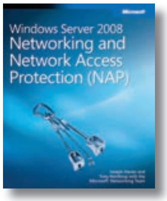

Windows Server 2008 Networking and Network Access Protection (NAP)

Joseph Davies, Tony Northrup,
Microsoft Networking Team

ISBN 9780735624221

Windows Server 2008 Active Directory Resource Kit

Stan Reimer et al. with
Microsoft Active Directory Team

ISBN 9780735625150

Windows® Administration Resource Kit: Productivity Solutions for IT Professionals

Dan Holme

ISBN 9780735624313

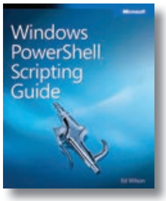

Windows Powershell Scripting Guide

Ed Wilson

ISBN 9780735622791

Internet Information Services (IIS) 7.0 Resource Kit

Mike Volodarsky et al.
with Microsoft IIS Team

ISBN 9780735624412

microsoft.com/mspress

Microsoft® Press

Windows Server 2008—Resources for Administrators

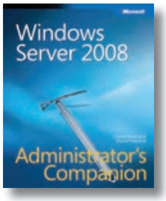

Windows Server® 2008 Administrator's Companion

Charlie Russel and Sharon Crawford

ISBN 9780735625051

Your comprehensive, one-volume guide to deployment, administration, and support. Delve into core system capabilities and administration topics, including Active Directory®, security issues, disaster planning/recovery, interoperability, IIS 7.0, virtualization, clustering, and performance tuning.

Windows Server 2008 Administrator's Pocket Consultant, Second Edition

William R. Stanek

ISBN 9780735627116

Portable and precise—with the focused information you need for administering server roles, Active Directory, user/group accounts, rights and permissions, file-system management, TCP/IP, DHCP, DNS, printers, network performance, backup, and restoration.

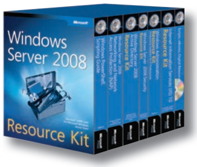

Windows Server 2008 Resource Kit

Microsoft MVPs with Microsoft Windows Server Team

ISBN 9780735623613

Six volumes! Your definitive resource for deployment and operations—from the experts who know the technology best. Get in-depth technical information on Active Directory, Windows PowerShell® scripting, advanced administration, networking and network access protection, security administration, IIS, and more—plus an essential toolkit of resources on CD.

Internet Information Services (IIS) 7.0 Administrator's Pocket Consultant

William R. Stanek

ISBN 9780735623644

This pocket-sized guide delivers immediate answers for administering IIS 7.0. Topics include customizing installation; configuration and XML schema; application management; user access and security; Web sites, directories, and content; and performance, backup, and recovery.

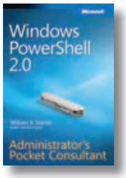

Windows PowerShell 2.0 Administrator's Pocket Consultant

William R. Stanek

ISBN 9780735625952

The practical, portable guide to using *cmdlets* and scripts to automate everyday system administration—including configuring server roles, services, features, and security settings; managing TCP/IP networking; monitoring and tuning performance; and other essential tasks.

ALSO SEE

Windows PowerShell 2.0 Best Practices

ISBN 9780735626461

Windows® Administration Resource Kit: Productivity Solutions for IT Professionals

ISBN 9780735624313

Windows Server 2008 Hyper-V™ Resource Kit

ISBN 9780735625174

Windows Server 2008 Security Resource Kit

ISBN 9780735625044

microsoft.com/mspress

Microsoft® Press