

OFFICIAL MICROSOFT LEARNING PRODUCT

# 20417A

Upgrading Your Skills to MCSA  
Windows Server® 2012

MCT USE ONLY. STUDENT USE PROHIBITED

Information in this document, including URL and other Internet Web site references, is subject to change without notice. Unless otherwise noted, the example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious, and no association with any real company, organization, product, domain name, email address, logo, person, place or event is intended or should be inferred. Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

The names of manufacturers, products, or URLs are provided for informational purposes only and Microsoft makes no representations and warranties, either expressed, implied, or statutory, regarding these manufacturers or the use of the products with any Microsoft technologies. The inclusion of a manufacturer or product does not imply endorsement of Microsoft of the manufacturer or product. Links may be provided to third party sites. Such sites are not under the control of Microsoft and Microsoft is not responsible for the contents of any linked site or any link contained in a linked site, or any changes or updates to such sites. Microsoft is not responsible for webcasting or any other form of transmission received from any linked site. Microsoft is providing these links to you only as a convenience, and the inclusion of any link does not imply endorsement of Microsoft of the site or the products contained therein.

© 2012 Microsoft Corporation. All rights reserved.

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

Product Number: 20417A

Part Number: X18-48638

Released: 08/2012



**MICROSOFT LICENSE TERMS**  
**OFFICIAL MICROSOFT LEARNING PRODUCTS**  
**MICROSOFT OFFICIAL COURSE Pre-Release and Final Release Versions**

---

These license terms are an agreement between Microsoft Corporation and you. Please read them. They apply to the Licensed Content named above, which includes the media on which you received it, if any. These license terms also apply to any updates, supplements, internet based services and support services for the Licensed Content, unless other terms accompany those items. If so, those terms apply.

**BY DOWNLOADING OR USING THE LICENSED CONTENT, YOU ACCEPT THESE TERMS. IF YOU DO NOT ACCEPT THEM, DO NOT DOWNLOAD OR USE THE LICENSED CONTENT.**

**If you comply with these license terms, you have the rights below.**

---

**1. DEFINITIONS.**

- a. "Authorized Learning Center" means a Microsoft Learning Competency Member, Microsoft IT Academy Program Member, or such other entity as Microsoft may designate from time to time.
- b. "Authorized Training Session" means the Microsoft-authorized instructor-led training class using only MOC Courses that are conducted by a MCT at or through an Authorized Learning Center.
- c. "Classroom Device" means one (1) dedicated, secure computer that you own or control that meets or exceeds the hardware level specified for the particular MOC Course located at your training facilities or primary business location.
- d. "End User" means an individual who is (i) duly enrolled for an Authorized Training Session or Private Training Session, (ii) an employee of a MPN Member, or (iii) a Microsoft full-time employee.
- e. "Licensed Content" means the MOC Course and any other content accompanying this agreement. Licensed Content may include (i) Trainer Content, (ii) software, and (iii) associated media.
- f. "Microsoft Certified Trainer" or "MCT" means an individual who is (i) engaged to teach a training session to End Users on behalf of an Authorized Learning Center or MPN Member, (ii) currently certified as a Microsoft Certified Trainer under the Microsoft Certification Program, and (iii) holds a Microsoft Certification in the technology that is the subject of the training session.
- g. "Microsoft IT Academy Member" means a current, active member of the Microsoft IT Academy Program.
- h. "Microsoft Learning Competency Member" means a Microsoft Partner Network Program Member in good standing that currently holds the Learning Competency status.
- i. "Microsoft Official Course" or "MOC Course" means the Official Microsoft Learning Product instructor-led courseware that educates IT professionals or developers on Microsoft technologies.

- j. "Microsoft Partner Network Member" or "MPN Member" means a silver or gold-level Microsoft Partner Network program member in good standing.
  - k. "Personal Device" means one (1) device, workstation or other digital electronic device that you personally own or control that meets or exceeds the hardware level specified for the particular MOC Course.
  - l. "Private Training Session" means the instructor-led training classes provided by MPN Members for corporate customers to teach a predefined learning objective. These classes are not advertised or promoted to the general public and class attendance is restricted to individuals employed by or contracted by the corporate customer.
  - m. "Trainer Content" means the trainer version of the MOC Course and additional content designated solely for trainers to use to teach a training session using a MOC Course. Trainer Content may include Microsoft PowerPoint presentations, instructor notes, lab setup guide, demonstration guides, beta feedback form and trainer preparation guide for the MOC Course. To clarify, Trainer Content does not include virtual hard disks or virtual machines.
2. **INSTALLATION AND USE RIGHTS.** The Licensed Content is licensed not sold. The Licensed Content is licensed on a one copy per user basis, such that you must acquire a license for each individual that accesses or uses the Licensed Content.
- 2.1 Below are four separate sets of installation and use rights. Only one set of rights apply to you.
- a. **If you are a Authorized Learning Center:**
- i. If the Licensed Content is in digital format for each license you acquire you may either:
    - 1. install one (1) copy of the Licensed Content in the form provided to you on a dedicated, secure server located on your premises where the Authorized Training Session is held for access and use by one (1) End User attending the Authorized Training Session, or by one (1) MCT teaching the Authorized Training Session, **or**
    - 2. install one (1) copy of the Licensed Content in the form provided to you on one (1) Classroom Device for access and use by one (1) End User attending the Authorized Training Session, or by one (1) MCT teaching the Authorized Training Session.
  - ii. You agree that:
    - 1. you will acquire a license for each End User and MCT that accesses the Licensed Content,
    - 2. each End User and MCT will be presented with a copy of this agreement and each individual will agree that their use of the Licensed Content will be subject to these license terms prior to their accessing the Licensed Content. Each individual will be required to denote their acceptance of the EULA in a manner that is enforceable under local law prior to their accessing the Licensed Content,
    - 3. for all Authorized Training Sessions, you will only use qualified MCTs who hold the applicable competency to teach the particular MOC Course that is the subject of the training session,
    - 4. you will not alter or remove any copyright or other protective notices contained in the Licensed Content,

5. you will remove and irretrievably delete all Licensed Content from all Classroom Devices and servers at the end of the Authorized Training Session,
6. you will only provide access to the Licensed Content to End Users and MCTs,
7. you will only provide access to the Trainer Content to MCTs, and
8. any Licensed Content installed for use during a training session will be done in accordance with the applicable classroom set-up guide.

**b. If you are a MPN Member.**

- i. If the Licensed Content is in digital format for each license you acquire you may either:
  1. install one (1) copy of the Licensed Content in the form provided to you on (A) one (1) Classroom Device, or (B) one (1) dedicated, secure server located at your premises where the training session is held for use by one (1) of your employees attending a training session provided by you, or by one (1) MCT that is teaching the training session, **or**
  2. install one (1) copy of the Licensed Content in the form provided to you on one (1) Classroom Device for use by one (1) End User attending a Private Training Session, or one (1) MCT that is teaching the Private Training Session.
- ii. You agree that:
  1. you will acquire a license for each End User and MCT that accesses the Licensed Content,
  2. each End User and MCT will be presented with a copy of this agreement and each individual will agree that their use of the Licensed Content will be subject to these license terms prior to their accessing the Licensed Content. Each individual will be required to denote their acceptance of the EULA in a manner that is enforceable under local law prior to their accessing the Licensed Content,
  3. for all training sessions, you will only use qualified MCTs who hold the applicable competency to teach the particular MOC Course that is the subject of the training session,
  4. you will not alter or remove any copyright or other protective notices contained in the Licensed Content,
  5. you will remove and irretrievably delete all Licensed Content from all Classroom Devices and servers at the end of each training session,
  6. you will only provide access to the Licensed Content to End Users and MCTs,
  7. you will only provide access to the Trainer Content to MCTs, and
  8. any Licensed Content installed for use during a training session will be done in accordance with the applicable classroom set-up guide.

**c. If you are an End User:**

You may use the Licensed Content solely for your personal training use. If the Licensed Content is in digital format, for each license you acquire you may (i) install one (1) copy of the Licensed Content in the form provided to you on one (1) Personal Device and install another copy on another Personal Device as a backup copy, which may be used only to reinstall the Licensed Content; or (ii) print one (1) copy of the Licensed Content. You may not install or use a copy of the Licensed Content on a device you do not own or control.

d. **If you are a MCT.**

- i. For each license you acquire, you may use the Licensed Content solely to prepare and deliver an Authorized Training Session or Private Training Session. For each license you acquire, you may install and use one (1) copy of the Licensed Content in the form provided to you on one (1) Personal Device and install one (1) additional copy on another Personal Device as a backup copy, which may be used only to reinstall the Licensed Content. You may not install or use a copy of the Licensed Content on a device you do not own or control.
- ii. **Use of Instructional Components in Trainer Content.** You may customize, in accordance with the most recent version of the MCT Agreement, those portions of the Trainer Content that are logically associated with instruction of a training session. If you elect to exercise the foregoing rights, you agree: (a) that any of these customizations will only be used for providing a training session, (b) any customizations will comply with the terms and conditions for Modified Training Sessions and Supplemental Materials in the most recent version of the MCT agreement and with this agreement. For clarity, any use of “customize” refers only to changing the order of slides and content, and/or not using all the slides or content, it does not mean changing or modifying any slide or content.

2.2 **Separation of Components.** The Licensed Content components are licensed as a single unit and you may not separate the components and install them on different devices.

2.3 **Reproduction/Redistribution Licensed Content.** Except as expressly provided in the applicable installation and use rights above, you may not reproduce or distribute the Licensed Content or any portion thereof (including any permitted modifications) to any third parties without the express written permission of Microsoft.

2.4 **Third Party Programs.** The Licensed Content may contain third party programs or services. These license terms will apply to your use of those third party programs or services, unless other terms accompany those programs and services.

2.5 **Additional Terms.** Some Licensed Content may contain components with additional terms, conditions, and licenses regarding its use. Any non-conflicting terms in those conditions and licenses also apply to that respective component and supplements the terms described in this Agreement.

3. **PRE-RELEASE VERSIONS.** If the Licensed Content is a pre-release (“beta”) version, in addition to the other provisions in this agreement, then these terms also apply:

- a. **Pre-Release Licensed Content.** This Licensed Content is a pre-release version. It may not contain the same information and/or work the way a final version of the Licensed Content will. We may change it for the final version. We also may not release a final version. Microsoft is under no obligation to provide you with any further content, including the final release version of the Licensed Content.
- b. **Feedback.** If you agree to give feedback about the Licensed Content to Microsoft, either directly or through its third party designee, you give to Microsoft without charge, the right to use, share and commercialize your feedback in any way and for any purpose. You also give to third parties, without charge, any patent rights needed for their products, technologies and services to use or interface with any specific parts of a Microsoft software, Microsoft product, or service that includes the feedback. You will not give feedback that is subject to a license that requires Microsoft to license its software, technologies, or products to third parties because we include your feedback in them. These rights

survive this agreement.

- c. **Term.** If you are an Authorized Training Center, MCT or MPN, you agree to cease using all copies of the beta version of the Licensed Content upon (i) the date which Microsoft informs you is the end date for using the beta version, or (ii) sixty (60) days after the commercial release of the Licensed Content, whichever is earliest ("**beta term**"). Upon expiration or termination of the beta term, you will irretrievably delete and destroy all copies of same in the possession or under your control.

4. **INTERNET-BASED SERVICES.** Microsoft may provide Internet-based services with the Licensed Content, which may change or be canceled at any time.

- a. **Consent for Internet-Based Services.** The Licensed Content may connect to computer systems over an Internet-based wireless network. In some cases, you will not receive a separate notice when they connect. Using the Licensed Content operates as your consent to the transmission of standard device information (including but not limited to technical information about your device, system and application software, and peripherals) for internet-based services.
- b. **Misuse of Internet-based Services.** You may not use any Internet-based service in any way that could harm it or impair anyone else's use of it. You may not use the service to try to gain unauthorized access to any service, data, account or network by any means.

5. **SCOPE OF LICENSE.** The Licensed Content is licensed, not sold. This agreement only gives you some rights to use the Licensed Content. Microsoft reserves all other rights. Unless applicable law gives you more rights despite this limitation, you may use the Licensed Content only as expressly permitted in this agreement. In doing so, you must comply with any technical limitations in the Licensed Content that only allows you to use it in certain ways. Except as expressly permitted in this agreement, you may not:

- install more copies of the Licensed Content on devices than the number of licenses you acquired;
- allow more individuals to access the Licensed Content than the number of licenses you acquired;
- publicly display, or make the Licensed Content available for others to access or use;
- install, sell, publish, transmit, encumber, pledge, lend, copy, adapt, link to, post, rent, lease or lend, make available or distribute the Licensed Content to any third party, except as expressly permitted by this Agreement.
- reverse engineer, decompile, remove or otherwise thwart any protections or disassemble the Licensed Content except and only to the extent that applicable law expressly permits, despite this limitation;
- access or use any Licensed Content for which you are not providing a training session to End Users using the Licensed Content;
- access or use any Licensed Content that you have not been authorized by Microsoft to access and use; or
- transfer the Licensed Content, in whole or in part, or assign this agreement to any third party.

6. **RESERVATION OF RIGHTS AND OWNERSHIP.** Microsoft reserves all rights not expressly granted to you in this agreement. The Licensed Content is protected by copyright and other intellectual property laws and treaties. Microsoft or its suppliers own the title, copyright, and other intellectual property rights in the Licensed Content. You may not remove or obscure any copyright, trademark or patent notices that appear on the Licensed Content or any components thereof, as delivered to you.

7. **EXPORT RESTRICTIONS.** The Licensed Content is subject to United States export laws and regulations. You must comply with all domestic and international export laws and regulations that apply to the Licensed Content. These laws include restrictions on destinations, End Users and end use. For additional information, see [www.microsoft.com/exporting](http://www.microsoft.com/exporting).
8. **LIMITATIONS ON SALE, RENTAL, ETC. AND CERTAIN ASSIGNMENTS.** You may not sell, rent, lease, lend or sublicense the Licensed Content or any portion thereof, or transfer or assign this agreement.
9. **SUPPORT SERVICES.** Because the Licensed Content is "as is", we may not provide support services for it.
10. **TERMINATION.** Without prejudice to any other rights, Microsoft may terminate this agreement if you fail to comply with the terms and conditions of this agreement. Upon any termination of this agreement, you agree to immediately stop all use of and to irretrievable delete and destroy all copies of the Licensed Content in your possession or under your control.
11. **LINKS TO THIRD PARTY SITES.** You may link to third party sites through the use of the Licensed Content. The third party sites are not under the control of Microsoft, and Microsoft is not responsible for the contents of any third party sites, any links contained in third party sites, or any changes or updates to third party sites. Microsoft is not responsible for webcasting or any other form of transmission received from any third party sites. Microsoft is providing these links to third party sites to you only as a convenience, and the inclusion of any link does not imply an endorsement by Microsoft of the third party site.
12. **ENTIRE AGREEMENT.** This agreement, and the terms for supplements, updates and support services are the entire agreement for the Licensed Content.
13. **APPLICABLE LAW.**
  - a. United States. If you acquired the Licensed Content in the United States, Washington state law governs the interpretation of this agreement and applies to claims for breach of it, regardless of conflict of laws principles. The laws of the state where you live govern all other claims, including claims under state consumer protection laws, unfair competition laws, and in tort.
  - b. Outside the United States. If you acquired the Licensed Content in any other country, the laws of that country apply.
14. **LEGAL EFFECT.** This agreement describes certain legal rights. You may have other rights under the laws of your country. You may also have rights with respect to the party from whom you acquired the Licensed Content. This agreement does not change your rights under the laws of your country if the laws of your country do not permit it to do so.
15. **DISCLAIMER OF WARRANTY. THE LICENSED CONTENT IS LICENSED "AS-IS," "WITH ALL FAULTS," AND "AS AVAILABLE." YOU BEAR THE RISK OF USING IT. MICROSOFT CORPORATION AND ITS RESPECTIVE AFFILIATES GIVE NO EXPRESS WARRANTIES, GUARANTEES, OR CONDITIONS UNDER OR IN RELATION TO THE LICENSED CONTENT. YOU MAY HAVE ADDITIONAL CONSUMER RIGHTS UNDER YOUR LOCAL LAWS WHICH THIS AGREEMENT CANNOT CHANGE. TO THE EXTENT PERMITTED UNDER YOUR LOCAL LAWS, MICROSOFT CORPORATION AND ITS RESPECTIVE AFFILIATES EXCLUDE ANY IMPLIED WARRANTIES OR CONDITIONS, INCLUDING THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT.**



16. **LIMITATION ON AND EXCLUSION OF REMEDIES AND DAMAGES. TO THE EXTENT NOT PROHIBITED BY LAW, YOU CAN RECOVER FROM MICROSOFT CORPORATION AND ITS SUPPLIERS ONLY DIRECT DAMAGES UP TO USD\$5.00. YOU AGREE NOT TO SEEK TO RECOVER ANY OTHER DAMAGES, INCLUDING CONSEQUENTIAL, LOST PROFITS, SPECIAL, INDIRECT OR INCIDENTAL DAMAGES FROM MICROSOFT CORPORATION AND ITS RESPECTIVE SUPPLIERS.**

This limitation applies to

- anything related to the Licensed Content, services made available through the Licensed Content, or content (including code) on third party Internet sites or third-party programs; and
- claims for breach of contract, breach of warranty, guarantee or condition, strict liability, negligence, or other tort to the extent permitted by applicable law.

It also applies even if Microsoft knew or should have known about the possibility of the damages. The above limitation or exclusion may not apply to you because your country may not allow the exclusion or limitation of incidental, consequential or other damages.

**Please note: As this Licensed Content is distributed in Quebec, Canada, some of the clauses in this agreement are provided below in French.**

**Remarque : Ce le contenu sous licence étant distribué au Québec, Canada, certaines des clauses dans ce contrat sont fournies ci-dessous en français.**

**EXONÉRATION DE GARANTIE.** Le contenu sous licence visé par une licence est offert « tel quel ». Toute utilisation de ce contenu sous licence est à votre seule risque et péril. Microsoft n'accorde aucune autre garantie expresse. Vous pouvez bénéficier de droits additionnels en vertu du droit local sur la protection des consommateurs, que ce contrat ne peut modifier. La ou elles sont permises par le droit locale, les garanties implicites de qualité marchande, d'adéquation à un usage particulier et d'absence de contrefaçon sont exclues.

**LIMITATION DES DOMMAGES-INTÉRÊTS ET EXCLUSION DE RESPONSABILITÉ POUR LES DOMMAGES.** Vous pouvez obtenir de Microsoft et de ses fournisseurs une indemnisation en cas de dommages directs uniquement à hauteur de 5,00 \$ US. Vous ne pouvez prétendre à aucune indemnisation pour les autres dommages, y compris les dommages spéciaux, indirects ou accessoires et pertes de bénéfices.

Cette limitation concerne:

- tout ce qui est relié au le contenu sous licence , aux services ou au contenu (y compris le code) figurant sur des sites Internet tiers ou dans des programmes tiers ; et
- les réclamations au titre de violation de contrat ou de garantie, ou au titre de responsabilité stricte, de négligence ou d'une autre faute dans la limite autorisée par la loi en vigueur.

Elle s'applique également, même si Microsoft connaissait ou devrait connaître l'éventualité d'un tel dommage. Si votre pays n'autorise pas l'exclusion ou la limitation de responsabilité pour les dommages indirects, accessoires ou de quelque nature que ce soit, il se peut que la limitation ou l'exclusion ci-dessus ne s'appliquera pas à votre égard.

**EFFET JURIDIQUE.** Le présent contrat décrit certains droits juridiques. Vous pourriez avoir d'autres droits prévus par les lois de votre pays. Le présent contrat ne modifie pas les droits que vous confèrent les lois de votre pays si celles-ci ne le permettent pas.

Revised December 2011

# Welcome!

Thank you for taking our training! We've worked together with our Microsoft Certified Partners for Learning Solutions and our Microsoft IT Academies to bring you a world-class learning experience—whether you're a professional looking to advance your skills or a student preparing for a career in IT.

- **Microsoft Certified Trainers and Instructors**—Your instructor is a technical and instructional expert who meets ongoing certification requirements. And, if instructors are delivering training at one of our Certified Partners for Learning Solutions, they are also evaluated throughout the year by students and by Microsoft.
- **Certification Exam Benefits**—After training, consider taking a Microsoft Certification exam. Microsoft Certifications validate your skills on Microsoft technologies and can help differentiate you when finding a job or boosting your career. In fact, independent research by IDC concluded that 75% of managers believe certifications are important to team performance<sup>1</sup>. Ask your instructor about Microsoft Certification exam promotions and discounts that may be available to you.
- **Customer Satisfaction Guarantee**—Our Certified Partners for Learning Solutions offer a satisfaction guarantee and we hold them accountable for it. At the end of class, please complete an evaluation of today's experience. We value your feedback!

We wish you a great learning experience and ongoing success in your career!

Sincerely,

Microsoft Learning  
[www.microsoft.com/learning](http://www.microsoft.com/learning)

**Microsoft** | Learning

<sup>1</sup> IDC, Value of Certification: Team Certification and Organizational Performance, November 2006



## Acknowledgments

Microsoft Learning would like to acknowledge and thank the following for their contribution towards developing this title. Their effort at various stages in the development has ensured that you have a good classroom experience.

### Stan Reimer – Content Developer

Stan Reimer is president of S. R. Technical Services Inc., and he works as a consultant, trainer, and author. Stan has extensive experience consulting on Active Directory® and Exchange Server deployments for some of the largest companies in Canada. Stan is the lead author for two Active Directory books for Microsoft Press. For the last nine years, Stan has been writing courseware for Microsoft Learning, specializing in Active Directory and Exchange Server courses. Stan has been a Microsoft Certified Trainer (MCT) for 12 years.

### Damir Dizdarevic – Subject Matter Expert/Content Developer

Damir Dizdarevic is an MCT, Microsoft Certified Solutions Expert (MCSE), Microsoft Certified Technology Specialist (MCTS), and a Microsoft Certified Information Technology Professional (MCITP). He is a manager and trainer of the Learning Center at Logosoft d.o.o., in Sarajevo, Bosnia and Herzegovina. Damir has more than 17 years of experience on Microsoft platforms and he specializes in Windows Server®, Exchange Server, security, and virtualization. He has worked as a subject-matter expert and technical reviewer on many Microsoft Official Courses (MOC) courses, and has published more than 400 articles in various IT magazines, such as *Windows ITPro* and *INFO Magazine*. He's also a frequent and highly rated speaker on most of Microsoft conferences in Eastern Europe. Additionally, he is a Microsoft Most Valuable Professional for Windows Server Infrastructure Management.

### Gary Dunlop – Subject Matter Expert

Gary Dunlop is based in Winnipeg, Canada and is a technical consultant and trainer for Broadview Networks. He has authored a number of Microsoft Learning titles and has been an MCT since 1997.

### Siegfried Jagott – Content Developer

Siegfried Jagott is a Principal Consultant and Team Lead for the Messaging and Collaboration team at Atos Germany. He is an award-winning author of *Microsoft Exchange Server 2010 Best Practices* (Microsoft Press), and has authored and technically reviewed several Microsoft Official Curriculum (MOC) courses on various topics such as MOC 10165: Updating Your Skills from Microsoft Exchange Server 2003 or Exchange Server 2007 to Exchange Server 2010 SP1. He has coauthored various books on Windows, Microsoft System Center Virtual Machine Manager, and Exchange, and is a frequent presenter on these topics at international conferences such as IT & Dev Connections Spring 2012 in Las Vegas. Siegfried has planned, designed, and implemented some of the world's largest Windows and Exchange Server infrastructures for international customers. He received an MBA from Open University in England, and has been an MCSE since 1997.

### Orin Thomas – Content Developer

Orin Thomas is an MVP, an MCT and has a string of Microsoft MCSE and MCITP certifications. He has written more than 20 books for Microsoft® Press and is a contributing editor at Windows IT Pro magazine. He has been working in IT since the early 1990s. He is a regular speaker at events such as TechED in Australia and around the world on Windows Server, Windows Client, System Center, and security topics. Orin founded and runs the Melbourne System Center Users Group.

## **Vladimir Meloski – Content Developer**

Vladimir is a Microsoft Certified Trainer, an MVP on Exchange Server, and consultant, providing unified communications and infrastructure solutions based on Microsoft Exchange Server, Lync Server, and System Center. Vladimir has 16 years of professional IT experience, and has been involved in Microsoft conferences in Europe and the United States as a speaker, moderator, proctor for hands-on labs, and technical expert. He has also been involved as a subject matter expert and technical reviewer for several Microsoft Official Curriculum courses.

# Contents

## Module 1: Installing and Configuring Servers Based on Windows Server 2012

<b>Lesson 1:</b> Installing Windows Server 2012	1-2
<b>Lesson 2:</b> Configuring Windows Server 2012	1-13
<b>Lesson 3:</b> Configuring Remote Management for Windows Server 2012 Servers	1-21
<b>Lab:</b> Installing and Configuring Servers Based on Windows Server 2012	1-25

## Module 2: Monitoring and Maintaining Windows Server 2012

<b>Lesson 1:</b> Reasons for Monitoring Servers	2-2
<b>Lesson 2:</b> Implementing Windows Server Backup	2-11
<b>Lesson 3:</b> Implementing Server and Data Recovery	2-15
<b>Lab:</b> Monitoring and Maintaining Windows 2012 Servers	2-19

## Module 3: Managing Windows Server 2012 by Using Windows PowerShell 3.0

<b>Lesson 1:</b> Overview of Windows PowerShell 3.0	3-2
<b>Lesson 2:</b> Using Windows PowerShell 3.0 to Manage AD DS	3-9
<b>Lesson 3:</b> Managing Servers by Using Windows PowerShell 3.0	3-20
<b>Lab:</b> Managing Servers Running Windows Server 2012 by Using Windows PowerShell 3.0	3-26

## Module 4: Managing Storage for Windows Server 2012

<b>Lesson 1:</b> New Features in Windows Server 2012 Storage	4-2
<b>Lesson 2:</b> Configuring iSCSI Storage	4-12
<b>Lesson 3:</b> Configuring Storage Spaces in Windows Server 2012	4-18
<b>Lab A:</b> Managing Storage for Servers Based on Windows Server 2012	4-23
<b>Lesson 4:</b> Configuring BranchCache in Windows Server 2012	4-25
<b>Lab:</b> Implementing BranchCache	4-36

## Module 5: Implementing Network Services

<b>Lesson 1:</b> Implementing DNS and DHCP Enhancements	5-2
<b>Lesson 2:</b> Implementing IP Address Management	5-10
<b>Lesson 3:</b> NAP Overview	5-14
<b>Lesson 4:</b> Implementing NAP	5-20
<b>Lab:</b> Implementing Network Services	5-25

## Module 6: Implementing DirectAccess

<b>Lesson 1:</b> Overview of DirectAccess	6-2
<b>Lesson 2:</b> Installing and Configuring DirectAccess Components	6-14
<b>Lab:</b> Implementing DirectAccess	6-24

**Module 7: Implementing Failover Clustering**

<b>Lesson 1:</b> Overview of Failover Clustering	7-2
<b>Lesson 2:</b> Implementing a Failover Cluster	7-13
<b>Lesson 3:</b> Configuring Highly Available Applications and Services on a Failover Cluster	7-18
<b>Lesson 4:</b> Maintaining a Failover Cluster	7-22
<b>Lesson 5:</b> Implementing a Multisite Failover Cluster	7-27
<b>Lab:</b> Implementing Failover Clustering	7-32

**Module 8: Implementing Hyper-V**

<b>Lesson 1:</b> Configuring Hyper-V Servers	8-2
<b>Lesson 2:</b> Configuring Hyper-V Storage	8-8
<b>Lesson 3:</b> Configuring Hyper-V Networking	8-16
<b>Lesson 4:</b> Configuring Hyper-V Virtual Machines	8-21
<b>Lab:</b> Implementing Server Virtualization with Hyper-V	8-27

**Module 9: Implementing Failover Clustering with Hyper-V**

<b>Lesson 1:</b> Overview of the Integration of Hyper-V with Failover Clustering	9-2
<b>Lesson 2:</b> Implementing Hyper-V Virtual Machines on Failover Clusters	9-7
<b>Lesson 3:</b> Implementing Hyper-V Virtual Machine Movement	9-14
<b>Lesson 4:</b> Managing Hyper-V Virtual Environments by Using System Center Virtual Machine Manager	9-19
<b>Lab:</b> Implementing Failover Clustering with Hyper-V	9-29

**Module 10: Implementing Dynamic Access Control**

<b>Lesson 1:</b> Overview of Dynamic Access Control	10-2
<b>Lesson 2:</b> Planning for a Dynamic Access Control Implementation	10-8
<b>Lesson 3:</b> Configuring Dynamic Access Control	10-13
<b>Lab:</b> Implementing Dynamic Access Control	10-22

**Module 11: Implementing Active Directory Domain Services**

<b>Lesson 1:</b> Deploying AD DS Domain Controllers	11-2
<b>Lesson 2:</b> Configuring AD DS Domain Controllers	11-11
<b>Lesson 3:</b> Implementing Service Accounts	11-16
<b>Lesson 4:</b> Implementing Group Policy in AD DS	11-19
<b>Lesson 5:</b> Maintaining AD DS	11-28
<b>Lab:</b> Implementing AD DS	11-35

**Module 12: Implementing Active Directory Federation Services**

<b>Lesson 1:</b> Overview of Active Directory Federation Services	12-2
<b>Lesson 2:</b> Deploying Active Directory Federation Services	12-11
<b>Lesson 3:</b> Implementing AD FS for a Single Organization	12-17
<b>Lesson 4:</b> Deploying AD FS in a Business-to-Business Federation Scenario	12-23
<b>Lab:</b> Implementing AD FS	12-28

**Lab Answer Keys**

<b>Module 1 Lab:</b> Installing and Configuring Servers Based on Windows Server 2012	L1-1
<b>Module 2 Lab:</b> Monitoring and Maintaining Windows 2012 Servers	L2-7
<b>Module 3 Lab:</b> Managing Servers Running Windows Server 2012 by Using Windows PowerShell 3.0	L3-15
<b>Module 4 Lab A:</b> Managing Storage for Servers Based on Windows Server 2012	L4-19
<b>Module 4 Lab B:</b> Implementing BranchCache	L4-26
<b>Module 5 Lab:</b> Implementing Network Services	L5-31
<b>Module 6 Lab:</b> Implementing DirectAccess	L6-43
<b>Module 7 Lab:</b> Implementing Failover Clustering	L7-55
<b>Module 8 Lab:</b> Implementing Server Virtualization with Hyper-V	L8-63
<b>Module 9 Lab:</b> Implementing Failover Clustering with Hyper-V	L9-71
<b>Module 10 Lab:</b> Implementing Dynamic Access Control	L10-77
<b>Module 11 Lab:</b> Implementing AD DS	L11-89
<b>Module 12 Lab:</b> Implementing AD FS	L12-97

**MCT USE ONLY. STUDENT USE PROHIBITED**

# About This Course

This section provides you with a brief description of the course, audience, suggested prerequisites, and course objectives.

## Course Description

**Note:** This first release ("A") Microsoft® Official Courses (MOC) version of course 20417A has been developed on Windows Server® 2012 RC. Microsoft Learning will release a "B" version of this course after the release-to-manufacturing (RTM) version of the software is available.

This course is designed primarily for people who want to upgrade their technical skills from Windows Server 2008 and Windows Server 2008 R2 to Windows Server 2012. It presumes a high level of knowledge about previous Windows Server versions. This course also serves as preparation for taking exam 70-417, on the upgrade path to a new MCSA: Windows Server 2012 certification.

## Audience

The primary audience for this course is Information Technology (IT) professionals who are experienced Windows Server 2008 Server Administrators, and who carry out day-to-day management and administrative tasks, and want to update their skills and knowledge to Windows Server 2012.

The secondary audience for this course includes candidates who hold existing credentials in Windows Server 2008 at Technology Specialist (TS) or Professional (PRO) level, and who want to migrate their current credentials to the new credential of Microsoft Certified Solutions Associate (MCSA) with Windows Server 2012.

## Student Prerequisites

In addition to their professional experience, students who attend this training should have the following technical knowledge:

- Two or more years of experience deploying and managing Windows Server 2008
- Experience with Windows networking technologies and implementation
- Experience with Active Directory® technologies and implementation
- Experience with Windows Server 2008 server virtualization technologies and implementation

Students attending this course are expected to have passed the following exams, or have equivalent knowledge:

- Exam 70-640: Windows Server 2008 Active Directory, Configuring
- Exam 70-642: Windows Server 2008 Network Infrastructure, Configuring
- Exam 70-646: Windows Server 2008, Server Administrator

## Course Objectives

After completing this course, students will be able to:

- Install and configure Windows Server 2012 servers.
- Monitor and maintain Windows Server 2012 servers.
- Use Windows PowerShell® 3.0 to manage Windows Server 2012 servers.
- Configure storage on Windows Server 2012 servers.
- Deploy and manage network services.
- Deploy and manage a DirectAccess infrastructure.
- Provide high availability for network services and applications by implementing failover clustering.
- Deploy and configure virtual machines on Hyper-V®.
- Deploy and manage Hyper-V virtual machines in a failover cluster.
- Configure Dynamic Access Control to manage and audit access to shared files.
- Implement the new features in Active Directory Domain Services (AD DS) for Windows Server 2012.
- Plan and implement an Active Directory Federation Services (AD FS) deployment.

## Course Outline

This section provides an outline of the course:

**Module 1,** Installing and Configuring Servers Based on Windows Server 2012

**Module 2,** Monitoring and Maintaining Windows Server 2012

**Module 3,** Managing Windows Server 2012 by Using Windows PowerShell 3.0

**Module 4,** Managing Storage for Windows Server 2012

**Module 5,** Implementing Network Services

**Module 6,** Implementing DirectAccess

**Module 7,** Implementing Failover Clustering

**Module 8,** Implementing Hyper-V

**Module 9,** Implementing Failover Clustering with Hyper-V

**Module 10,** Implementing Dynamic Access Control

**Module 11,** Implementing Active Directory Domain Services

**Module 12,** Implementing Active Directory Federation Services



## Exam/Course Mapping

This course, 20417A: *Upgrading Your Skills to MCSA Windows Server 2012*, has a direct mapping of its content to the objective domain for the Microsoft exam 70-417: *Upgrading Your Skills to MCSA Windows Server 2012*.

The below table is provided as a study aid that will assist you in preparation for taking this exam and to show you how the exam objectives and the course content fit together. The course is not designed exclusively to support the exam but rather provides broader knowledge and skills to allow a real-world implementation of the particular technology. The course will also contain content that is not directly covered in the examination and will use the unique experience and skills of your qualified Microsoft Certified Trainer.



**Note:** The exam objectives are available online at the following URL:

<http://www.microsoft.com/learning/en/us/exam.aspx?ID=70-417&locale=en-us#tab2>.

Exam Objective Domains		Course Content		
Exam 70-410: Installing and Configuring Windows Server 2012				
Install and Configure Servers		Module	Lesson	Lab
Install servers.	This objective may include but is not limited to: Plan for a server installation; plan for server roles; plan for a server upgrade; install Server Core; optimize resource utilization by using Features on Demand; migrate roles from previous versions of Windows Server	Mod 1	Lesson 1/2	Mod 1 Ex 1
Configure servers.	This objective may include but is not limited to: Configure Server Core; delegate administration; add and remove features in offline images; deploy roles on remote servers; convert Server Core to/from full GUI; configure services; configure NIC teaming	Mod 1	Lesson 2/3	Mod 1 Ex 2/3
Configure local storage.	This objective may include but is not limited to: Design storage spaces; configure basic and dynamic disks; configure MBR and GPT disks; manage volumes; create and mount virtual hard disks (VHDs); configure storage pools and disk pools	Mod 4	Lesson 3	Mod 4 Ex 2/3
Configure Server Roles and Features				
Configure servers for remote management.	This objective may include but is not limited to: Configure WinRM; configure down-level server management; configure servers for day-to-day management tasks; configure multi-server management; configure Server Core; configure Windows Firewall	Mod 1	Lesson 1/2/3	Mod 1 Ex 1/2

Exam Objective Domains		Course Content		
Exam 70-410: Installing and Configuring Windows Server 2012 (continued)				
Configure Hyper-V				
Create and configure virtual machine settings.	This objective may include but is not limited to: Configure dynamic memory; configure smart paging; configure Resource Metering; configure guest integration services	Mod 8	Lesson 1/4	Mod 8 Ex 3
Create and configure virtual machine storage.	This objective may include but is not limited to: Create VHDs and VHDX; configure differencing drives; modify VHDs; configure pass-through disks; manage snapshots; implement a virtual Fibre Channel adapter	Mod 8	Lesson 2	Mod 8 Ex 2/3
Create and configure virtual networks.	This objective may include but is not limited to: Implement Hyper-V Network Virtualization; configure Hyper-V virtual switches; optimize network performance; configure MAC addresses; configure network isolation; configure synthetic and legacy virtual network adapters	Mod 8	Lesson 3	
Install and Administer Active Directory				
Install domain controllers.	This objective may include but is not limited to: Add or remove a domain controller from a domain; upgrade a domain controller; install Active Directory Domain Services (AD DS) on a Server Core installation; install a domain controller from Install from Media (IFM); resolve DNS SRV record registration issues; configure a global catalog server	Mod 11	Lesson 1/2	Mod 11 Ex 2/3
Exam 70-411: Administering Windows Server 2012				
Deploy, Manage, and Maintain Servers				
Monitor servers.	This objective may include but is not limited to: Configure Data Collector Sets (DCS); configure alerts; monitor real-time performance; monitor virtual machines (VMs); monitor events; configure event subscriptions; configure network monitoring	Mod 2	Lesson 1	Mod 2 Ex 1
Configure Network Services and Access				
Configure DirectAccess.	This objective may include but is not limited to: Implement server requirements; implement client configuration; configure DNS for Direct Access; configure certificates for Direct Access	Mod 6	Lesson 1/2	Mod 6 Ex 1/2/3

Exam Objective Domains		Course Content		
Exam 70-411: Administering Windows Server 2012 (continued)				
Configure a Network Policy Server Infrastructure				
Configure Network Access Protection (NAP).	This objective may include but is not limited to: Configure System Health Validators (SHVs); configure health policies; configure NAP enforcement using DHCP and VPN; configure isolation and remediation of non-compliant computers using DHCP and VPN; configure NAP client settings	Mod 5	Lesson 4	Mod 5 Ex 3
Configure and Manage Active Directory				
Configure Domain Controllers.	This objective may include but is not limited to: Configure Universal Group Membership Caching (UGMC); transfer and seize operations masters; install and configure a read-only domain controller (RODC); configure Domain Controller cloning	Mod 11	Lesson 1/2	Mod 11 Ex 1
Maintain Active Directory.	This objective may include but is not limited to: Back up Active Directory and SYSVOL; manage Active Directory offline; optimize an Active Directory database; clean up metadata; configure Active Directory snapshots; perform object- and container-level recovery; perform Active Directory restore	Mon 11	Lesson 5	
Configure and Manage Group Policy				
Configure Group Policy processing.	This objective may include but is not limited to: Configure processing order and precedence; configure blocking of inheritance; configure enforced policies; configure security filtering and WMI filtering; configure loopback processing; configure and manage slow-link processing; configure client-side extension (CSE) behavior	Mod 11	Lesson 4	Mod 11 Ex 2

Exam Objective Domains		Course Content		
Exam 70-412: Configuring Advanced Windows Server 2012 Services				
Configure and Manage High Availability				
Configure failover clustering.	This objective may include but is not limited to: Configure Quorum; configure cluster networking; restore single node or cluster configuration; configure cluster storage; implement Cluster Aware Updating; upgrade a cluster	Mod 7	Lesson 1/2/4	Mod 7 Ex 1/2/4
Manage failover clustering roles.	This objective may include but is not limited to: Configure role-specific settings including continuously available shares; configure VM monitoring; configure failover and preference settings	Mod 7	Lesson 3/4	Mod 7 Ex 2
Manage Virtual Machine (VM) movement.	This objective may include but is not limited to: Perform live migration; perform quick migration; perform storage migration; import, export, and copy VMs; migrate from other platforms (P2V and V2V)	Mod 8	Lesson 4	
		Mod 9	Lesson 3/4	Mod 9 Ex 3
Configure File and Storage Solutions				
Implement Dynamic Access Control (DAC).	This objective may include but is not limited to: Configure user and device claim types; implement policy changes and staging; perform access-denied remediation; configure file classification	Mod 10	Lesson 1/2/3	Mod 10 Ex 2/3/4/5
Implement Business Continuity and Disaster Recovery				
Configure and manage backups.	This objective may include but is not limited to: Configure Windows Server backups; configure Windows Online backups; configure role-specific backups; manage VSS settings using VSSAdmin; create System Restore snapshots	Mod 2	Lesson 2	Mod 2 Ex 2/3/4
Configure site-level fault tolerance.	This objective may include but is not limited to: Configure Hyper-V Replica including Hyper-V Replica Broker and VMs; configure multi-site clustering including network settings, Quorum, and failover settings	Mod 9	Lesson 1/3	Mod 9 Ex 1
Configure Network Services				
Deploy and manage IPAM.	This objective may include but is not limited to: Configure IPAM manually or by using Group Policy; configure server discovery; create and manage IP blocks and ranges; monitor utilization of IP address space; migrate to IPAM; delegate IPAM administration; manage IPAM collections	Mod 5	Lesson 2	Mod 5 Ex 2

Exam Objective Domains		Course Content		
Exam 70-412: Configuring Advanced Windows Server 2012 Services				
Configure Identity and Access Solutions				
Implement Active Directory Federation Services 2.1 (AD FSv2.1).	This objective may include but is not limited to: Implement claims-based authentication including Relying Party Trusts; configure Claims Provider Trust rules; configure attribute stores including Active Directory Lightweight Directory Services (AD LDS); manage AD FS certificates; configure AD FS proxy, Integration with Cloud Services	Mod 12	Lesson 1/2/3	Mod 12 Ex 1/2/3/4



**Important** Attending this course in itself will not successfully prepare you to pass any associated certification exams.

The taking of this course does not guarantee that you will automatically pass any certification exam. In addition to attendance at this course, you should also have the following:

- Experience with implementing, managing and administering a Windows Server 2008 and Windows Server 2008 R2 environment
- Knowledge equivalent to the MCSA: Windows Server 2008 credential
- Minimum of one to two years real world, hands-on experience Installing and configuring a Windows Server Infrastructure
- Additional study outside of the content in this handbook

There may also be additional study and preparation resources, such as practice tests, available for you to prepare for this exam. Details of these are available at the following URL:

<http://www.microsoft.com/learning/en/us/exam.aspx?ID=70-417&locale=en-us#tab3>

You should familiarize yourself with the audience profile and exam prerequisites to ensure you are sufficiently prepared before taking the certification exam. The complete audience profile for this exam is available at the following URL:

<http://www.microsoft.com/learning/en/us/exam.aspx?ID=70-417&locale=en-us#tab1>

The exam/course mapping table outlined above is accurate at the time of printing, however it is subject to change at any time and Microsoft bears no responsibility for any discrepancies between the version published here and the version available online and will provide no notification of such changes.

## Course Materials

The following materials are included with your kit:

- **Course Handbook** A succinct classroom learning guide that provides all the critical technical information in a crisp, tightly-focused format, which is just right for an effective in-class learning experience.
  - **Lessons:** Guide you through the learning objectives and provide the key points that are critical to the success of the in-class learning experience.
  - **Labs:** Provide a real-world, hands-on platform for you to apply the knowledge and skills learned in the module.
  - **Module Reviews and Takeaways:** Provide improved on-the-job reference material to boost knowledge and skills retention.
  - **Lab Answer Keys:** Provide step-by-step lab solution guidance at your fingertips when it's needed.
- **Course evaluation** At the end of the course, you have the opportunity to complete an online evaluation to provide feedback on the course, training facility, and instructor.
  - To provide additional comments or feedback on the course, send email to [support@mscourseware.com](mailto:support@mscourseware.com). To inquire about the Microsoft Certification Program, send email to [mcphelp@microsoft.com](mailto:mcphelp@microsoft.com).

## Virtual Machine Environment

This section provides the information for setting up the classroom environment to support the business scenario of the course.

### Virtual Machine Configuration

In this course, you will use Microsoft Hyper-V to perform the labs.

**Important** At the end of each lab, you must revert the virtual machines to a snapshot. You can find the instructions for this procedure at the end of each lab. For the Module 8 lab, you should leave the virtual machines running for the Module 9 lab.

The following table shows the role of each virtual machine used in this course:

Virtual machine	Role
20417A-LON-DC1	Domain controller that is running Windows Server 2012 in the Adatum.com domain
20417A-LON-SVR1	Windows Server 2012 server, member of Adatum.com domain
20417A-LON-SVR2	Windows Server 2012 server, member of Adatum.com domain
20417A-LON-SVR3	Windows Server 2012 server, member of Adatum.com domain
20417A-LON-SVR4	Windows Server 2012 server, member of Adatum.com domain
20417A-LON-SVR5	Server with blank vhd

Virtual machine	Role
20417A-LON-TMG	Threat Management Gateway server in Adatum.com domain
20417A-MUN-DC1	Domain controller that is running Windows Server 2012 in the TreyResearch.com
20417A-LON-CL1	Client computer running Windows® 8 and Office 2010 Service Pack 1 (SP1) in the Adatum.com domain
20417A-LON-CL2	Client computer running Windows 8 and Office 2010 SP1 in the Adatum.com domain

## Software Configuration

The following software is installed on each virtual machine:

- Windows Server 2012 RC
- Windows 8 RP

## Classroom Setup

Each classroom computer will have the same virtual machine configured in the same way.

## Course Hardware Level

To ensure a satisfactory student experience, Microsoft Learning requires a minimum equipment configuration for trainer and student computers in all Microsoft Certified Partner for Learning Solutions (CPLS) classrooms in which Official Microsoft Learning Product courseware are taught.

### Hardware Level 6

- Intel Virtualization Technology (Intel VT) or AMD Virtualization (AMD-V) processor
- Dual 120 gigabyte (GB) hard disks 7200 RM SATA or better\*
- 8 GB random access memory (RAM) or higher
- DVD drive
- Network adapter
- Super VGA (SVGA) 17-inch monitor
- Microsoft Mouse or compatible pointing device
- Sound card with amplified speakers

**MCT USE ONLY. STUDENT USE PROHIBITED**



# Module 1

## Installing and Configuring Servers Based on Windows Server 2012

### Contents:

Module Overview	1-1
Lesson 1: Installing Windows Server 2012	1-2
Lesson 2: Configuring Windows Server 2012	1-13
Lesson 3: Configuring Remote Management for Windows Server 2012 Servers	1-21
Lab: Installing and Configuring Servers Based on Windows Server 2012	1-25
Module Review and Takeaways	1-30

## Module Overview

Knowing the capabilities of the Windows Server® 2012 operating system enables you to use it effectively, and to take complete advantage of what it can offer your organization. Some of the many improvements to Windows Server 2012 include:

- Increased scalability and performance
- Virtualization features, such as Hyper-V Replica
- Improved Windows PowerShell® and scripting support
- High performance SMB 3.0 file shares

This module introduces you to Windows Server 2012, how to install it, how to perform post-installation configuration tasks, and how to configure it to support remote management.

### Objectives

After completing this module, you will be able to:

- Describe the installation requirements for Windows Server 2012.
- Configure Windows Server 2012.
- Configure Windows Remote Management.
- Install the Windows Server 2012 operating system on servers.

## Lesson 1

# Installing Windows Server 2012

You must have a firm understanding of your organization's requirements so that you can deploy the appropriate edition of Windows Server 2012. You must also understand which hardware configuration is appropriate for Windows Server 2012, whether a virtual deployment might be more suitable than a physical deployment, and which installation source enables you to deploy Windows Server 2012 efficiently.

This lesson provides an overview of the different Windows Server 2012 editions, hardware requirements, deployment options, and installation process.

### Lesson Objectives

After completing this lesson you will be able to:

- Describe the different editions of Windows Server 2012.
- Determine whether a particular hardware configuration is appropriate for Windows Server 2012.
- Explain how to perform a physical or a virtual deployment of Windows Server 2012.
- Select an appropriate installation source for a Windows Server 2012 deployment.
- Determine when you can upgrade and when you must migrate to Windows Server 2012.
- Decide between a Server Core installation and full installation.
- Install Windows Server 2012.
- Perform post-installation configuration tasks.

### Windows Server 2012 Editions

There are several editions of Windows Server 2012. Organizations can select the edition of Windows Server 2012 that best meets their needs. Systems Administrators can save costs by selecting the appropriate edition when deploying a server for a specific role. The editions of Windows Server 2012 are listed in the following table.

- Windows Server 2012 Standard edition
- Windows Server 2012 Datacenter edition
- Windows Server 2012 Foundation edition
- Windows Server 2012 Essentials
- Microsoft Hyper-V Server 2012
- Windows Storage Server 2012 Workgroup
- Windows Storage Server 2012 Standard
- Windows MultiPoint Server 2012 Standard
- Windows MultiPoint Server 2012 Premium

Edition	Description
Windows Server 2012 Standard edition	<ul style="list-style-type: none"> <li>• Provides all roles and features available on the Windows Server 2012 platform.</li> <li>• Supports up to 64 sockets and up to 4 terabytes (TB) of RAM.</li> <li>• Includes 2 virtual machine licenses.</li> </ul>

Edition	Description
Windows Server 2012 Datacenter edition	<ul style="list-style-type: none"> <li>• Provides all roles and features that are available on the Windows Server 2012 platform.</li> <li>• Supports 64 sockets, up to 640 processor cores, and up to 4 terabytes of RAM.</li> <li>• Includes unlimited virtual machine licenses for virtual machines run on the same hardware.</li> </ul>
Windows Server 2012 Foundation edition	<ul style="list-style-type: none"> <li>• Allows only 15 users and cannot be joined to a domain.</li> <li>• Supports one processor core and up to 32 GB of RAM.</li> <li>• Includes limited server roles.</li> </ul>
Windows Server 2012 Essentials	<ul style="list-style-type: none"> <li>• Serves as the next edition of Small Business Server.</li> <li>• Cannot function as a Hyper-V, failover clustering, server core, or remote desktop services server.</li> <li>• Supports up to 25 users, 50 devices.</li> <li>• Supports 2 processor cores and 64 GB of RAM.</li> <li>• Must be root server in domain.</li> </ul>
Microsoft Hyper-V Server 2012	<ul style="list-style-type: none"> <li>• Stand-alone Hyper-V platform for virtual machines with no UI.</li> <li>• No licensing cost for host OS, virtual machines to be licensed normally.</li> <li>• Supports 64 sockets and 4 TB of RAM.</li> <li>• Supports domain join.</li> <li>• Does not support other Windows Server 2012 roles other than limited file services features.</li> </ul>
Windows Storage Server 2012 Workgroup	<ul style="list-style-type: none"> <li>• Entry-level unified storage appliance.</li> <li>• Supports up to 50 users.</li> <li>• Supports one processor core, 32 GB of RAM.</li> <li>• Supports domain join.</li> </ul>
Windows Storage Server 2012 Standard	<ul style="list-style-type: none"> <li>• Supports 64 sockets, but is licensed on a 2 socket incrementing basis.</li> <li>• Supports 4 TB of RAM.</li> <li>• Includes 2 virtual machine licenses.</li> <li>• Supports domain join.</li> <li>• Supports some roles, including DNS and DHCP Server roles, but does not support others, including Active Directory® Domain Services (AD DS), Active Directory Certificate Services (AD CS), and Active Directory Federation Services (AD FS).</li> </ul>

Edition	Description
Windows MultiPoint Server 2012 Standard	<ul style="list-style-type: none"> <li>• Supports multiple users accessing the same host computer directly using separate mouse, keyboard, and monitors.</li> <li>• Supports one socket, 32 GB of RAM and a maximum of 12 sessions.</li> <li>• Supports some roles, including DNS and DHCP Server roles, but does not support others including, AD DS, AD CS, and AD FS.</li> <li>• Does not support domain join.</li> </ul>
Windows MultiPoint Server 2012 Premium	<ul style="list-style-type: none"> <li>• Supports multiple users accessing the same host computer directly using separate mouse, keyboard, and monitors.</li> <li>• Limited to 2 sockets, 4 TB of RAM and a maximum of 22 sessions.</li> <li>• Supports some roles, including DNS and DHCP Server roles, but does not support others, including AD DS, AD CS, and AD FS.</li> <li>• Supports domain join.</li> </ul>



**Additional Reading:** For more information about the differences between Windows Server 2012 editions, see <http://www.windowsservercatalog.com/svvp.aspx>.

## Hardware Requirements for Installing Windows Server 2012

Hardware requirements define the absolute minimum required to run the server software. The actual hardware requirements depend on the services that the server is hosting, the load on the server, and how responsive you want the server to be.

The services and features of each role put a unique load on network, disk I/O, processor, and memory resources.

Virtualized deployments of Windows Server 2012 must match the same hardware specifications as physical deployments. Windows Server 2012 is supported on Hyper-V® and certain third-party virtualization platforms.

Windows Server 2012 has the following minimum hardware requirements:

- Processor architecture x86-64
- Processor speed 1.4 GHz
- Memory (RAM) 512 MB
- HDD space 32 GB
  - More if the server has more than 16 GB of RAM



The minimum hardware requirements for Windows Server 2012 are shown in the following table.

Component	Requirement
Processor architecture	x86-64
Processor speed	1.4 GHz
Memory (RAM)	512 MB
Hard disk drive space	32 GB, or more if the server has more than 16 GB of RAM



**Additional Reading:** For more information about the Windows Server Virtualization Validation Program, see <http://www.windowsservercatalog.com/svvp.aspx>.

## Considerations for Deploying Physical or Virtual Machines

With virtualization you can be more efficient in the way that you allocate resources to servers. Instead of allocating separate hardware to a server that minimally uses resources, you can virtualize that server and enable those minimally used hardware resources to be shared with other virtual machines.

When determining whether to deploy a server physically or virtually, you must determine how that server uses hardware resources. Consider these points:

- Servers that constantly put hardware under resource pressure are poor candidates for virtualization. This is because virtual machines share resources. A single virtual machine that uses a disproportionate amount of hypervisor resources can have an adverse effect on other virtual machines hosted on the same hypervisor.
- Servers that put minimal pressure on hardware resources are good candidates for virtualization. These servers are unlikely to monopolize the host resources, ensuring that each virtual machine hosted on the hypervisor can access enough hardware resources to perform adequately.

For example, a particular database server that heavily uses disk and network resources would be better deployed on a physical computer. If it were deployed as a virtual machine, other virtual machines on the same hypervisor would have to compete for access to those heavily-used disk and network resources. Alternatively, allocating a physical platform to a server that requires minimal hardware resources, such as a server running Certificate Services, means that powerful hardware is underused.

Other things to consider when determining whether to deploy a server virtually or physically are:

- **High Availability.** After you have built a highly available virtual machine cluster, any virtual machine deployed to that cluster also becomes highly available. This is simpler than setting up separate failover clusters for physical servers that host the same role.

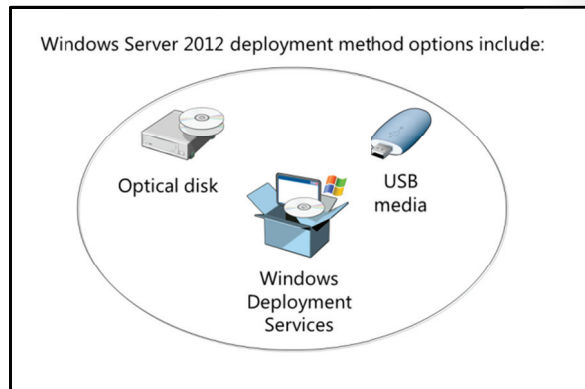
Deploy Physically	Deploy Virtually
Server software puts hardware under resource pressure.	Server uses minimal hardware resources.
Server role requires exclusive access to hardware resources.	Server configured for high availability.
	Server used for scale-up as a new host.

- **Scalability.** Moving a virtual machine with its associated applications and data to a new host platform is significantly simpler than migrating a physically deployed server, its applications, and data to a new host platform. If you must quickly scale-up capacity, you can also migrate a virtual machine to a cloud provider, something that is far more difficult to do with a physically deployed server.

## Windows Server 2012 Installation Sources

Microsoft distributes Windows Server 2012 either on optical media or in an .iso image format.

You can install Windows Server 2012 by using several methods, including those listed in the following table.



Method	Notes
Optical media	<ul style="list-style-type: none"> <li>• Requires that the computer has access to a DVD drive.</li> <li>• Optical media is usually slower than USB media.</li> <li>• You cannot update the installation image without replacing the media.</li> <li>• You can only perform one installation per DVD at a time.</li> </ul>
USB media	<ul style="list-style-type: none"> <li>• Requires the administrator to perform special steps to prepare USB media from ISO file.</li> <li>• All computers support booting from USB media.</li> <li>• Image can be updated as new software updates and drivers become available.</li> <li>• Answer file can be stored on USB drive, reducing the interaction that the administrator must perform.</li> </ul>
Mounted ISO image	<ul style="list-style-type: none"> <li>• Virtualization software enables you to directly mount the ISO image.</li> <li>• Does not require writing the ISO image to optical media.</li> </ul>
Network share	<ul style="list-style-type: none"> <li>• Deploy from installation files on network share.</li> <li>• Requires you boot the server off a boot device (DVD or USB drive) and install from installation files hosted on a network share.</li> <li>• Much slower than using Windows Deployment Services (WDS).</li> <li>• If you already have access to a DVD or USB media, it is simpler to use those tools for operating system deployment.</li> </ul>
Windows Deployment Services (WDS)	<ul style="list-style-type: none"> <li>• WDS let you deploy Windows Server 2012 from Windows Imaging Format (WIM) image files or specially prepared VHD files.</li> <li>• You can use the Windows Automated Installation Kit to configure lite-touch deployment.</li> </ul>

Method	Notes
	<ul style="list-style-type: none"> <li>• Clients perform a Pre-Boot Execution Environment (PXE) boot to contact the WDS server. The operating system image is then transmitted to the server over the network.</li> <li>• WDS supports multiple concurrent installations of Windows Server 2012 using multicast network transmissions.</li> </ul>
System Center Configuration Manager	<ul style="list-style-type: none"> <li>• Microsoft® System Center Configuration Manager enables you to fully automate the deployment of Windows Server 2012 to “bare metal” servers.</li> <li>• Enables Zero Touch deployment.</li> </ul>
Virtual Machine Manager templates	<ul style="list-style-type: none"> <li>• Requires Virtual Machine Manager (VMM) in System Center.</li> <li>• Enables rapid deployment of Windows Server 2012 in private cloud scenarios.</li> <li>• Can be used to enable self-service deployment of Windows Server 2012 virtual machines.</li> </ul>

Microsoft distributes Windows Server 2012 either on optical media or in an .iso image format.

You can install Windows Server 2012 by using several methods, including those listed in the following table.

Method	Notes
Optical media	<ul style="list-style-type: none"> <li>• Requires that the computer has access to a DVD drive.</li> <li>• Optical media is usually slower than USB media.</li> <li>• You cannot update the installation image without replacing the media.</li> <li>• You can only perform one installation per DVD at a time.</li> </ul>
USB media	<ul style="list-style-type: none"> <li>• Requires the administrator to perform special steps to prepare USB media from ISO file.</li> <li>• All computers support booting from USB media.</li> <li>• Image can be updated as new software updates and drivers become available.</li> <li>• Answer file can be stored on USB drive, reducing the interaction that the administrator must perform.</li> </ul>
Mounted ISO image	<ul style="list-style-type: none"> <li>• Virtualization software enables you to directly mount the ISO image.</li> <li>• Does not require writing the ISO image to optical media.</li> </ul>
Network share	<ul style="list-style-type: none"> <li>• Deploy from installation files on network share.</li> <li>• Requires you boot the server off a boot device (DVD or USB drive) and install from installation files hosted on a network share.</li> <li>• Much slower than using Windows Deployment Services (WDS).</li> <li>• If you already have access to a DVD or USB media, it is simpler to use those tools for operating system deployment.</li> </ul>

Method	Notes
Windows Deployment Services (WDS)	<ul style="list-style-type: none"> <li>• WDS let you deploy Windows Server 2012 from Windows Imaging Format (WIM) image files or specially prepared VHD files.</li> <li>• You can use the Windows Automated Installation Kit to configure lite-touch deployment.</li> <li>• Clients perform a Pre-Boot Execution Environment (PXE) boot to contact the WDS server. The operating system image is then transmitted to the server over the network.</li> <li>• WDS supports multiple concurrent installations of Windows Server 2012 using multicast network transmissions.</li> </ul>
System Center Configuration Manager	<ul style="list-style-type: none"> <li>• Microsoft® System Center Configuration Manager enables you to fully automate the deployment of Windows Server 2012 to “bare metal” servers.</li> <li>• Enables Zero Touch deployment.</li> </ul>
Virtual Machine Manager templates	<ul style="list-style-type: none"> <li>• Requires Virtual Machine Manager (VMM) in System Center.</li> <li>• Enables rapid deployment of Windows Server 2012 in private cloud scenarios.</li> <li>• Can be used to enable self-service deployment of Windows Server 2012 virtual machines.</li> </ul>

## Options for Upgrading and Migrating to Windows Server 2012

When considering whether to upgrade or migrate a server to Windows Server 2012, consider the options described in the following table.

### Upgrading to Windows Server 2012:

- Can upgrade from x64 editions of Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2.
- Can only upgrade to same or higher edition.
- Requires same processor architecture.

### Migrating to Windows Server 2012:

- Must migrate from x86 version of Windows Server.
- Can use Windows Server Migration Tools feature.

Installation option	Description
Upgrade	An upgrade preserves the files, settings, and applications installed on the original server. You perform an upgrade when you want to keep all these items and want to continue using the same server hardware. Upgrade requires an x64 processor architecture and an x64 edition of the Windows Server operating system. You can only upgrade to Windows Server 2012 from x64 versions of Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, and Windows Server 2008 R2. You can only upgrade to an equivalent or a later edition of Windows Server 2012. You start an upgrade by running Setup.exe from the original operating system.



Installation option	Description
Migration	Use migration when you migrate from an x86 version of Windows Server 2003, Windows Server 2003 R2, or Windows Server 2008. Use migration when you want to replace the original server with one running an earlier edition, for example replacing Windows Server 2008 R2 Enterprise edition with Windows Server 2012 Standard edition. You can use the Windows Server Migration Tools feature in Windows Server 2012 to transfer files and settings from computers running the Windows Server 2003, Windows Server 2003 R2, Windows Server 2008, Windows Server 2008 R2 and Windows Server 2012 operating systems.

## Choosing Between Server Core and Full Installation

Server Core is a minimal installation option for Windows Server 2012. With Server Core, you perform management tasks locally from the command-line or remotely from another computer. Server Core is the default installation option for Windows Server 2012. Server Core has the following advantages over a traditional deployment of Windows Server 2012:

- Reduced update requirements. Because Server Core installs fewer components, Server Core deployments require the application of fewer software updates. This reduces the time that is required for an administrator to service Server Core.
- Reduced hardware footprint. Server Core computers require less RAM and less hard disk space. This means that when virtualized, more servers can be deployed on the same host.

### Server Core in Windows Server 2012:

- Default minimal installation option
- Reduces update requirements
- Requires fewer hardware resources
- Supports most Windows Server 2012 roles and features
- Installation options:
  - Server Core standard deployment
  - Server Core with management

Increasing numbers of Microsoft server applications are designed to run on computers that have Server Core installations. Microsoft SQL Server® 2012 can be installed on computers running the Server Core version of Windows Server 2008 R2.

There are two options for installing the Server Core, as described in the following table.

Option	Description
Server Core	This is the standard deployment of Server Core. By default all graphical administration components are in a Removed state. Simply stated, Removed components occupy no disk space on the server. Server Core systems are managed locally by using command-line interface only, or can be managed by a remote system using graphical administration tools. You can convert to the full version of Windows Server 2012 that includes the graphical administration components only if you have access to an installation source with all server files, such as a mounted WIM image. Any Server Core component in a Removed state can only be installed by using an installation source.
Server Core with Management	This is also known as Server Core-Full Server. This works the same as a deployment of Windows Server 2012 with the graphical components. With this installation option the graphical administration components are not in a Removed state. Instead, these components are available (they are located on the server's disk), but not installed into the OS. You can convert between Server Core with Management and Windows Server 2012 with a graphical interface by installing the graphical features, but without having to specify an installation source.

On a local connection, you can use the tools described in the following table to manage Server Core installations of Windows Server 2012.

Tool	Function
Cmd.exe	Enables you to run traditional command-line utilities, such as ping.exe, ipconfig.exe, and netsh.exe.
PowerShell.exe	Enables you to start a Windows PowerShell session on the Server Core deployment. You can then perform Windows PowerShell tasks as usual.
Sconfig.cmd	Command-line menu driven administrative tool that enables you to perform most common server administrative tasks.
Notepad.exe	Enables you to use the Notepad.exe Text Editor in the Server Core environment.
Registry Editor	Provides registry access within the Server Core environment.
Msiinfo32.exe	Enables you to view system information about the server core deployment.
Taskmgr.exe	Starts the Task Manager.



**Note:** If you accidentally close the Command Prompt window on a computer running Server Core, you can restore it using this procedure:

1. Press Ctrl+Alt+Delete.
2. On the menu, click **Task Manager**.
3. On the **File** menu, click **New Task (Run...)**.
4. Type **cmd.exe** and then press Enter.

Server Core supports most, but not all, Windows Server 2012 roles and features. You cannot install the following roles on a computer running Server Core:

1. AD FS
2. Application Server
3. Network Policy and Access Services
4. Windows Deployment Services

Even if a role is available to a computer running the Server Core installation option, a specific role service associated with that role may not be.



**Note:** You can check which roles are not available on Server Core by running the following query.

```
Get-WindowsFeature | where-object {$_.InstallState -eq Removed}
```

The Windows Server 2012 administration model focuses on managing many servers from one console instead of the traditional method of managing each server separately. When you want to perform an administrative task, you are more likely to manage multiple computers running the Server Core operating system from one computer than you are to connect to each computer individually. You can enable remote management of a computer running Server Core by using **sconfig.cmd** or by executing the command:

```
Netsh.exe firewall set service remoteadmin enable ALL
```

## Installation Process for Windows Server 2012

In a typical installation of Windows Server 2012, if you do not have an existing answer file, you perform the following steps:

1. Connect to the installation source. Some options for this include:
  - Inserting a DVD-ROM that has the Windows Server 2012 installation files and booting from the DVD-ROM.
  - Connecting a USB drive that is made bootable and contains a copy of the Windows Server 2012 installation files.
  - Performing a PXE boot from the computer that Windows Server 2012 will be installed on to, and connecting to a WDS server.
2. On the first page of the Windows Setup Wizard, select the following:
  - Language to install
  - Time and currency format
  - Keyboard or input method
3. On the second page of the Windows Setup Wizard, click **Install now**. You can also use this page to select **Repair Your Computer**. Use this option if an installation has become corrupted and you can no longer boot into Windows Server 2012.
4. On the **Select The Operating System You Want To Install** page of the Windows Setup Wizard, select from the available operating system installation options. The default option is Server Core installation.
5. On the **License Terms** page of the Windows Setup Wizard, review the terms of the operating system license. You must accept the license terms before you can continue with the installation process.
6. On the **Which Type Of Installation Do You Want** page of the Windows Setup Wizard, you have the following options:
  - **Upgrade**. Select this option if you have an existing Windows Server installation that you want to upgrade to Windows Server 2012. You should start upgrades from the earlier version of Windows Server instead of booting from the installation source.
  - **Custom**. Select this option if you want to perform a new installation.
7. On the **Where do you want to install Windows** page of the Windows Setup Wizard, select an available disk on which to install Windows. You can also choose to repartition and reformat disks from this page. When you click **Next**, the installation process will copy files and restart the computer several times. This part of the installation can take several minutes, depending on the speed of the platform on which you are installing Windows Server 2012.

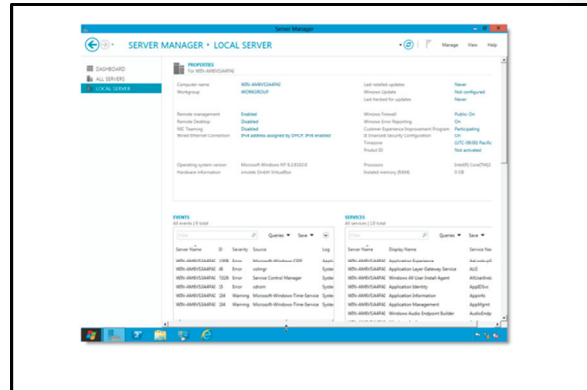
8. On the **Settings** page, provide a password for the local Administrator account. After you have provided this password, you can log on to the server and begin performing post installation configuration tasks.

## Post-Installation Tasks

In earlier versions of Windows operating systems, the installation required you to configure network connections, computer name, user account, and domain membership information. The Windows Server 2012 installation process reduces the number of questions that you have to answer. The only information that you provide during installation is the password that is used by the default local Administrator account.

After it is installed, all the following steps can be performed when you select the **Local Server** node in the Server Manager console:

- Configure the IP address
- Set the computer name
- Join an Active Directory domain
- Configure the time zone
- Enable automatic updates
- Add roles and features
- Enable remote desktop
- Configure Windows Firewall settings



## Lesson 2

# Configuring Windows Server 2012

By correctly configuring a server first, you can avoid significant problems later. When planning to configure a server, you must determine what roles to deploy. You must also assess whether roles can be co-located on the same server or if you deploy certain roles on separate servers.

### Lesson Objectives

After completing this lesson you will be able to:

- Describe Windows Server 2012 server roles.
- Install roles and use the Best Practice Analyzer to check role configuration.
- Configure a computer running the Server Core installation option.
- Switch a computer between Server Core and the full GUI installation option.
- Configure networking and network interface teaming.

### Demonstration: Exploring Server Manager in Windows Server 2012

In this demonstration, you will see how to use Server Manager to perform the following tasks:

- Log on to Windows Server 2012.
- View the Windows Server 2012 desktop.
- Start the Server Manager console.
- Add a server role or feature.
- View role related events.
- Run the Best Practice Analyzer for a role.
- List the tools available from Server Manager.
- Open the **Start** menu.
- Log off the currently logged on user.
- Restart Windows Server 2012.

### Demonstration Steps

1. On LON-DC1, open the Add Roles and Features Wizard from the Server Manager Console.
2. Start the Add Roles and Features Wizard and select the following options:
  - Role-based or feature-based installation
  - LON-DC1
  - FAX Server role
  - BranchCache feature
3. Use the notification area to review the messages.
4. On the Dashboard, view DNS Events.

5. Configure the DNS - Events Detail View with the following settings:
  - Time period: **12 hours**
  - Event Sources: **All**
6. View the DNS Best Practice Analyzer (BPA) with the following settings:
  - Severity Levels: **All**
7. Use the **Tools** menu to view the tools that are installed on LON-DC1.
8. Demonstrate log off LON-DC1 and then log back on.
9. Open Windows PowerShell and then use the shutdown command to shut the server down.

## Server Roles in Windows Server 2012

Roles and their associated Role Services are still a primary function of a server. Similarly, if you install the Web Server (IIS) role, Windows Server 2012 by default only selects critical services that are required for the role to function. If you want to use additional components with the Web Server (IIS) role, such as Windows Authentication, you must select and install that component as a role service.

Windows Server 2012 supports the roles described in the following table.

Windows Server 2012 supports these roles:

- |   |                                      |
|---|--------------------------------------|
| • Active Directory Certificate Services           | • Fax Server                         |
| • Active Directory Domain Services (AD DS)        | • File and Storage Services          |
| • Active Directory Federation Services            | • Hyper-V                            |
| • Active Directory Lightweight Directory Services | • Network Policy and Access Services |
| • Active Directory Rights Management Services     | • Print and Document Services        |
| • Application Server                              | • Remote Access                      |
| • DHCP Server                                     | • Remote Desktop Services            |
| • DNS Server                                      | • Volume Activation Services         |
|   | • Web Server (IIS)                   |
|   | • Windows Deployment Services        |
|   | • Windows Server Update Services     |

Role	Function
Active Directory Certificate Services	Enables the deployment of certification authorities and related role services.
AD DS	Centralized store of information about network objects including user and computer accounts. Used for authentication and authorization.
AD FS	Provides web single sign-on (SSO) and secured identify federation support.
Active Directory Lightweight Directory Services (AD LDS)	Supports storage of application specific data for directory-aware applications that do not require the full infrastructure of AD DS.
Active Directory Rights Management Services(AD RMS)	Enables you to prevent unauthorized access to sensitive documents by applying rights management policies.
Application Server	Supports centralized management and hosting of high-performance distributed business applications, such as those built with the .NET Framework 4.5 and Enterprise Services.
DHCP Server	Provisions client computers on the network with temporary IP addresses.
DNS Server	Provides name resolution for TCP/IP networks.

Role	Function
Fax Server	Supports sending and receiving of faxes. Also enables you to manage fax resource on the network.
File and Storage Services	Supports the storage of management of shared folders, Distributed File System, and network storage.
Hyper-V	Enables you to host virtual machines on computers running Windows Server 2012.
Network Policy and Access Services	Authorization infrastructure for remote connections, including Health Registration Authority for Network Access Protection.
Print and Document Services	Supports centralized management of document tasks, including network scanners and networked printers.
Remote Access	Supports Seamless Connectivity, Always On, Always Managed features based on DirectAccess. Also supports Remote Access through VPN and dial-up.
Remote Desktop Services	Supports access to virtual desktops, session-based desktops, and RemoteApp programs.
Volume Activation Services	New to Windows Server 2012. Enables you to automate and simplify the management of volume license keys and volume key activation. Also enables you to manage a Key Management Service host or configure AD DS-based activation for computers that are members of the domain.
Web Server (IIS)	The Windows Server 2012 web server component.
Windows Deployment Services	Enables you to deploy server operating systems to clients over the network.
Windows Server Update Services	Provides a method of deploying updates for Microsoft products to computers on the network.

When you deploy a role, Windows Server 2012 automatically configures aspects of the server's configuration, such as firewall settings, to support the role. When you deploy a role, Windows Server 2012 automatically deploys role dependencies at the same time. For example, when you install the Windows Server Update Services role, Windows Server 2012 installs the Web Server (IIS) role components that are required to support the Web Server role.

You add and remove roles using the Add Roles and Features Wizard, available from the Server Manager console. You can also add and remove roles using the **Install-WindowsFeature** and **Remove-WindowsFeature** Windows PowerShell cmdlets.

## Demonstration: Installing and Optimizing Server Roles in Windows Server 2012

In this demonstration you will see how to install and optimize a server role in Windows Server 2012.

### Demonstration Steps

1. Use the Add Roles and Features Wizard to add the **Application Server** role to LON-DC1.
2. View App Server Performance.
3. View DHCP BPA results.

## Configuring Server Core in Windows Server 2012

You must perform several aspects of post-installation configuration of server core operating systems from the command-line. You can perform most post-installation configuration tasks using the menu-driven command prompt utility **sconfig.cmd**. By using this utility, you minimize the possibility of the Administrator making syntax errors when you use more complex command-line utilities. You can use **sconfig.cmd** to perform the following tasks:

- Configure Domain and Workgroup information
- Configure the computer's name
- Add local Administrator accounts
- Configure Remote Management
- Enable Windows Update
- Download and install updates
- Enable Remote Desktop
- Configure Network Address information
- Set the date and time
- Perform Windows Activation
- Enable the Graphic User Interface
- Log off
- Restart the server
- Shut down the server

- Use **sconfig.cmd** for common administrative tasks
- Use Windows PowerShell cmdlets for more complex administrative tasks

### Configure IP Address Information

You can configure the IP address and DNS information by using **sconfig.cmd** or **netsh.exe**. To configure IP address information by using **sconfig.cmd**, perform the following steps:

1. Run **sconfig.cmd** from the command-line.
2. Select **option 8** to configure Network Settings.



3. Select the index number of the network adapter to which you want to assign an IP address.
4. In the Network Adapter Settings area, select between one of the following options:
  - **Set Network Adapter Address**
  - **Set DNS Servers**
  - **Clear DNS Server Settings**
  - **Return to Main Menu**

### Change Server Name

You can change the server name using the **netdom** command with the **renamecomputer** option. For example, to rename a computer to Melbourne, type the following command:

```
Netdom renamecomputer %computername% /newname:Melbourne
```

You can change a server name using **sconfig.cmd** by performing the following steps:

1. Run **sconfig.cmd** from the command-line.
2. Select **option 2** to configure the computer name.
3. Type the new computer name and then press Enter.

You must restart a server for the configuration change to take effect.

### Joining the Domain

You can join a Server Core computer to a domain using the **netdom** command with the join option. For example, to join the adatum.com domain using the Administrator account, and to be prompted for a password, issue the command:

```
Netdom join %computername% /domain:adatum.com /UserD:Administrator /PasswordD:*
```

To join a server core computer to the domain using **sconfig.cmd**, perform the following steps:

1. Run **sconfig.cmd** from the command-line.
2. Select **option 1** to configure Domain/Workgroup.
3. Type **D** and press Enter to select the Domain option.
4. Type the name of the domain to which you want to join the computer.
5. Provide the details of an account authorized to join the domain in domain\username format.
6. Type the password associated with that account.

To complete a domain join operation you must restart the computer.



**Note:** Before joining the domain, verify that you can ping the DNS server by host name.

### Add Roles and Features Using Windows PowerShell

You can add and remove roles and features to a computer running the Server Core installation option by using the **Get-WindowsFeature**, **Install-WindowsFeature**, and **Remove-WindowsFeature** Windows PowerShell cmdlets. These cmdlets are available after you load the Server Manager module.

For example, you can view a list of roles and features that are installed by executing the following Windows PowerShell command:

```
Get-WindowsFeature | Where-Object {$_.InstallState -eq "Installed"}
```

You can install a Windows role or feature using the **Install-WindowsFeature** cmdlet. For example, to install the Network Load Balancing feature, execute the command:

```
Install-WindowsFeature NLB
```

Not all features are directly available for installation on a computer running the Server Core operating system. You can determine which features are not directly available for installation by running the following command:

```
Get-WindowsFeature | Where-Object {$_.InstallState -eq Removed}
```

You can add a role or feature that is not available for installation by using the **-Source** parameter of the **Install-WindowsFeature** cmdlet. You must specify a source location that hosts a mounted installation image that includes the full version of Windows Server 2012. You can mount an installation image using the **DISM.exe** command prompt utility.

## Switching Between Server Core, Full, and Minimal Server Interface Options

Windows Server 2012 offers the option of switching between Server Core and the full installation. When you install Server Core, the necessary components to convert to the full version are not installed. You can install these if you have access to a mounted image of the full version of the Windows Server 2012 installation files.

You can switch from Server Core to the graphical version of Windows Server 2012 by running the following Windows PowerShell cmdlet, where **c:\mount** is the root directory of a mounted image that hosts the full version of the Windows Server 2012 installation files:

```
Import-Module ServerManager
Install-WindowsFeature -IncludeAllSubFeature User-Interfaces-Infra -Source c:\mount
```

This gives you the option of performing administrative tasks using the graphical tools. You can also add the graphical tools using the **sconfig.cmd** menu-driven command prompt utility.

After you have performed the necessary administrative tasks, you can return the computer to its original Server Core configuration. You can switch a computer that has the graphical version of Windows Server 2012 to Server Core by removing the following features:

- Graphical Management Tools and Infrastructure
- Server Graphical Shell

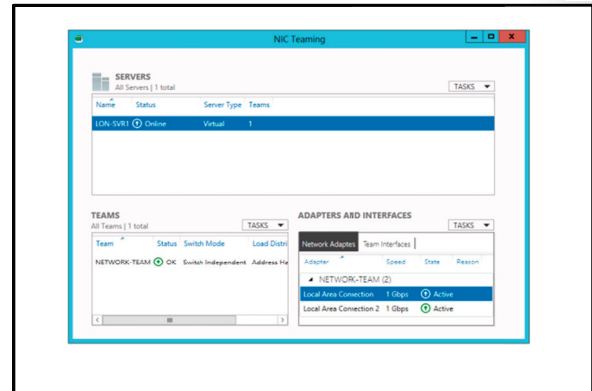
- You must have access to an installation source to convert from Server Core to the full version unless you have already installed the necessary components.
- You can install Server Core from **sconfig.cmd**
- You can remove the full version from Windows Server 2012 by removing the following features:
  - Graphical Management Tools and Infrastructure
  - Server Graphical Shell

The Minimal Server interface differs from Server Core because it has all components available and does not require you to provide access to a mounted directory that contains the full version of the Windows Server 2012 installation files. You can use the **Install-WindowsFeature** command without specifying a source location when you convert the Minimal Server interface to the full installation of Windows Server 2012. The advantage of the Server Core installation option over Minimal Server is that, even though they look similar, Server Core requires a smaller amount of hard disk space as it does not have all components available for installation.

## Configuring Networking and Network Interface Teaming

Configuring the network involves setting or verifying the server's IP address configuration. By default, a newly-deployed server tries to obtain IP address information from a DHCP server. You can view a server's IP address configuration by clicking the **Local Server** node in Server Manager.

If the server has an IPv4 address in the Automatic Private Internet Protocol Addressing (APIPA) range of 169.254.0.1 to 169.254.255.254, the server has not been configured with an IP address from a DHCP server. This may be because a DHCP server has not been configured on the network, or because there is a problem with the network infrastructure that blocks the adapter from receiving an address.



**Note:** If you are using a purely IPv6 network, an IPv4 address in this range is not a problem, and IPv6 address information is still configured automatically. You will learn more about implementing IPv6 in Module 8, "Implementing IPv6."

### Configuration Using Server Manager

To manually configure IP address information for a server, perform the following steps:

1. In the Server Manager console, click the address next to the network adapter that you want to configure. This will open the Network Connections window.
2. Right-click the network adapter that you want to configure an address for, and then click **Properties**.
3. In the **Adapter Properties** dialog box, click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
4. In the **Internet Protocol Version 4 (TCP/IPv4) Properties** dialog box, enter the following IPv4 address information, and then click **OK**, and then click **OK** again:
  - IP address
  - Subnet Mask
  - Default Gateway
  - Preferred DNS server
  - Alternative DNS server

## Command-Line IPv4 Address Configuration

You can manually set IPv4 address information from an elevated command prompt by using the **netsh.exe** command from the interface ipv4 context. For example, to configure the adapter named **Local Area Connection** with the IPv4 address 10.10.10.10 and subnet mask 255.255.255.0, type the following command:

```
Netsh interface ipv4 set address "Local Area Connection" static 10.10.10.10 255.255.255.0
```

You can use the same context of the netsh.exe command to configure DNS configuration. For example, to configure the adapter named **Local Area Connection** to use the DNS server at IP address 10.10.10.5 as the primary DNS server, type the following command:

```
Netsh interface ipv4 set dnsservers "Local Area Connection" static 10.10.10.5 primary
```

## Network Card Teaming

Network Card Teaming is a new feature in Windows Server 2012. With Network Card Teaming you can increase the availability of a network resource. When you configure Network Card Teaming, a computer uses one network address for multiple cards. If one of the cards fails, the computer continues communicating with other hosts on the network that are using that shared address. This enables you to provide hardware redundancy for a server's network cards. Network Card Teaming does not require that the network cards be the same model or use the same driver.

Windows Server 2012 supports up to 32 network adapters in a team. When a computer has separate network adapters that are not part of a team, incoming and outgoing traffic may not be balanced across those adapters. Network Card Teaming also provides bandwidth aggregation, ensuring that traffic is balanced across network interfaces as a way to increase effective bandwidth.

To team network cards, perform the following steps:

1. Ensure that the server has more than one network adapter.
2. In Server Manager, click the **Local Server** node.
3. Click **Disabled** next to Network Adapter Teaming. This opens the **NIC Teaming** dialog box.
4. In the **NIC Teaming** dialog box, press the Ctrl key, and then click each network adapter that you want to add to the team.
5. Right-click these selected network adapters, and then click **Add to New Team**.
6. In the **New Team** dialog box, enter a name for the team, and then click **OK**.

## Lesson 3

# Configuring Remote Management for Windows Server 2012 Servers

When you want to perform an administration task, it is more efficient to manage multiple servers from a single console than to connect to each server separately. You should spend time ensuring that newly deployed servers are configured so that you can manage them centrally. This enables you to spend more time at your desk administering those servers, instead of having to trek into the datacenter to start a direct connection.

### Lesson Objectives

After completing this lesson you will be able to:

- Describe the different Windows Server 2012 remote management technologies.
- Configure Windows Server 2012 to support Remote Management.
- Collect servers into Server Groups.
- Deploy roles and features remotely.

### What Is Remote Management?

With Windows Remote Management, you can use Remote Shell, remote Windows PowerShell, and remote management tools to remotely manage a computer. Remote Shell enables you to run command-line utilities against correctly configured remote servers as long as the command prompt utility is present on the remote server. Remote Windows PowerShell lets you run Windows PowerShell commands or scripts against correctly configured remote servers when the script is hosted on the local server. Remote Windows PowerShell also lets you load Windows PowerShell modules, such as Server Manager locally and execute the cmdlets available in that module against suitably configured remote servers. Remote Management is enabled by default on computers running Windows Server 2012.

You can enable Remote Management using:

- Server Manager console
- WinRM -qc
- Sconfig.cmd

Use Remote Desktop when administrators are using Windows XP or third-party operating systems to manage servers running Windows Server 2012.

You can enable and disable Remote Management from Server Manager by clicking the text next to the Remote Management item when you have the Local Server node selected in the Server Manager console.

To enable remote management from the command-line, type the command **WinRM qc**. The "qc" is an abbreviation of Quick Configuration. You can disable Remote Management by using the same method that you use to enable it.

To disable remote management on a computer running the Server Core installation option, use **sconfig.cmd**.

Remote Desktop is still a necessary Windows Server 2012 remote management technology because some environments have not upgraded their administrator's workstations from Windows® XP and other environments may have Windows Server 2012 deployed even when the users in those environments primarily use third-party operating systems. You can configure Remote Desktop on a computer running the full version of Windows Server 2012 by performing the following steps:

1. In the Server Manager console, click the **Local Server** node.
2. Click **Disabled** next to Remote Desktop.
3. On the **Remote** tab of the **System Properties** dialog box, select between one of the following options:
  - **Don't allow connections to this computer.** The default state of remote desktop is disabled.
  - **Allow connections from computers running any version of Remote Desktop.** Enables connections from Remote Desktop clients that do not support Network Level Authentication
  - **Allow Connections only from Computers running Remote Desktop with Network Level Authentication.** Enables secure connections from computers running Remote Desktop clients that support network level authentication.

You can enable and disable Remote Desktop on computers running the Server Core installation option by using the **sconfig.cmd** menu-driven command prompt utility.

## How Remote Management Works In Windows Server 2012

Windows Remote Management (WinRM) is a collection of technologies that enables administrators to manage server hardware when logged on directly or over the network. Windows Server 2012 uses WinRM to enable management of multiple computers concurrently through a single Server Manager console. Windows Remote Management includes the following components:

- **WS-Management protocol.** A SOAP-based firewall-aware protocol that enables computers to exchange management information. SOAP uses XML messages when transmitting information.
- **WinRM Scripting API.** This scripting API enables systems to obtain data from remote computers through WS-Management protocol operations.
- **Winrm.cmd.** Command-line systems management tool that enables you to configure WinRM. For example, you can use this tool to enable Windows Remote Management on a server.
- **Winrs.exe.** Tool that enables you to execute most **cmd.exe** commands on remote servers.

For example, to obtain the IP address information and list of running tasks on server LON-SVR1, issue the command:

```
Winrs -r:lön-svr1 ipconfig;tasklist
```



**Note:** You can learn more about Windows Remote Management at:  
[http://msdn.microsoft.com/en-us/library/windows/desktop/aa384291\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/windows/desktop/aa384291(v=vs.85).aspx).

Use winrm qc to enable Windows Remote Management. This will:

- Configure the WinRM service for automatic startup
- Start the WinRM service
- Configure a listener
- Configure firewall rules

Windows PowerShell Remoting uses the following cmdlets:

- Invoke-Command. Allows you to run commands remotely
- Enter-PSession. Start a remote PowerShell session.

You can enable Windows Remote Management by issuing the following command:

```
Winrm qc
```

Running this command does the following:

1. Configures the WinRM service to with the Automatic startup type.
2. Starts the WinRM service.
3. Configures a listener that will accept WinRM requests on any IP address.
4. Creates a firewall exception for WS-Management traffic using the HTTP protocol.

If you do not know whether a server is configured for Windows Remote Management, you can run the following command to obtain Windows Remote Management configuration information:

```
Winrm get winrm/config
```



**Additional Reading:** You can learn more about configuring Windows Remote Management by reading the following Performance Team post: <http://blogs.technet.com/b/askperf/archive/2010/09/24/an-introduction-to-winrm-basics.aspx>.

You can use Remote Windows PowerShell to run commands against a correctly configured remote server. There are several methods that you can use to accomplish this. You can use the **Invoke-Command** cmdlet to run a command or a script. For example, to view the list of installed roles and features on LON-SVR1 and LON-SVR2 when the ServerManager module is loaded and both are configured for Windows Remote Management, issue the command:

```
Invoke-Command -Computers LON-SVR1, LON-SVR2 -scriptblock {Get-WindowsFeature | Where-Object {$_.InstallState -eq "Installed"}}
```

You can also start a remote Windows PowerShell session by using the **Enter-PSSession** cmdlet. To end the session, run the **Exit-PSSession** cmdlet. For example, to start a remote Windows PowerShell session to LON-SVR1, issue the command:

```
Enter-PSSession -computername LON-SVR1
```



**Additional Reading:** You can learn more about Remote Windows PowerShell at: [http://msdn.microsoft.com/en-us/library/windows/desktop/ee706585\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/windows/desktop/ee706585(v=vs.85).aspx).

## Demonstration: Configuring Servers for Remote Management

In this demonstration you will disable and enable Remote Management from Server Manager.

### Demonstration Steps

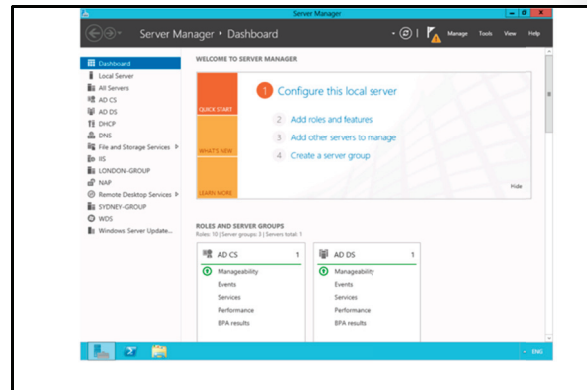
1. Use Server Manager on LON-DC1 to disable Remote Management.
2. Use the **winrm qc** command from a Windows PowerShell prompt to re-enable remote management on LON-DC1.
3. Use Server Manager to verify that Remote Management is re-enabled.

## Managing Server Groups in Server Manager

Server Manager in Windows Server 2012 automatically groups servers by role. This enables you to perform role-based tasks across all servers that host that role in the organization. For example, rather than connecting to each DNS server in the domain to perform a particular task, you can select the DNS node, select all servers that host DNS that you want to perform the task on, and then perform the task against that selection of servers.

A benefit to administrators is that servers in your organization are automatically grouped by role. For example, all servers that host the IIS or NAP roles are automatically grouped under the category nodes for those roles in the Server Manager console.

You can also use the Server Manager console to create custom server groups. A custom server group is a user-defined group of servers rather than a group of servers that share a specific role.



## Demonstration: Managing Remote Servers by Using Server Manager

In this demonstration you will see how to create a server group. You will then perform a remote management task on both servers that are members of the group using a single action.

### Demonstration Steps

1. On LON-DC1, use Server Manager to create a server group named LONDON-GROUP that has LON-DC1 and LON-SVR4 as members.
2. Use the group node as a method of starting the performance counters on both servers using the one action, rather than enabling performance counters on each server individually.
3. Use the Manageability column to verify that both LON-DC1 and LON-SVR5 are listed as **Online**.



# Lab: Installing and Configuring Servers Based on Windows Server 2012

## Scenario

A. Datum is an engineering and manufacturing company. The organization is based in London, England. The organization is quickly expanding the London location as well as internationally. Because the company has expanded, some business requirements are changing as well. To address some business requirements, A. Datum has decided to deploy Windows Server 2012 on an existing network populated with servers running the Windows Server 2008 and Windows Server 2008 R2 operating systems.

As one of the experienced Windows Server 2008 administrators, you are responsible for implementing many of the new features on Windows Server 2012. To become familiar with the new operating system, you plan to install a new Windows Server 2012 server running the Server Core version and complete the initial configuration tasks. You also plan to configure and explore the remote management features that are available in Windows Server 2012.

## Objectives

- Install Windows Server 2012 server core.
- Configure a Windows Server 2012 server core.
- Configure remote management for Windows Server 2012 Servers.

## Lab Setup

Estimated time: **60 minutes**

Virtual Machines	20417A-LON-DC1 20417A-LON-SVR5
User Name	Adatum\Administrator
Password	Pa\$\$w0rd

For this lab, you will use the available virtual machine environment. Before you begin the lab, you must complete the following steps:

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, click **20417A-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on using the following credentials:
  - User name: **Adatum\Administrator**
  - Password: **Pa\$\$w0rd**

## Exercise 1: Install Windows Server 2012 Server Core

### Scenario

After having problems effectively deploying and configuring the Server Core version of Windows Server 2008, A. Datum is interested in using the Server Core installation of Windows Server 2012 when possible because of the reduced hardware footprint and minimized update requirements. To become familiar with the new operating system, you plan to install and configure a new Windows Server 2012 server running the Server Core version as a way to determine whether the product is more easily managed than the earlier version.

The main tasks in this exercise are:

1. Install Windows Server 2012.
2. Convert a Windows Server 2012 server core installation to a full installation.
3. Convert a Windows Server 2012 full installation to a server core installation.

### ► Task 1: Install Windows Server 2012

1. In the Hyper-V Manager console, open the settings for **20417A-LON-SVR5**.
2. Configure the DVD drive to use the Windows Server 2012 image file named **Win2012\_RC.ISO**. This file is located at **C:\Program Files\Microsoft Learning\20417\Drives**.
3. Start **20417A-LON-SVR5**. On the **Windows Server 2012** page of the Windows Setup Wizard, verify the following settings, click **Next**, and then click **Install Now**:
  - Language to install: **English (United States)**
  - Time and currency format: **English (United States)**
  - Keyboard or input method: **US**
4. Select to install the **Windows Server 2012 Release Candidate Datacenter (Server Core Installation)** operating system.
5. Accept the license terms and then select **Custom: Install Windows Only (Advanced)**.
6. Install Windows Server 2012 on Drive 0.
  - Depending on the speed of the host computer, the installation will take approximately 20 minutes.
  - The virtual machine will restart several times during this process.
7. On the log on page, click **OK** and then enter **Pa\$\$w0rd** in both the **Password** and **Confirm password** boxes.
8. Click **OK** to complete the installation and log on.

### ► Task 2: Convert a Windows Server 2012 Server Core Installation to a Full Installation

1. On LON-SVR5 at the command prompt type:

```
mkdir c:\mount
```

2. Issue the following command and press Enter:

```
dism.exe /mount-image /ImageFile:d:\sources\install.wim /Index:4 /Mountdir:c:\mount /readonly
```

3. Start Windows PowerShell by typing the following command:

```
PowerShell.exe
```

4. From Windows PowerShell issue the following commands, pressing Enter after each:

```
Import-Module ServerManager  
  
Install-WindowsFeature -IncludeAllSubfeature User-Interfaces-Infra -  
Source:c:\mount\windows
```

5. When prompted, restart the server and then log on as **Administrator** with the password of **Pa\$\$w0rd** to verify the presence of the full GUI components.

### ► Task 3: Convert a Windows Server 2012 Full Installation to a Server Core Installation

1. Log on to LON-SVR5 and attempt to start Internet Explorer.
2. Start Windows PowerShell and issue the following commands:

```
Import-Module ServerManager  
  
Uninstall-WindowsFeature User-Interfaces-Infra  
  
Shutdown /r /t 5
```

3. Log on to LON-SVR5 as **Administrator** with the password of **Pa\$\$w0rd** and verify that it now configured to use the Server Core configuration.

## Exercise 2: Configure a Computer Running a Server Core Installation of Windows Server 2012

### Scenario

After you install Server Core, you want to configure some basic network and firewall settings and join computer to domain. During this initial deployment, you plan to perform these steps manually from the command-line.

The main tasks for this exercise are as follows:

1. Configure the network.
2. Add the server to the domain.
3. Configure Windows Firewall.

### ► Task 1: Configure the network

1. On LON-SVR5 in the command prompt, type **sconfig**.
2. Set the computer name **LON-SVR5**.
3. Restart the server as prompted and log on to LON-SVR5 as **Administrator** with the password of **Pa\$\$w0rd**.
4. Use the **hostname** command to verify the name change.
5. Start **sconfig** and configure **Network Settings**.
6. Select the index number of the network adapter that you want to configure.

7. Set the **Network Adapter Address** to the following:
  - o IP address: **172.16.0.111**.
  - o Subnet Mask: **255.255.0.0**.
  - o Default gateway **172.16.0.1**.
8. Set the preferred DNS server to **172.16.0.10**. Do not configure an alternative DNS server address.
9. Exit `sconfig` and verify network connectivity to `lon-dc1.adatum.com` using the ping utility.

### ► Task 2: Add the server to the domain

1. Use `sconfig` to switch to configure **Domain/Workgroup**.
2. Join the domain **adatum.com** using account **adatum\administrator** and the password of **Pa\$\$w0rd**.
3. Restart the server.
4. Log on to LON-SVR5 with the **adatum\administrator** account and a password of **Pa\$\$w0rd**.

### ► Task 3: Configure Windows Firewall

1. Use `sconfig.cmd` to Enable Remote Management.
2. At the command prompt, type **PowerShell.exe**.
3. Issue the following command to view the enabled Firewall rules that allow traffic:

```
Get-NetFirewallRule | Where-Object {$_.Action -eq "Allow"} | Format-Table -Property DisplayName
```

4. Issue the following command to view all disabled Firewall rules:

```
Get-NetFirewallRule | Where-Object {$_.Enabled -eq "False"} | Format-Table -Property Displayname
```

5. Issue the following command to view all Windows PowerShell cmdlets related to NetFirewallRule:

```
Get-Command -Noun NetFirewallRule
```

6. View the status of the Remote Desktop inbound firewall rule by issuing the following command:

```
Get-NetFirewallRule RemoteDesktop-UserMode-In-TCP
```

7. Issue the following command to enable the Remote Desktop Inbound Firewall rule:

```
Enable-NetFirewallRule RemoteDesktop-UserMode-In-TCP
```

8. Issue the following command to verify that the Remote Desktop Inbound Firewall rule is enabled:

```
Get-NetFirewallRule RemoteDesktop-UserMode-In-TCP
```

9. Issue the following command to disable the Remote Desktop Inbound Firewall Rule:

```
Disable-NetFirewallRule RemoteDesktop-UserMode-In-TCP
```

10. Verify that the Remote Desktop Inbound Firewall Rule is disabled.

```
Get-NetFirewallRule RemoteDesktop-UserMode-In-TCP
```

## Exercise 3: Configure Remote Management for Servers Running Windows Server 2012

### Scenario

IT management at A. Datum expects that many servers running Windows Server 2012 will be deployed in remote offices or as part of an online services deployment. To ensure that these servers can all be managed from a central location, you must configure the server for remote management. You must also verify the remote management functionality, and use Server Manager to manage multiple servers.

The main tasks for this exercise are as follows:

1. Validate the WinRM configuration.
2. Configure Server Manager for multiple server management.
3. Deploy a feature to the Server Core server.
4. To prepare for next module.

#### ► Task 1: Validate the WinRM configuration

1. On LON-DC1 use Server Manager to disable Remote Management.
2. Close the Server Manager console.
3. Open Windows PowerShell and issue the command **winrm qc**. When you are prompted, type **Y** and press Enter.
4. Open the Server Manager console and verify that Remote Management is now enabled.

#### ► Task 2: Configure Server Manager for multiple server management

1. On LON-DC1 in Server Manager, create a server group named **LONDON-GROUP** that has **LON-DC1** and **LON-SVR5** as members.
2. In the details pane, select both servers.
3. Scroll down to the **Performance** section, select both listed servers, right-click **LON-DC1**, and then click **Start Performance Counters**.
4. Scroll up and verify that in the **Manageability** column, both **LON-DC1** and **LON-SVR5** are listed as **Online**.

#### ► Task 3: Deploy a feature to the Server Core server

1. In the Server Manager console on LON-DC1, click **LONDON-GROUP**.
2. Add the **Windows Server Backup** feature to LON-SVR5.
3. In Server Manager, click the Flag and verify that the remote installation of Windows Server Backup has occurred.

#### ► Task 4: To prepare for next module

- When you are finished with the lab, revert the virtual machines to their initial state.

## Module Review and Takeaways

### Best Practices

- Unless you must have a full installation to support roles and features, deploy Server Core.
- Use Windows Remote Management to manage multiple servers from a single server using the Server Manager console.
- Use Windows PowerShell remoting to run remote Windows PowerShell sessions rather than logging on locally to perform the same task.

### Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Remote management connections fail	
Windows PowerShell commands not available	
Cannot install GUI features on Server Core Deployment	
Unable to restart a computer running Server Core	
Unable to join the domain	

### Review Question

Why is the Server Core installation the default installation option for Windows Server 2012 installations?

### Real-world Issues and Scenarios

Unless a particular role requires it, consider using the Server Core installation option as your default server deployment option. You can always install the GUI later if required.

Understand what roles and features you must deploy on a server prior to deploying that server, rather than deploying roles and features to servers without planning.

You should plan to manage many servers from one console, rather than logging on to each server individually.

# Module 2

## Monitoring and Maintaining Windows Server 2012

### Contents:

Module Overview	2-1
Lesson 1: Monitoring Windows Server 2012	2-2
Lesson 2: Implementing Windows Server Backup	2-11
Lesson 3: Implementing Server and Data Recovery	2-15
Lab: Monitoring and Maintaining Windows 2012 Servers	2-19
Module Review and Takeaways	2-26

## Module Overview

After you deploy Windows Server® 2012, you must ensure that it continues to run optimally by maintaining a healthy and stable environment. As in earlier versions of Windows Server, to maintain a healthy and stable environment, you must monitor Windows Server 2012 performance and make adjustments as required. Additionally, you must identify your important data and create backup copies. Finally, you must know how to restore your important data and servers by using the backup copies that you have created.

### Objectives

After completing this module, you will be able to:

- Monitor Windows Server 2012.
- Implement Windows Server Backup.
- Restore data and servers by using Windows Server Backup.

## Lesson 1

# Monitoring Windows Server 2012

When a system failure or an event that affects system performance occurs, you must be able to repair the problem or resolve the issue quickly and efficiently. With so many variables and possibilities in the modern network environment, the ability to determine the cause quickly frequently depends on having an effective performance monitoring methodology and tool set.

You can use performance-monitoring tools to identify components that require additional tuning and troubleshooting. By identifying components that require additional tuning, you can improve the efficiency of your servers. In addition to monitoring system performance, Windows Server 2012 provides tools for resource management. In this lesson, you will learn about tools in Windows Server 2012 that you can use for performance and resource monitoring and management.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe the reasons for monitoring servers.
- Describe the typical performance bottlenecks.
- Describe the tools for monitoring in Windows Server 2012.
- Create data collector sets.
- Describe the most common performance counters.
- Describe the use of alerts.
- Describe the use of event subscriptions.
- Configure event subscriptions.
- Describe how to monitor a network.

### Reasons for Monitoring Servers

Monitoring servers provides several benefits, and you might monitor a Windows-based server for several reasons. Some reasons include:

- To monitor the health of the IT infrastructure.
- To monitor service-level agreements (SLAs).
- To plan for future requirements.
- To identify issues.

#### IT Infrastructure Health

The effective operation of the server infrastructure is frequently critical to your organization's business goals.

The key factors in maintaining the consistency of server operation include correctly functioning and configured hardware, and sufficient use and assignment of resources.

Using performance-monitoring tools, you can record performance statistics that you can use to determine when a server is slower at responding to user requests, instead of relying on user perception of slow and

- Health of the IT infrastructure:
  - Normal activity
  - Abnormal activity
- SLA monitoring:
  - Monitor SLA areas to prevent issues
- Planning for future requirements:
  - Capacity
  - Reallocation
- Identifying issues:
  - Reactive
  - Proactive





fast response times. You can use these statistics to determine which component or components of the server infrastructure may be the source of performance-related issues.

### SLA Monitoring

Many organizations maintain SLAs that dictate the required availability for servers and server-hosted applications. These SLAs may contain stipulations about server availability (for example, the LON-DC1 server must be available 99.995 percent of business hours), or they may specify performance-related requirements (for example, the average query time for this database server must be less than five seconds for any given day).

Frequently, violation of an SLA results in reduction of payment for services or similar penalties. Therefore, you want to ensure that the SLAs imposed upon your environment are met on a continuing basis.

You can use performance-monitoring tools to monitor the specific areas related to your SLAs and help you identify issues that could affect your SLA before they become a problem.

### Planning for Future Requirements

The business and technical needs of your organization are subject to change. New initiatives may require new servers to host new applications or increased storage within your environment. Monitoring these areas over time enables you to assess effectively how the server resources are being used currently. Then, you can make an informed decision on how the server environment has to grow or change to meet future requirements.

### Identifying Issues

Troubleshooting problems that arise in the server environment can be tedious. Issues that affect users have to be resolved as quickly as possible and with minimal effect on the business needs of your organization.

Troubleshooting an issue only on the symptoms provided by users or anecdotal evidence frequently leads to misdiagnosis and wasted time and resources. Monitoring the server environment lets you take a more informed and proactive approach to troubleshooting. When you have an effective monitoring solution implemented, you can identify issues within your infrastructure before they cause a problem for the end-users. You can also have more concrete evidence of reported issues and narrow the cause of problems, saving you investigative time.

**Question:** List four troubleshooting procedures that would benefit from server monitoring.

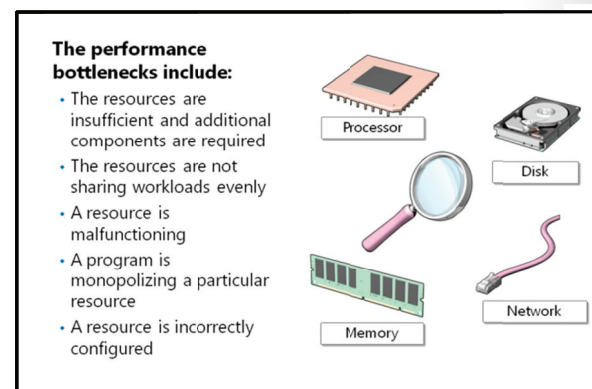
### Typical Performance Bottlenecks

Analysis of your monitoring data can reveal problems such as excessive demand on certain hardware resources that result in bottlenecks.

#### Causes of Bottlenecks

Demand on certain hardware resources may become extreme enough to cause resource bottlenecks for the following reasons:

- The resources are insufficient, and additional or upgraded components are required.
- The resources are not sharing workloads evenly and have to be balanced.
- A resource is malfunctioning and has to be replaced.



- A program is monopolizing a particular resource. This might require substituting another program, having a developer rewrite the program, adding or upgrading resources, or running the program during periods of low demand.
- A resource is configured incorrectly and configuration settings have to be changed.
- A security issue, such as viruses or Denial of Service attacks can be the reason for a bottleneck.

By monitoring the basic hardware components of your servers, you can determine the most likely bottleneck that is affecting the performance of your servers. By adding additional capacity to components, you can tune the servers to overcome initial limitations. The following table lists suggestions for improving performance on various types of hardware.

Hardware	Suggestion
Processors	<ul style="list-style-type: none"> <li>• You may be able to overcome performance bottlenecks that occur with processors by:</li> <li>• Adding processors.</li> <li>• Increasing the speed of processors.</li> <li>• Reducing or controlling process or affinity, or the number of processor cores an application uses. Limiting an application to only some processor cores frees the remaining cores for other applications to use.</li> </ul>
Disks	<ul style="list-style-type: none"> <li>• You may be able to increase disk performance by:</li> <li>• Adding faster disks.</li> <li>• Performing routine maintenance tasks such as defragmenting.</li> <li>• Moving data, applications, and the page files onto separate disks.</li> </ul>
Memory	<p>You can improve memory bottlenecks by adding additional physical memory. If the memory requested exceeds the physical memory, information will be written to virtual memory, which is slower than physical memory.</p> <p>However, increasing a computer's virtual memory could enable applications that consume a large amount of memory to run on a computer that has limited physical memory.</p> <p>Or, you can reduce the load on the server by reducing the number of users on the server or through application tuning.</p>
Networks	<p>You can reduce network bottlenecks by:</p> <ul style="list-style-type: none"> <li>• Upgrading network infrastructure, including network adapters to support increased network bandwidth.</li> <li>• Installing multiple network adapters in a server to distribute network load.</li> <li>• Reducing the traffic.</li> </ul> <p>You should consider the limitations of network bandwidth and segment networks, where appropriate. You can increase network throughput by tuning the network adapter and other network devices such as switches, firewalls, and routers.</p>

## Tools for Monitoring in Windows Server 2012

Several tools are available to help you in monitoring the server environment, both historical and real time. The following is a list of tools to help you in monitoring the server environment.

### Tools:

- Event Viewer: Collects information that relates to server operations
- Task Manager: Provides information related to hardware performance and applications that are currently running on the server
- Resource Monitor: Provides real-time performance-related information of the server
- Performance Monitor: Provides both real-time and historical monitoring of the server's performance
- Reliability Monitor: Provides a historical view of the server's reliability-related information



Tool	Description
Event Viewer	Event Viewer collects information that relates to server operations. This information can help identify performance issues on a server. You should search for specific events in the event log file to locate and identify problems.
Task Manager	Task Manager helps you monitor the real-time aspects of the server. You can view information related to hardware performance and the applications and processes that are currently running on the server.
Resource Monitor	Resource Monitor helps you to look deeper into the real-time performance of the server. It provides performance information related to the CPU, memory, hard disk, and network components of the server.
Performance Monitor	Performance Monitor is the most robust monitoring tool in Windows Server 2012. It enables both real-time and historical monitoring of the server's performance and configuration data.
Reliability Monitor	Reliability Monitor provides a historical view of the server's reliability-related information such as event log errors and warnings.

## Demonstration: Creating Data Collector Sets

The data collector set is a custom set of performance counters, event traces, and system configuration data.

A data collector set organizes multiple data-collection points into a single, portable component. You can use a data collector set on its own, group it with other data collector sets, and incorporate it into logs, or view it in the Performance Monitor. You can configure a data collector set to generate alerts when it reaches thresholds.

You can also configure a data collector set to run at a scheduled time, for a specific length of time, or until it reaches a predefined size. For example, you can run the data collector set for ten minutes every hour during your working hours to create a performance baseline. You can also set the data collector to restart when set limits are reached so that a separate file is created for each interval.

After you have created a combination of data collectors that describe useful system information, you can save them as a data collector set, and then run the set and view the results.

In this demonstration, you will create a data collector set.

## Demonstration Steps

### Create a new data collector set named Windows Server Monitoring

1. On LON-SVR1, open the Performance Monitor, and create a data collector set named **Windows Server Monitoring**.
2. Configure the data collector set to include the Performance counter data logs for Processor/% Processor Time, Memory/ Available Mbytes, and Logical Disk/% Free Disk Space.

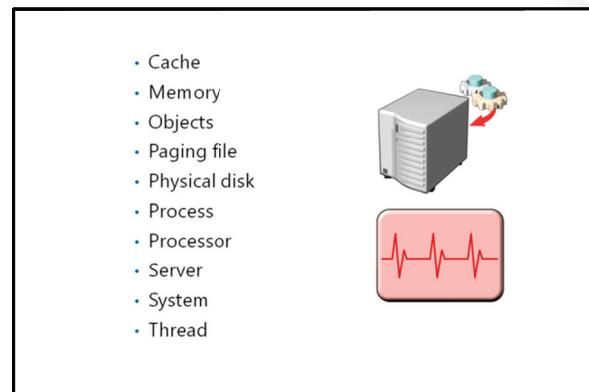
### Verify that the data collector set works correctly

1. Start the Windows Server Monitoring data collector set, and let it run for one minute.
2. Stop the Windows Server Monitoring data collector set, and then review the latest report.

## Most Common Performance Counters

Specific server roles install a range of performance objects and associated counters. The common performance counters include:

- *Cache counters.* These counters monitor the file system cache. The cache is an area of physical memory that is used to store recently-used data to enable access to the data without having to read from the disk.
- *Memory counters.* These counters monitor physical, random access memory (RAM), virtual memory, and disks, including paging, which is the movement of pages of code and data between disk and physical memory.
- *Counters for objects.* These counters monitor logical objects in the system, including threads and processes.
- *Paging file counters.* Paging file is the reserved space on the disk that complements committed physical memory.
- *Physical disk counters.* These counters monitor the physical disks such as hard drivers or fixed drives. The drives that appear in the Disk Management console are monitored by these counters. Hardware redundant array of independent disks (RAID) may not be visible to these counters.
- *Process counters.* These counters monitor running applications and system processes. All the threads in a process share the same address space and have access to the same data.
- *Processor counters.* These counters measure aspects of processor activity. Each processor is represented as an instance of the object.
- *Server counters.* These counters measure communication between the local computer and network.
- *System counters.* These counters apply to more than one instance of component processes on the computer.
- *Thread counters.* These counters measure aspects of thread behavior. A thread is the basic object that runs instructions on a processor. All running processes have at least one thread.



Windows Server 2012 uses server roles to improve server efficiency and security. Only the performance objects and counters that are relevant to the installed server role are available to monitor.

You can enable missing performance objects and counters by installing additional server roles or adding features. Additional performance objects that are installed with each server role can help with server monitoring. The following table identifies common server roles and the performance objects that can be monitored to assess performance.

Server role	Performance counters to monitor
Active Directory® Domain Services (AD DS)	<p>If you notice slow write or read operations, under the Physical Disk category, check the following disk I/O counters to see whether many queued disk operations exist:</p> <ul style="list-style-type: none"> <li>• Avg. Disk Queue Length</li> <li>• Avg. Disk Read Queue Length</li> <li>• Avg. Disk Write Queue Length</li> </ul> <p>If Local Security Authority Subsystem or lsass.exe uses lots of physical memory, under the Database category, check the following Database counters to see how much memory is used to cache the database for Active Directory Domain Services:</p> <ul style="list-style-type: none"> <li>• Database Cache % Hit</li> <li>• Database Cache Size (MB)</li> </ul>
File Server	<p>File Servers are typically heavily dependent on their physical disk systems for file read and write operations. You should measure the following counters to ensure that the PhysicalDisk subsystem is keeping up with server demand:</p> <ul style="list-style-type: none"> <li>• % Disk Time</li> <li>• Avg. Disk Queue Length</li> <li>• Avg. Disk Bytes/Transfer</li> </ul> <p>Network performance is also a primary component of file server performance. You should monitor the following counters to ensure that required network bandwidth is available to the file server:</p> <ul style="list-style-type: none"> <li>• Bytes Received Per Second</li> <li>• Bytes Sent Per Second</li> <li>• Output Queue Length</li> </ul>
Hyper-V® (virtualization)	<p>Performance troubleshooting and tuning can be difficult on virtualized servers. Virtual hardware provides a less consistent monitoring environment than physical hardware.</p> <p>Two layers of performance monitoring are usually recommended in a virtualized scenario. One at the physical or host server level to monitor key physical hardware components, and one at the virtualized server level to monitor the virtual hardware and its effect on the operating system and applications of the virtual server.</p>
Web Server (IIS)	<p>Network-related performance counters are an important tool in measuring web server performance.</p> <p>Additionally, processor related counters can be helpful in identifying issues in which web server applications are running processor intensive processes.</p> <p>The Web Service performance counters provide valuable information about requests to the web server, bandwidth consumed, and web server-specific statistics like page not found errors.</p>

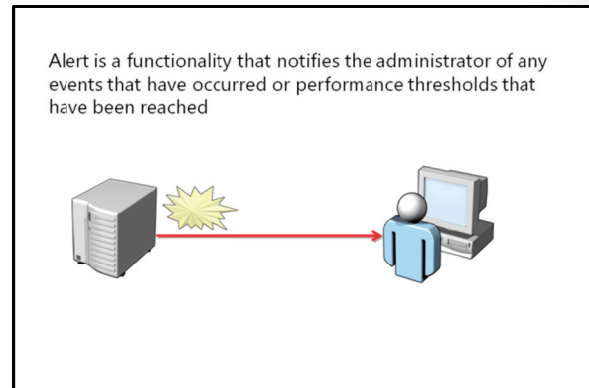
## What Are Alerts?

Alert is a functionality in Windows Server 2012 that notifies you when certain events have occurred or when certain performance thresholds are reached. You can configure alerts in Windows Server 2012 as network messages or as events that are logged in the application event log. You can also configure alerts to start applications and performance logs.

You can configure alerts when you create data collectors, by selecting the **Performance** Counter Alert type of the data collector.

When you create the alert, configure the following settings:

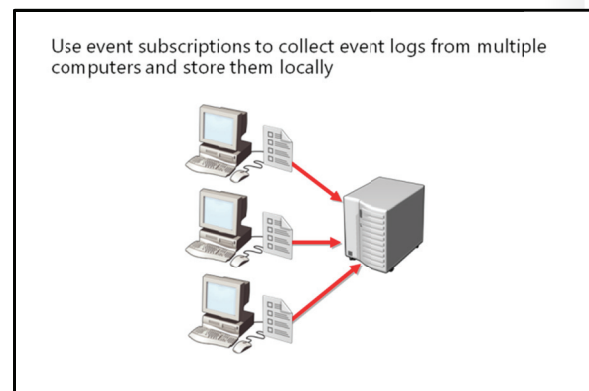
- *Alert when.* This is the alert threshold setting for a specific performance counter.
- *Alert Action.* This setting specifies whether to log an entry in the application event log, or start another data collector set.
- *Alert Task.* This setting specifies which command task should be triggered and when alert threshold is reached. In addition, you may specify command parameters, if applicable.



## What Are Event Subscriptions?

Event log subscriptions is a feature when it is configured, enables a single server to collect copies of events from multiple systems. Using WinRM and the Windows Event Collector service, you can collect events in the event logs of a centralized server, where you can analyze them together with the event logs of other computers that are being collected on the same central server.

Subscriptions can be either collector-initiated or source computer-initiated:



- *Collector-initiated.* A collector-initiated subscription, or a pull subscription identifies all the computers that the collector will receive events from, and will typically pull events from these computers. In a collector-initiated subscription, the subscription definition is stored and maintained on the collector computer. You use pull subscriptions when much of the computers have to be configured to forward the same types of events to a central location. In this manner, only one subscription definition has to be defined and specified to apply to all computers in the group.
- *Source computer-initiated.* In a source computer-initiated subscription, or push subscription, source computers push events to the collector. In a source computer-initiated subscription, the subscription definition is created and managed on the source computer, which is the computer that is sending events to a central source. You can define these subscriptions manually, or by using Group Policy. You create push subscriptions when each server is forwarding a different set of event than other servers, or when control over the event forwarding process has to be maintained at the source computer; possibly when frequent changes have to be made to the subscription.

## Event Subscription Requirements

To implement event subscriptions in your environment, several prerequisites must be met:

- You must enable and configure WinRM on both the source and the collector computers by using the following command.

```
winrm qc
```

- You must start and configure the Windows Event Collector (Wecutil) service to receive events on the collector computer. You can achieve this by running the following command.

```
Wecutil qc
```

Events that are collected by a subscription can be collected into any of the collector computer's default event logs, or they can be collected into an event log specifically created to host collected events.

## Demonstration: Configuring Event Subscriptions

Event subscription is a cost-effective and customizable tool to get a consolidated view of monitored activities and events in target servers, and timely issue alerts. In Windows Server 2012, subscribing and forwarding events with triggers to send out alerts is a straight-forward process.

### Demonstration Steps

#### Configure the source computer

1. Switch to LON-SVR1.
2. At the command prompt, run the **winrm quickconfig** command to enable the administrative changes that are required on a source computer.
3. Add the LON-DC1 computer to the local Administrators group.

#### Configure the collector computer

1. Switch to LON-DC1.
2. At the command prompt, run the **wecutil qc** command to enable the administrative changes that are required on a collector computer.

#### Create a subscribed log

1. Open Event Viewer.
2. Create a new subscription with the following properties:
  - Computers: **LON-SVR1**
  - Name: **LON-SVR1 Events**
  - Type of subscription: **Collector Initiated**
  - Events: **Critical, Warning, Information, Verbose, and Error**
  - Logged: **last 7 days**
  - Logs: **Windows Logs**

#### Check the subscribed log

1. Switch to LON-DC1.
2. In Performance Monitor, check for events in the subscribed Application log.




## Monitoring a Network

Because network infrastructure services are an important foundation of many other server-based services, you must make sure that they are configured correctly and are running optimally.

Collecting performance-related data on the network infrastructure services benefits your organization in:

- Helping to optimize network infrastructure server performance. By providing performance baseline and trend data, you can help your organization optimize network infrastructure server performance.
- Troubleshooting servers. Where server performance has decreased, either over time or during periods of peak activity, you can help identify possible causes and take corrective action to ensure that you can bring the service back within the limits of your SLA.

You can use Performance Monitor to collect and analyze the relevant data.



### Monitoring Domain Name System DNS

Domain Name System (DNS) provides name resolution services on the network. You can monitor the DNS Server role of Windows Server 2012 to determine the following aspects of your DNS infrastructure:

- General DNS server statistics, including the number of overall queries and responses that are processed by the DNS server
- User Datagram Protocol (UDP) or Transmission Control Protocol (TCP) counters, for measuring DNS queries and responses that are processed respectively by using either of these transport protocols
- Dynamic update and secure dynamic update counters, for measuring registration and update activity that is generated by dynamic clients
- Memory usage counters, for measuring system memory usage and memory allocation patterns that are created by operating the server as a DNS server
- Recursive lookup counters, for measuring queries and responses when the DNS service uses recursion to look up and fully resolve DNS names on behalf of requesting clients
- Zone transfer counters, including specific counters for measuring the following: all zone transfer (AXFR), incremental zone transfer (IXFR), and DNS zone update notification activity

### Monitoring DHCP

The Dynamic Host Configuration Protocol (DHCP) service provides dynamic IP configuration services on the network. You can monitor the Windows Server 2012 DHCP Server role to determine the following aspects of your DHCP server:

- The Average Queue Length indicates the current length of the internal message queue of the DHCP server. This number represents the number of unprocessed messages that are received by the server. A large number might indicate heavy server traffic.
- The Milliseconds per packet (Avg.) counter is the average time in milliseconds that is used by the DHCP server to process each packet it receives. This number varies, depending on the server hardware and its I/O subsystem. A spike could indicate a problem, either with the I/O subsystem becoming slower or because of a processing overhead on the server.



## Lesson 2

# Implementing Windows Server Backup

In order to protect critical data, every organization must perform a backup regularly. Having a well-defined and tested backup strategy ensures that companies can restore data if there is any unexpected failures or data loss. This lesson describes the Windows Server Backup feature in Windows Server 2012 and the Microsoft Online Backup Service for Windows Server 2012.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe the features of Windows Server Backup.
- Describe the Microsoft Online Backup Service.
- Describe the methods for backing up server roles running Windows Server 2012.
- Back up Windows Server 2012 by using Windows Server Backup.

### Features of Windows Server Backup in Windows 2012

The Windows Server Backup feature in Windows Server 2012 consists of a Microsoft Management Console (MMC) snap-in and command-line tools. You can use wizards in the Windows Server Backup feature to guide you through running backups and recoveries. You can use Windows Server Backup 2012 to back up:

- Full server (all volumes)
- Selected volumes
- Select specific items for backup

You can use Windows Server Backup to:

- Back up full server (all volumes)
- Back up selected volumes
- Back up select specific items for backup
- Perform a bare-metal recovery
- Perform a system state
- Back up individual files and folders
- Exclude selected files or file types during backup
- Select from more storage locations for the backup
- Use the Microsoft Online Backup Service

In addition, Windows Server Backup 2012 lets you:

- Perform a bare-metal restore. Bare-metal restore includes all volumes that are required for Windows to run. You can use this backup type together with the Windows Recovery Environment to recover from a hard disk failure, or if you have to recover the whole computer image to new hardware.
- Use system state. System state is the ability to use the GUI interface to create a system state backup.
- Recover individual files and folders. The Individual files and folders option enables you to back up selected files and folders, instead of just full volumes.
- Exclude selected files or file types. For example, you can exclude .tmp files.
- Select from more storage locations. You can store backups on remote shares or non-dedicated volumes.
- Use the Microsoft Online Backup Service. The Microsoft Online Backup Service is a cloud-based backup solution for Windows Server 2012 which enables files and folders to be backed up and recovered from the cloud to provide off-site backup.

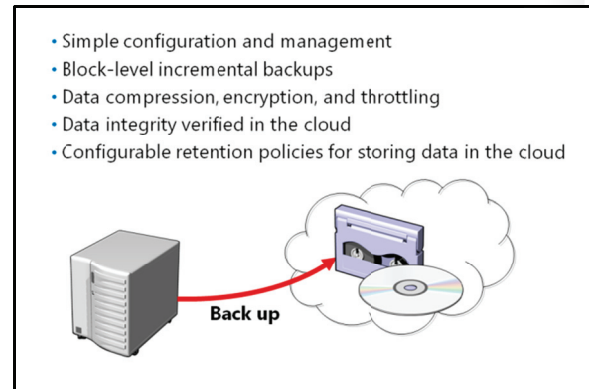
If there are disasters such as hard disk failures, you can perform system recovery by using a full server backup and the Windows Recovery Environment—this will restore your complete system onto the new hard disk.

The ability to take just a system state backup is not exposed in the GUI interface of backup. If you want to take just a system state backup, you must use the `wbadmin.exe` utility. `WBadmin.exe` is a command prompt utility.

## What Is Microsoft Online Backup Service?

The Microsoft Online Backup Service is a cloud-based backup solution for Windows Server 2012 managed by Microsoft. You can use this service to back up files and folders and recover them from the cloud to provide off-site protection against data loss caused by disasters. You can use this service to back up and protect critical data from any location.

This service is built on the Windows Azure® platform and uses Windows Azure blob storage for storing customer data. Windows Server 2012 uses the downloadable Microsoft Online Backup Agent to transfer file and folder data securely to the Microsoft Online Backup Service. After you install the Microsoft Online Backup Agent, the Microsoft Online Backup Service Agent integrates its functionality through the familiar Windows Server Backup interface.



### Key Features

The key features that Windows Server 2012 provides through the Microsoft Online Backup service include:

- *Simple configuration and management.* Integration with the familiar Windows Server Backup utility provides a seamless backup and recovery experience to a local disk, or to the cloud. Other features include:
  - Simple user interface to configure and monitor the backups
  - Integrated recovery experience to recover files and folders from local disk or from cloud
  - Easily recover any data that was backed up onto any server of your choice
  - Scripting capability that is provided by the Windows PowerShell command-line interface
- *Block-level incremental backups.* The Microsoft Online Backup Agent performs incremental backups by tracking file and block-level changes and only transferring the changed blocks, therefore, reducing the storage and bandwidth usage. Different point-in-time versions of the backups use storage efficiently by only storing the changed blocks between these versions.
- *Data compression, encryption and throttling.* The Microsoft Online Backup Agent ensures that data is compressed and encrypted on the server before it is sent to the Microsoft Online Backup Service on the network. Therefore, the Microsoft Online Backup Service only stores encrypted data in the cloud storage. The encryption passphrase is not available to the Microsoft Online Backup Service, and therefore, the data is never decrypted in the service. Also, users can set up throttling and configure how the Microsoft Online Backup service uses the network bandwidth when backing up or restoring information.
- *Data integrity verified in the cloud.* In addition to the secure backups, the backed up data is also automatically checked for integrity after the backup is finished. Therefore, any corruptions which may arise because of data transfer can be easily identified and they are fixed in next backup automatically.

- *Configurable retention policies for storing data in the cloud.* The Microsoft Online Backup Service accepts and implements retention policies to recycle backups that exceed the desired retention range, thereby meeting business policies and managing backup costs.



**Additional Reading:** Windows Azure Storage

<http://www.windowsazure.com/en-us/home/features/storage/>

## Methods to Back Up Server Roles

You can back up most services on computers running Windows Server 2012 by performing a system state backup. Some services also enable configuration and data backup from their respective management console.

The following table lists the methods that you can use to back up specific roles on computers running Windows Server 2012.

Component	Backup Strategy
DHCP	<ul style="list-style-type: none"> <li>• System state</li> <li>• DHCP console</li> </ul>
Certificate Services	<ul style="list-style-type: none"> <li>• System state</li> <li>• Certification Authority console</li> </ul>
IIS	<ul style="list-style-type: none"> <li>• System state</li> <li>• Website files and folders backup</li> <li>• Appcmd.exe</li> </ul>
Network Policy and Access Services (NPAS)	<ul style="list-style-type: none"> <li>• System state</li> </ul>
DNS	<ul style="list-style-type: none"> <li>• System state</li> <li>• Dnscmd.exe for zone export and import</li> </ul>
File and Print Services	<ul style="list-style-type: none"> <li>• Volumes</li> <li>• File and folder backup</li> <li>• System state (for shares and permissions)</li> </ul>

Role	Method
DHCP	<ul style="list-style-type: none"> <li>• System state backup backs up all scopes and options.</li> <li>• DHCP console backup backs up individual scopes or all scopes.</li> </ul>
Certificate	<ul style="list-style-type: none"> <li>• System state backup backs up whole configuration and certificate services database.</li> <li>• Certification Authority console backup backs up certificate services data and settings.</li> </ul>
Internet Information Services (IIS)	<ul style="list-style-type: none"> <li>• System state backup enables the back up of IIS data and settings.</li> <li>• Appcmd.exe lets you back up IIS components.</li> <li>• Website files and folders have to be backed up. When backing up IIS components, ensure that the website files and folders are also backed up. These are not backed up by a system state backup.</li> </ul>
Network Policy and Access Services (NPAS)	<ul style="list-style-type: none"> <li>• System state backup enables the back up of NPAS configuration.</li> </ul>
DNS	<ul style="list-style-type: none"> <li>• System state backup backs up all DNS configurations and zones stored on the server.</li> <li>• Dnscmd.exe lets you export and import zones.</li> </ul>
File and Print Services	<ul style="list-style-type: none"> <li>• System state backs up shared folder permissions and settings.</li> <li>• Volume backup enables a back up of all files and folders that are located on that volume.</li> <li>• File and folder backup backs up content of shared folders.</li> </ul>

## Demonstration: Backing Up Windows Server 2012 by Using Windows Server Backup

In this demonstration, you will see how to use the backup wizard to back up a folder.

### Demonstration Steps

1. On **LON-SVR1**, start **Windows Server Backup**.
2. Run the Backup Once Wizard to back up the **C:\HR Data** folder to the remote folder, **\\LON-DC1\Backup**.

## Lesson 3

# Implementing Server and Data Recovery

Every organization might experience losing some of its data, because of reasons, such as hardware failures, file system corruption, or when a user unintentionally deletes critical data. Therefore, organizations must have well-defined and tested recovery strategies that will help them to bring their servers and data back to a healthy and operational state, in the fastest time possible. This lesson describes how to restore data and servers by using Windows Server Backup feature in Windows Server 2012 and Microsoft Online Backup Service in Windows Server 2012.

### Lesson Objectives

- Describe the options for server recovery.
- Describe the option for server restore.
- Describe the considerations for data recovery.
- Perform a restore with Windows Server Backup.
- Describe how to perform a restore with online backup.

### Options for Server Recovery

Windows Server Backup in Windows Server 2012 provides the following recovery options:

- *Files and folders.* You can back up individual files or folders as long as the backup is on an external disk or in a remote shared folder.
- *Applications and data.* You can recover applications and data if the application has a Volume Shadow Copy Service writer and is registered with Windows Server Backup.
- *Volumes.* Restoring a volume always restores all the contents of the volume. You cannot restore individual files or folders.
- *Operating system.* You can recover the operating system through Windows Recovery Environment (WinRE).
- *Full server.* You can recover the full server through WinRE.
- *System state.* System state creates a point-in-time backup that you can use to restore a server to a previous working state.

The options for server recovery include:

- Files and folders
- Applications and data
- Volumes
- Operating system
- Full server
- System state

The Windows Server Backup Recovery Wizard provides several options for managing file and folder recovery. They are:

- *Recovery Destination.* Under Recovery Destination, you can select any one of the following options:
  - *Original location.* The original location restores the data to the location it was backed up originally.
  - *Another location.* Another location restores the data to a different location.

- **Conflict Resolution.** Restoring data from a backup frequently conflicts with existing versions of the data. Conflict resolution lets you determine how those conflicts will be handled. When these conflicts occur, you have the following options:
  - Create copies and have both versions
  - Overwrite existing version with recovered version
  - Do not recover items if they already exist in the recovery location
- **Security Settings.** You can use this option to restore permissions to the data being recovered.

## Options for Server Restore

You perform server restore by starting the computer from the Windows Server 2012 installation media, selecting the computer repair option, and then selecting the full server restore option.

When you perform full server restore, consider the following aspects:

- **Bare-metal restore.** Bare-metal restore is the process during which you restore an existing server in its entirety to new or replacement hardware. When you perform a bare-metal restore, the restore proceeds and the server restarts. Later, the server becomes operational. In some cases, you may have to reset the computer's Active Directory account because these can sometimes become desynchronized.
- **Same or larger disk drives.** The server hardware that you are restoring to must have disk drives that are the same size or larger than the drives of the original host server. If this is not the case, the restore will fail. It is possible, although not advised, to successfully restore to hosts that have slower processors and less RAM.
- **Importing to Hyper-V.** Because server backup data is written to the VHD format, which is also the format that is used for virtual machine hard disks, it is possible, with some care, to use full server backup data as the basis of creating a virtual machine. Doing this gives you the option of ensuring business continuity while sourcing the appropriate replacement hardware.

The server restore locations include:

- Original host: Bare-metal restore
- New host: Bare-metal restore
- Hyper-V Server: Virtual machine restore
- Alternate boot-to-VHD

## Considerations for Data Recovery

There are several strategies that you can pursue in developing a data recovery procedure. Data is the most frequently recovered component of an IT infrastructure.

Consider the following components in a data recovery strategy:

- Letting users recover their own data by using the earlier version's functionality (volume shadow copy)

The four options to recover data include:

- Allowing users to recover their data
- Recovering data to an alternate location
- Recovering data to the original location
- Performing a full volume recovery

- Performing a recovery to an alternative location
- Performing a recovery to the original location
- Performing a full volume recovery

### Earlier Versions of Files: Users Recover Their Own Data

The most common form of data recovery performed by IT departments is the recovery of files and folders that users have deleted, lost, or in some way made corrupted. The Previous Versions of Files functionality, which you can enable on all computers running Windows Server 2012 lets users recover their own files. After end-users are trained to do this, the IT department spends time recovering more important data.

From a planning perspective, you should consider increasing the frequency at which snapshots for previous versions of files are generated. This gives users more options when they try to recover files that have recently become deleted or corrupted.

### Recovering Data to an Alternative Location

A common recovery problem is the unintentional replacement of important data when recovering from backup. This can occur when recovery is performed to a location with live data, instead of to a separate location where the necessary data can be located and the unnecessary data discarded.

When you perform a recovery to an alternative location, always ensure that permissions are also restored. A common problem is administrators recovering data that includes restricted material to a location where important permissions are not applied, enabling unintended access to data for those that should not have it.

### Recovering Data to the Original Location

During some types of failures, such as data corruption or deletion, you have to restore data to the original location, because applications or users who access those data are preconfigured with the information on where the data is located.

### Recovering Volumes

If a disk fails, the quickest way to recover the data sometimes is to do a volume recovery, instead of a selective recovery of files and folders. When you do a volume recovery, you must check whether any shared folders are configured for the disks, and if the quotas and File Server Resource Manager management policies are still in effect.

## Demonstration: Restoring with Windows Server Backup

In this demonstration, you will see how to use the Recovery Wizard to restore a folder.

### Demonstration Steps

1. On LON-SVR1, delete the **C:\HR Data** folder.
2. In the Windows Server Backup MMC, run Recovery Wizard and specify the following information:
  - Getting Started: **A backup stored on another location**
  - Specify Location type: **Remote Shared Folder**
  - Specify Remote Folder: **\\LON-DC1\Backup**
  - Select Backup Date: Default value, **Today**
  - Select Recovery Type: Default value, **Files and Folders**

- Select Items to Recover: **LON-SVR1\Local Disk (C:)\HR Data**
  - Specify Recovery Options: **Another Location (C:)**
3. Locate **C:\** and ensure that the files are restored.

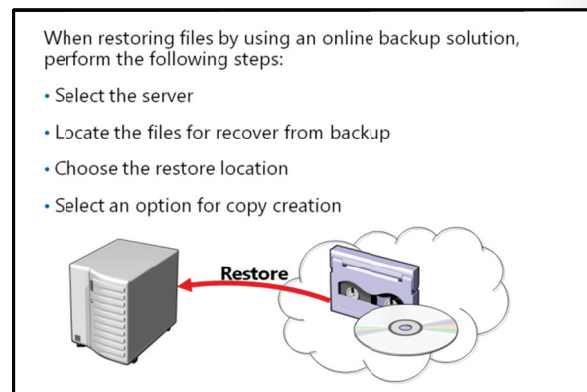
## Restoring with an Online Backup Solution

You can use Microsoft Online Backup Service only on Windows Server 2012 servers. You do not have to restore data on the same server that you backed up. You can restore data on some other server, instead.

You can recover files and folders by using both Microsoft Online Backup MMC in Server Manager, or Windows PowerShell® by performing the following steps:

1. Select the server where backup data was originally created, that is, whether it is a local server or another server. If you select Another server option, you must provide your Microsoft Online Backup Service Administrator credentials.
2. Browse for files that have to be restored can be browsed or search for them in the Microsoft Online Backup Service.
3. After you locate the files, select them for recovery, and select a location where the files will be restored.
4. When restoring files, select from the following options:
  - Create copies so that you have both the restored file and original file in the same location. The restored file has its name in the following format: Recovery Date+Copy of+Original File Name
  - Overwrite the existing versions with the recovered version
  - Do not recover the items that already exist on the recovery destination

After you complete the restore procedure, the files will be restored on Windows Server 2012 located in your site.





## Lab: Monitoring and Maintaining Windows 2012 Servers

### Scenario

To obtain accurate information about server usage, it is important to establish a performance baseline with a typical load for the new Windows Server 2012 servers. In addition, to make the process of monitoring and troubleshooting easier, IT management wants to implement centralized monitoring of event logs.

Much of the data that is stored on the A. Datum network is very valuable to the organization. Losing this data permanently would be a very significant loss to the organization. Also, several servers that run on the network provide very valuable services for the organization; losing these servers for a significant time would also result in losses to the organization. Because of the significance of the data and services, it is important that they can be restored even if there is any disaster.

One of the options that A. Datum is considering is backing up some critical data to a cloud-based service. A. Datum is considering this as an option for small branch offices that do not have a full data center infrastructure.

As one of the senior network administrators at A. Datum, you are responsible for planning and implementing a monitoring and system recovery solution that will meet the management and business requirements.

### Objectives

After completing this lab, you will be able to:

- Configure centralized monitoring for Windows 2012 servers.
- Back up Windows Server 2012 Servers.
- Restore files by using Windows Server Backup.
- Perform an online backup and restore for Windows Server 2012 servers.

### Lab Setup

Estimated time: **75 minutes**

Virtual Machine(s)	20417A-LON-DC1 20417A-LON-SVR1
User Name	Adatum\Administrator
Password	Pa\$\$w0rd

Virtual Machine(s)	MSL-TMG1
User Name	Administrator
Password	Pa\$\$w0rd

## Lab Setup

For this lab, you will use the available virtual machine environment. Before you begin the lab, you must complete the following steps:

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, click **20417A-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on using the following credentials:
  - a. User name: **Adatum\Administrator**
  - b. Password: **Pa\$\$w0rd**
5. Repeat steps 2-4 for **20417A-LON-SVR1**.
6. Repeat steps 2-3 for **MSL-TMG1**. Log on as **Administrator** with the password of **Pa\$\$w0rd**.

## Exercise 1: Configuring Centralized Monitoring for Windows Server 2012 servers

### Scenario

The management at A.Datum has asked for a monthly report on server performance. To provide a monthly report, you plan to establish centralized monitoring of the server. You decide to configure Server Manager to monitor all servers from a single console. You also decide to configure performance monitoring for some critical resources, and to collect events from several business-critical servers at a central location.

The main tasks for this exercise are as follows:

1. Configure Server Manager to monitor multiple servers.
2. Configure a data collector set.
3. Configure an event subscription.

### ► Task 1: Configure Server Manager to monitor multiple servers

1. Switch to LON-SVR1.
2. In the Server Manager console, in the navigation pane, click **All Servers**.
3. In the Server Manager console add LON-DC1 as another server to be monitored.
4. In the Actions pane, start the performance counters for both LON-SVR1 and LON-DC1.

### ► Task 2: Configure a data collector set

1. On LON-SVR1, open the Performance Monitor, and create a data collector set named **Windows Server Monitoring**.
2. Configure the data collector set to include the Performance counter data logs for Processor/% Processor Time, Memory/ Available MBytes and Logical Disk/% Free Disk Space.
3. Start the Windows Server Monitoring data collector set, and let it run for one minute.
4. Stop the Windows Server Monitoring data collector set, and then review the latest report.

### ► Task 3: Configure an event subscription

1. Switch to LON-SVR1.
2. At the command prompt, run the **winrm quickconfig** command to enable the administrative changes that are required on a source computer.
3. Add the LON-DC1 computer to the local Administrators group.
4. Switch to LON-DC1.
5. At the command prompt, run the **wecutil qc** command to enable the administrative changes that are required on a collector computer.
6. Open Event Viewer.
7. Create a new subscription with the following properties:
  - Computers: **LON-SVR1**
  - Name: **LON-SVR1 Events**
  - Type of subscription: **Collector Initiated**
  - Events: **Critical, Warning, Information, Verbose, and Error**
  - Logged: **last 7 days**
  - Logs: **Windows Logs**
8. Expand Event Viewer, expand Windows Logs, and then click Forwarded Events. Verify that events are forwarded from LON-SVR1.

**Results:** After completing this exercise, you will have configured Server Manager to monitor multiple servers, configured a data collector set, and configured an event subscription.

## Exercise 2: Backing up Windows Server 2012

### Scenario

The LON-SVR1 server contains financial data that must be backed up regularly. This data is important to the organization. You decide to use Windows Server Backup to back up critical data. You plan to install this feature and configure a scheduled backup.

The main tasks for this exercise are as follows:

1. Install the Windows Server Backup feature.
2. Configure a scheduled backup.
3. Complete an on-demand backup.

### ► Task 1: Install the Windows Server Backup feature

1. Switch to LON-SVR1.
2. Open Server Manager and install the Windows Server Backup role.
3. Install the role on **LON-SVR1** and then accept the default values on the Add Role wizard.

### ► Task 2: Configure a scheduled backup

1. On LON-SVR1, start Windows Server Backup.
2. Configure Backup Schedule with the following options:
  - Backup Configuration: Full server (recommended).
  - Backup Time: **Once a day, 1:00 AM.**
  - Destination Type: **Back up to a shared network folder**
  - Remote Shared Folder: **\\LON-DC1\Backup.**
    - Register Backup Schedule: Username: **Administrator**
    - Password: **Pa\$\$w0rd**
3. Close Windows Server Backup.

### ► Task 3: Complete an on-demand backup

To prepare for this task, you need to create a folder on LON-SVR1, with a name **Financial Data** on drive **C:** and within **Financial Data** folder you need to create a text file with a name **Financial Report.txt**.

To complete an on-demand backup, perform the following steps:

1. On LON-SVR1, start Windows Server Backup.
2. Run the Backup Once Wizard to back up the C:\Financial Data folder to the remote folder, \\LON-DC1\Backup.

**Results:** After completing this exercise, you will have installed the Windows Server Backup feature, configured a scheduled backup, and ran an on demand backup.

## Exercise 3: Restoring files by using Windows Server Backup

### Scenario

To ensure that the financial data can be restored, you must validate the procedure for restoring the data to an alternative location. You may also have to restore different versions of the data. For this purpose, you may have to use the Vssadmin tool to review backups.

The main tasks for this exercise are as follows:

1. Delete a file from the file server.
2. View the available restores by using the Vssadmin command.
3. Restore the file from backup.

### ► Task 1: Delete a file from the file server

- On LON-SVR1, delete the C:\Financial Data folder.

### ► Task 2: View the available restores by using the Vssadmin command

1. On LON-SVR1, run Windows PowerShell.
2. At the Windows PowerShell prompt, run **Vssadmin list shadows** command to list existing volume shadow copies.

### ► Task 3: Restore the file from backup

1. In the Windows Server Backup MMC, run the Recovery Wizard and specify the following information:
  - Getting Started: **A backup stored on another location**
  - Specify Location type: **Remote Shared Folder**
  - Specify Remote Folder: **\\LON-DC1\Backup**
  - Select Backup Date: Default value, **Today**
  - Select Recovery Type: Default value, **Files and Folders**
  - Select Items to Recover: **LON-SVR1\Local Disk (C:)\Financial Data**
  - Specify Recovery Options: **Another Location (C:)**
2. Locate **C:\** and ensure that the files are restored.

**Results:** After completing this exercise, you will have deleted a folder to simulate data loss, viewed available resources, and then restored the folder the backup that you created.

## Exercise 4: Implementing Microsoft Online Backup and Restore

### Scenario

A. Datum has to protect critical data in small branch offices. Those offices do not have backup hardware and full data center infrastructure. Therefore A. Datum has decided to back up the critical data in branch offices to a cloud-based service by using Microsoft Online Backup Service in Windows Server 2012.

The main tasks for this exercise are as follows:

1. Install the Microsoft Online Backup Service component.
2. Register the server with Microsoft Online Backup.
3. Configure an online backup.
4. Restore files by using the online backup.
5. Unregister the server from the Microsoft Online Backup Service.

### ► Task 1: Install the Microsoft Online Backup Service component

1. On LON-SVR1, in drive E, locate the installation file of the Microsoft Online Sign-in Assistant, **msoidcli.msi**. Install the application.
2. On LON-SVR1, in drive E, locate the installation file of the Microsoft Online Backup Agent, **OBSInstaller.exe**.
3. Start the installation of Microsoft Online Backup Agent by double-clicking the installation file **OBSInstaller.exe**.
4. Complete the setup by specifying the following information:
  - Installation Folder: **C:\Program Files**
  - Cache Location: **C:\Program Files\Microsoft Online Backup Service Agent**
  - Microsoft Update Opt-In: **I don't want to use Microsoft Update.**

5. Verify the installation; ensure you receive the following message: **Microsoft Online Backup Service Agent installation has completed successfully**. Clear the **Check for newer updates** check box, and then click **Finish**.
6. On the **Start** screen, verify the installation by clicking **Microsoft Online Backup Service** and **Microsoft Online Backup Service Shell**.

### ► Task 2: Register the server with Microsoft Online Backup

Before you start this task, you should rename **LON-SVR1** to *YOURCITYNAME-YOURNAME*, for example **NEWYORK-ALICE**. This is because this exercise will be performed online, and therefore the computer names used in this lab should be unique. If there is more than one student in the classroom with the same name, add a number at the end of the computer name, such as **NEWYORK-ALICE-1**.

To rename **LON-SVR1**, perform the following steps:

1. In the Server Manager window, rename **LON-SVR1** as *YOURCITYNAME-YOURNAME*, and then restart *YOURCITYNAME-YOURNAME*.
2. Wait until *YOURCITYNAME-YOURNAME* is restarted, and then log on as **Adatum\Administrator** with password **Pa\$\$w0rd**.

To register the server with Microsoft Online Backup, perform the following steps:

1. In the Microsoft Online Backup Service console, register LON-SVR1 by specifying the following information:
  - Account Credentials:
    - Username: **holuser@onlinebackupservice.onmicrosoft.com**,
    - Password: **Pa\$\$w0rd**
  - **Note:** In real-life scenario, you would type username and password of your Microsoft Online Backup Service subscription account.
  - Encryption Settings:
    - Enter passphrase: **Pa\$\$w0rdPa\$\$w0rd**
    - Confirm passphrase: **Pa\$\$w0rdPa\$\$w0rd**
2. Verify that you receive the following message: **Microsoft Online Backup Service is now available for this server**.

### ► Task 3: Configure an online backup

1. Switch to the Microsoft Online Backup Service console.
2. Configure an online backup by using the following options:
  - Select Items to back up: **C:\Financial Data**
  - Specify Backup Time: **Saturday, 1:00AM**
  - Specify Retention Setting: Default values
3. In the Microsoft Online Backup Service console, start the backup by clicking **Backup Now**.

► **Task 4: Restore files by using the online backup**

1. Switch to the Microsoft Online Backup Service console.
2. Restore files and folders by using the **Recover Data** option and specify the following information:
  - Identify the server on which the backup was originally created: **This server**
  - Select Recovery Mode: **Browse for files**
  - Select Volume and Date: **C:\** and **date and time of the latest backup**.
  - Select Items to Recover: **C:\Financial Data**
  - Specify Recovery Options: **Original location** and **Create copies so that you have both versions**

► **Task 5: Unregister the server from the Microsoft Online Backup Service**

1. Switch to the Microsoft Online Backup Service console.
2. Unregister the server from the Microsoft Online Backup Service using the following credentials:
  - Username: **holuser@onlinebackupservice.onmicrosoft.com**,
  - Password: **Pa\$\$w0rd**

**Results:** After completing this exercise, you will have installed the Microsoft Online Backup Service agent, registered the server with Microsoft Online Backup Service, configured a scheduled backup, and performed a restore by using Microsoft Online Backup Service.

► **Task: To prepare for next module**

When you finish the lab, revert the virtual machines to their initial state. To do this, complete the following steps.

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20417A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20417A-LON-SVR1** and **MSL-TMG1**.

## Module Review and Takeaways

### Review Questions

**Question:** Why is monitoring important?

**Question:** You want to create a strategy on how to back up different technologies that are used in your organization such as DHCP, DNS, Active Directory, and SQL Server. What should you do?

**Question:** How frequently should we perform backup on critical data?

### Best Practices

- Create an end-to-end monitoring strategy for your IT infrastructure. Monitoring should focus on proactively detecting potential failures or performance issues.
- When monitoring, estimate the baseline of system utilizations for each server. This will help you determine whether the system is performing well or is overused.
- Analyze your important infrastructure resources and mission-critical and business-critical data. Based on that analysis, create a backup strategy that will protect the company's critical infrastructure resources and business data.
- Identify with the organization's business managers the minimum recovery time for business-critical data. Based on that information, create an optimal restore strategy.
- Always test backup and restore procedures regularly, even if data loss or system failures never occur. Perform testing in a non-production and isolated environment.

### Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
During monitoring, multiple sources are concurrently reporting different problems.	
The server has suffered a major failure on its components.	
You must have a way to back up and restore your data quickly on a different company's locations. You do not have backup media or backup hardware in each site	
You must restore your data because of failure of the disk system. However, you find that your backup media is corrupted.	

### Real-world Issues and Scenarios

Your organization needs information on which data to back up, how frequently to back up different types of data and technologies, where to store backed up data (onsite or in the cloud), and how fast they can restore backed up data if a failure were to occur? Also, what is your suggestion to improve your organization's ability to efficiently restore data when it is necessary?



**Tools**

Tool	Use for	Where to find it
Server Manager Dashboard	Monitoring multiple servers	Server Manager
Performance Monitor	Monitoring services and application and hardware performance data	Server Manager/Tools
Resource Monitor	Controlling how your system resources are being used by processes and services	Server Manager/Tools
Windows Server Backup	Performing on demand or scheduled backup and restoring data and servers	Server Manager/Tools
Microsoft Online Backup Service	Performing on demand or schedule backup to the cloud and restoring data from the backup located in the cloud	Server Manager/Tools

**MCT USE ONLY. STUDENT USE PROHIBITED**

# Module 3

## Managing Windows Server 2012 by Using Windows PowerShell 3.0

### Contents:

Module Overview	3-1
<b>Lesson 1:</b> Overview of Windows PowerShell 3.0	3-2
<b>Lesson 2:</b> Using Windows PowerShell 3.0 to Manage AD DS	3-9
<b>Lesson 3:</b> Managing Servers by Using Windows PowerShell 3.0	3-20
<b>Lab:</b> Managing Servers Running Windows Server 2012 by Using Windows PowerShell 3.0	3-26
Module Review and Takeaways	3-31

## Module Overview

Windows PowerShell is a core feature of Windows Server® 2012 that enables command line management and configuration of the operating system. It is a standardized, task-based command-line shell and scripting language that offers administrators more flexibility and choice in how they manage computers running Windows®.

Windows PowerShell 3.0, included in Windows Server 2012, has more functionality and features than earlier versions. You can now use Windows PowerShell® to manage all the Windows Server roles and features. This enables administrators to quickly automate configuration tasks with a single tool, instead of having to use multiple tools, such as batch scripts, Microsoft Visual Basic® Script Edition scripts (VBScript), and manual configuration steps.

In this module, you will learn key Windows PowerShell concepts and new Windows PowerShell 3.0 features. This module will also describe how to practically use Windows PowerShell in your daily activities.

### Objectives

After completing this module, you will be able to:

- Describe the Windows PowerShell command-line interface.
- Use Windows PowerShell to manage Active Directory® Domain Service (AD DS).
- Manage servers by using Windows PowerShell.

## Lesson 1

# Overview of Windows PowerShell 3.0

As a Windows Server administrator, you can use Windows PowerShell to install and configure native Windows Server 2012 roles and features and to administer software such as Microsoft Exchange Server and Microsoft System Center 2012. Although you can use a graphical user interface (GUI) for administration, using Windows PowerShell with these applications enables bulk administration. This provides the ability to create automation scripts for administration and access to configuration options that are not available when you use a GUI. Some tasks that you can perform in Windows PowerShell will already be familiar to you, such as listing the contents of a directory. To use Windows PowerShell effectively, you must have a basic understanding of Windows PowerShell.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe Windows PowerShell.
- Describe the Windows PowerShell syntax.
- Describe cmdlet aliases.
- Use the Windows PowerShell Integrated Scripting Environment (ISE).
- Access Help in Windows PowerShell.
- Describe Windows PowerShell modules.
- Describe Windows PowerShell remoting.
- Describe the new features in Windows PowerShell 3.0.

### What Is Windows PowerShell?

Windows PowerShell is a command-line management interface that you can use to configure Windows Server 2012 and products such as System Center 2012, Exchange Server 2010, and Microsoft SharePoint® Server 2010. This management interface provides an alternative to the GUI management that enables administrators to:

- Create automation scripts.
- Perform batch modifications.
- Access settings that might be unavailable or more difficult to configure in the GUI.

- PowerShell is an object-based management environment
- PowerShell is an engine that enables administrators to:
  - Create automation scripts
  - Perform batch modifications
  - Access unavailable settings
- PowerShell provides a foundation upon which the GUI-based administrative tools of Microsoft can rest:
  - Actions can be accomplished in its command-line console
  - Actions can be invoked within GUIs by running PowerShell commands in the background

A GUI can guide you through complex operations, and can help you understand your choices and. However, a GUI can be inefficient for tasks that you have to perform repeatedly, such as creating new user accounts. By building administrative functionality in the form of Windows PowerShell commands, Microsoft lets you select the right method for a given task.

As you become more comfortable with Windows PowerShell, you may use it in place of other low-level administrative tools that you may have used. For example, Windows PowerShell has access to the same features that VBScript does, but in many cases provides easier ways to perform the same tasks.

Windows PowerShell may also change the way you use Windows Management Instrumentation (WMI). Windows PowerShell can wrap task-specific commands around the underlying WMI functionality. When you use Windows PowerShell with WMI, your work is simplified because Windows PowerShell provides easy to use, task-based commands.

## Windows PowerShell Syntax

Windows PowerShell has rules for naming and implementing functions. For example, Windows PowerShell commands, known as *cmdlets*, use a naming convention of verb or action, followed by a hyphen and a noun or subject. For example, to retrieve a list of virtual machines (VMs), you would use the cmdlet **Get-VM**. This standardization helps you more easily learn how to perform administrative tasks. For example, to change settings of a VM, you would use the cmdlet **Set-VM**.

Optionally, one or more parameters can be used with a cmdlet to modify its behavior or specify settings. Parameters are written after the cmdlet. Each parameter that is used is separated by a space, and begins with a hyphen. Not all cmdlets use the same parameters. Some cmdlets have parameters that are unique to its functionality. For example, the **Move-Item** cmdlet has the *Destination* parameter to specify the location to move the object; whereas the **Get-ChildItem** has the *-Recurse* switch parameter. There are several kinds of parameters, including the following:

- **Named.** Named parameters are most common. They are parameters that can be specified and require a value or modifier. For example, by using the **Move-Item** cmdlet, you would specify the *-Destination* parameter along with the exact destination to move the item.
- **Switch.** Switch parameters modify the behavior of the cmdlet, but do not require any additional modifiers or values. For example, you can specify the *-Verbose* parameter without specifying a value of \$True.
- **Positional.** Positional parameters are parameters that can be omitted and can still accept values based on where the information is specified in the command. For example, you could run **Get-EventLog -EventLog System** to retrieve information from the System event log. However, because the *-EventLog* positional parameter accepts values for the first position, you can also run **Get-EventLog System** to get the same results. When the *-EventLog* parameter is not present, the cmdlet still accepts the value of System because it is the first item after the cmdlet name.

Parameters that are common to many cmdlets include options to test the actions of the cmdlet or to generate verbose information about the execution of cmdlet. Common parameters include:

- **-Verbose.** This parameter displays detailed information about the performed command. You should use this parameter to obtain more information about the execution of the command.
- **-WhatIf.** This parameter displays the outcome of running the command without running it. This is helpful when testing a new cmdlet or script and you do not want the cmdlet to run.
- **-Confirm.** This parameter displays a confirmation prompt before executing the command. This is helpful when you are running scripts and you want to prompt the user before executing a specific step in the script.

- Verb-Noun pair naming is as follows:

Verb	Noun	Cmdlet
Get	EventLog	Get-EventLog
Set	ExecutionPolicy	Set-ExecutionPolicy
New	VM	New-VM

- Use cmdlet parameters to modify actions and provide configuration information. Parameters include:

- **Named.** *-EventLog System, -UserName John*
- **Switch.** *-Verbose, -Debug, -Confirm*

- **Positional.**

- *Get-EventLog System*
- *Get-EventLog -LogName System*

- Common parameters: *-WhatIf, -Debug, -Verbose, -Confirm*



### Additional Reading: Cmdlet Verbs

[http://msdn.microsoft.com/en-us/library/windows/desktop/ms714428\(v=vs.85\).asp](http://msdn.microsoft.com/en-us/library/windows/desktop/ms714428(v=vs.85).asp)

## Cmdlet Aliases

Although the standard naming convention used by cmdlets facilitate learning, the names themselves can be very long, and sometimes do not match common terminology associated with performing a task. For example, you may be familiar with the **dir** command which lists the contents of a directory (or folder). The Windows PowerShell cmdlet for this task, however, is **Get-ChildItem**. To make using cmdlets easier, Windows PowerShell enables aliases to be created for cmdlets. There is an alias created by default for **dir** that points to **Get-ChildItem**.

You can use aliases for:

- Backward compatibility
- Shorten scripts
- Easier discoverability

Common Aliases:

- **cd** -> Set-Location
- **dir** -> Get-Child-Item
- **ls** -> Get-Child-Item
- **copy** -> Copy-Item
- **kill** -> Stop-Process
- **rm** -> Remove-Item
- **type** -> Get-Content
- **help** -> Get-Help

You can create new aliases for your common cmdlets, scripts, and programs by using the **New-Alias** cmdlet. Default aliases include:

- **cd** -> **Set-Location**
- **copy** -> **Copy-Item**
- **kill** -> **Stop-Process**
- **move** -> **Move-Item**
- **rm** -> **Remove-Item**
- **type** -> **Get-Content**
- **help** -> **Get-Help**

## Demonstration: Using the Windows PowerShell ISE

The Windows PowerShell ISE application is a graphical tool that enables you to write and test Windows PowerShell scripts similar to the way a developer would write an application by using Microsoft Visual Studio®. The Windows PowerShell ISE for Windows PowerShell 3.0 includes IntelliSense to provide instance suggestions on the correct script syntax and available cmdlet parameters. Windows PowerShell ISE is divided into two main parts: the Script pane and the Console pane.

### Demonstration Steps

1. Logon to LON-DC1 as the domain administrator.
2. Open Windows PowerShell ISE as an administrator and review the Script pane and the Console pane.
3. Follow the steps in the following demonstration script: **E:\ModXA\Democode\Using Windows PowerShell ISE.ps1**.

## Accessing Help in Windows PowerShell

Whether you are an experienced professional or new to Windows PowerShell, the cmdlet Help documentation is rich source of information. To access the Help documentation, use the **Get-Help** cmdlet or its alias *help* followed by the cmdlet name. **Get-Help** has parameters to adjust the Help content that is displayed. The parameters are:

- *-Detailed*. This parameter displays more detailed help than the default option.
- *-Examples*. This parameter displays only the examples for using the cmdlet.
- *-Full*. This parameter displays detailed help and usage examples.
- *-Online*. This parameter opens a Web browser to the cmdlet documentation on the Microsoft website.

Windows PowerShell 3.0 includes the ability to download the latest help document from Microsoft for use locally. To do this, use the **Update-Help** cmdlet. Also, new in Windows PowerShell 3.0 is the **Show-Command** cmdlet. This helps PowerShell beginning users interact with the input and output options for a cmdlet by using a graphical interface.

The **Get-Command** cmdlet returns a list of all locally available cmdlets, functions, and aliases. You can use it to discover new cmdlets by using wildcard searches. For example, to return a list of all cmdlets that include VM in them, you could run **Get-Command \*VM\***.

- To access the Help documentation, run **Get-Help** or the alias *help* followed by the cmdlet name:

```
Get-Help Get-EventLog
Get-EventLog -help
```

- **Get-Help** has parameters to adjust the amount of help displayed. The parameters are:
  - *-detailed*
  - *-examples*
  - *-full*
  - *-online*
- Other cmdlets that you can use for accessing help: **Update-Help**, **Show-Command**, **Get-Command**, and tab completion

## Using Windows PowerShell Modules

Windows PowerShell is designed to be extensible. Adding new cmdlets and functions in Windows PowerShell 3.0 is performed in part through modules.



**Note:** In earlier versions of Windows PowerShell, extensibility was provided by using snap-ins. For backward compatibility, Windows PowerShell 3.0 continues to support snap-ins.

Windows PowerShell uses the

**Microsoft.PowerShell.Management** module which provides basic functionality. When you install additional roles on a server, additional Windows PowerShell modules are installed and registered. For example, you install the Microsoft Hyper-V® Role and also choose to install the Hyper-V module for Windows PowerShell. To manage Hyper-V from Windows PowerShell, you must import the Hyper-V module into the Windows PowerShell session. To import the Hyper-V module, run the following command:

```
Import-Module Hyper-V
```

### Windows PowerShell is extended through modules

- You can import modules by using the **Import-Module** cmdlet:

```
Import-Module Hyper-V
```

- You can list loaded modules by running the following command:

```
Get-Module
```

- Modules can be of the following types:

- Script
- Binary

Run the following command to list all modules that are imported:

```
Get-Module
```

It is not always necessary to manually import modules. For example, the Windows PowerShell module for Exchange Server 2010 is automatically imported during product installation. However, if you cannot run cmdlets for a specific Windows Role or application, it may indicate that you have to import the appropriate Windows PowerShell module.

There are two basic module types:

- **Binary.** A binary module is created by using the .NET Framework and is frequently provided with a product to provide Windows PowerShell support. Binary modules many times add cmdlets that consists of noun or subject types that are newly created in the AD DS schema to support the product. An example is the **New-Mailbox** cmdlet of Exchange Server 2010.
- **Script.** A script module is composed of Windows PowerShell cmdlets that already exist in the environment. These scripts can provide additional functions and variables to automate repetitive or tedious tasks. You may want to create your own module that includes functions or variables specific to your environment as a timesaving or configuration management measure.



#### **Additional Reading:** Windows PowerShell Modules

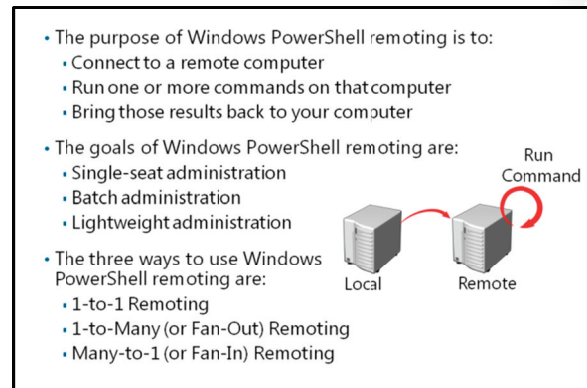
[http://msdn.microsoft.com/en-us/library/windows/desktop/dd878324\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/windows/desktop/dd878324(v=vs.85).aspx)

## What Is Windows PowerShell Remoting?

The purpose of Windows PowerShell remoting is to connect to remote computers, to run commands on those computers, and to direct the results back to your local computer. This enables single-seat administration, or the ability to manage the computers on the network from the client computer, instead of having to physically visit each computer. A key goal of Windows PowerShell remoting is to enable batch administration, which lets you run commands on a whole set of remote computers concurrently.

There are three main ways to use remoting:

- **One-to-One remoting.** In this scenario, you connect to a single remote computer and run shell commands on it, exactly as if you had logged into the console and opened a Windows PowerShell window.
- **One-to-Many remoting, or Fan-Out remoting.** In this scenario, you issue a command that will be executed on one or more remote computers in parallel. You are not working with each remote computer interactively. Instead, your commands are issued and executed in a batch and the results are returned to your computer for your use.
- **Many-to-One remoting, or Fan-In remoting.** In this scenario, multiple administrators make remote connections to a single computer. Typically, those administrators will have different permissions on the remote computer and might be working in a restricted runspace within the shell. This scenario usually requires custom development of the restricted runspace and will not be covered further in this course.





Remoting requires both Windows PowerShell and Windows Remote Management (WinRM) utilities on your local computer and on any remote computers to which you want to connect. WinRM is a Microsoft implementation of Web Services for Management, or WS-MAN, which is a set of protocols that is widely-adopted across different operating systems. As the name implies, WS-MAN and WinRM use web-based protocols. An advantage to these protocols is that they use a single, definable port. This makes them easier to pass through firewalls than older protocols that randomly selected a port. WinRM communicates by using the Hypertext Transfer Protocol (HTTP). By default, WinRM and Windows PowerShell remoting uses TCP port 5985 for incoming connections that are not encrypted and TCP port 5986 for incoming encrypted connections. Applications that use WinRM, such as Windows PowerShell, can also apply their own encryption to the data that is passed to the WinRM service. WinRM supports authentication and, by default, uses the Active Directory native Kerberos protocol in a domain environment. Kerberos does not pass credentials over the network and it supports mutual authentication to ensure that incoming connections are coming from valid computers.

Establishing a One-to-One remoting session by using Windows PowerShell ISE is performed by clicking the **New Remote PowerShell** tab on the **File** menu. You can also establish a remote Windows PowerShell session by using the **Enter-PSSession** cmdlet. For example, to open a Remote PowerShell session on a computer named LON-SVR2, you would use the following syntax:

```
Enter-PSSession -ComputerName LON-SVR2
```

One-to-Many remoting is primarily performed by using the **Invoke-Command** cmdlet. To run the **Get-EventLog** cmdlet against the computers named LON-SVR1 and LON-SVR2, use the following command:

```
Invoke-Command -ScriptBlock { Get-EventLog System -Newest 5 } -Computers LON-SVR1, LON-SVR2
```



**Note:** Unlike in earlier versions, Windows Server 2012 has Windows PowerShell remoting and WinRM enabled by default.

## What Is New in Windows PowerShell 3.0?

Windows PowerShell 3.0 has new features that facilitate managing larger groups of servers through better scaling, additional functionality, and better management. Windows PowerShell 3.0 includes the following new features:

- *Windows PowerShell Workflow*. This enables coordination of complex parallel and sequenced commands.
- *Windows PowerShell Web Access*. This feature enables encrypted and authenticated access to Windows PowerShell by using a Web browser on any device.
- *Scheduled Jobs*. This feature enables scheduling of Windows PowerShell commands and scripts to automatically run administrative tasks.

### Windows PowerShell 3.0 improvements include:

- Over 260 core cmdlets
- Management of all Windows Roles and Features
- Windows PowerShell Workflow
- Windows PowerShell Web Access
- Scheduled Jobs
- Enhanced Online Help
- ISE Autosense
- Robust Session Connectivity

- *Enhanced Online Help.* You can now download the latest Help files from Microsoft by using the **Update-Help** cmdlet and view the latest help online. This guarantees you are getting the latest information about how to use Windows PowerShell.
- *Windows PowerShell ISE Autosense.* Windows PowerShell ISE provides hints for cmdlets, including valid parameters that make it easier than ever to use Windows PowerShell.
- *Robust Session Connectivity.* These connections enable you to connect to a remote server and if connectivity is lost or you intentionally disconnect, you can resume the connection at the point it was disconnected. Previously, if connection to a session was lost, all the session data, variables, and command history would be lost.

## Lesson 2

# Using Windows PowerShell 3.0 to Manage AD DS

Active Directory is the technology that many administrators spend most of their time using, completing day-to-day administrative tasks such as adding users and updating directory objects. With the number of Active Directory–focused cmdlets in Windows Server 2012, those administrators can save time and energy by using Windows PowerShell to automate many of their more time-consuming or repetitive tasks. Automation can also help improve security and consistency because it is less prone to repeated human error than manual administration. If you are already comfortable performing common Active Directory administrative tasks in other tools, you should quickly be able to learn to perform equivalent tasks in Windows PowerShell.

This lesson will help you understand the approach used by the Active Directory cmdlets. It will help you develop the skills that you must have to discover, explore, learn, and use other add-in commands, whether they are included with Windows Server 2012 or with another Microsoft or third-party software product.

### Lesson Objectives

After completing this lesson, students will be able to:

- Describe the Active Directory modules for Windows PowerShell.
- Describe how to use variables.
- Describe how to use pipelines and scripts.
- Describe how to format output from a Windows PowerShell command.
- Describe how to create and run Windows PowerShell scripts.
- Describe how to use Windows PowerShell loops and conditional expressions.
- Manage AD DS with Windows PowerShell.
- Describe how to obtain the Windows PowerShell history information from Active Directory Administrative Center.

## Using the Active Directory Module for Windows PowerShell

You may be comfortable managing AD DS by using the common graphical tools such as Active Directory Users and Computers. Another option that you may not be as comfortable with is the Windows PowerShell cmdlets. Using the AD DS cmdlets to perform common tasks will help you learn how to use Windows PowerShell.

The Active Directory PowerShell module included in Windows Server 2012, provides over 130 cmdlets for managing Active Directory objects such as computer and user accounts, groups, trusts, and policies.

**The Active Directory PowerShell Module included in Windows Server 2012, provides over 130 cmdlets for managing Active Directory objects, such as:**

- Computer Accounts
- User Accounts
- Service Accounts
- Groups
- Organizational Units
- Replication
- Trusts
- Central Access Policies
- Password Policies

## Using Windows PowerShell Variables

Windows PowerShell enables you to retrieve, modify, and filter data from many different sources. In some cases, you may want to store data for comparison or use. For example, you may want to retrieve a list of the members of a particular security group and then modify the description field of each of the users. Variables are used to store and retrieve data in memory during a Windows PowerShell session. A variable always begins with a dollar (\$) sign and can then be named with descriptive text or numbers, such as \$Variable1, \$x, and \$MemberList. Windows PowerShell variables are typed. This means that they are created to store a specific type of data whether it is text, numbers, objects, time, arrays, or other defined object.

- A variable is a temporary holding place in memory for a value, object, or collection of objects
- Variables are named, and their names are preceded with a dollar sign



You can declare a variable in one of two ways, the first of which is using the **Set-Variable** cmdlet. For example to declare a variable named \$ADDS and assign it the object returned from **Get-ADDomain** by using the **Set-Variable** cmdlet, use the following command:

```
Set-Variable -Name ADDS -Value (Get-ADDomain)
```

You will notice you do not specify the \$ symbol when you use the **Set-Variable** cmdlet to declare variables. The second way to create a variable is by declaring it, and then assigning a value to it. To do this, start the command with the name of the variable followed by an equal sign and then the command, commands, or value to assign. For example to declare a variable named \$ADDS and assign it the object returned from **Get-ADDomain** use the following command:

```
$ADDS = Get-ADDomain
```

The \$ADDS variable now holds a copy of the object output by the **Get-ADDomain** cmdlet. The output object takes on the type that is defined in the relevant class and the variable maintains that structure. You can now read and manipulate the variable as similar to how you would a .NET object. To obtain information about the properties or to run methods, you can use dotted notation on the variable. For example, to determine the domain functional level reported by the **DomainMode** property of **Get-ADDomain**, you can use the following command:

```
> $ADDS.DomainMode
Windows2008R2Domain
```

You can also access methods or actions from a variable. For example, to determine the **BaseType** of \$ADDS, you can use the **GetType()** method by running the following command:

```
> $ADDS.GetType().BaseType
Microsoft.ActiveDirectory.Management.ADPartition
```

When you use methods, you must follow the method with () to distinguish that it is a method and not a property. You can also use variables in calculations, for example, you can add the contents of two variables. To declare two variables and then add them together, use the following commands:

```
> $A = 1
> $B = 2
> $A + $B
3
```

When you use variables in calculations, make sure that they are typed correctly because typing them incorrectly could lead to unexpected results. For example, notice when variables are typed as string data instead of numbers:

```
> $C = "3"
> $D = "4"
> $C + $D
34
```

Instead of adding the two values numerically, they are concatenated together. When you mix types together, there is more potential for unexpected results because Windows PowerShell will automatically cast or convert some data types. For example, see how the data is cast in the following example:

```
> $A + $C
4
> $C + $A
31
```

In these examples, the type of the first variable is used to cast the other variables for the calculation. To better control how data is cast, you can specify the data type for each variable. To control how each variable is cast, see the following example:

```
> [string] $A + $C
13
> [int] $C + $A
2
```



**Additional Reading:** about\_Variables

<http://technet.microsoft.com/en-us/library/dd347604.aspx>

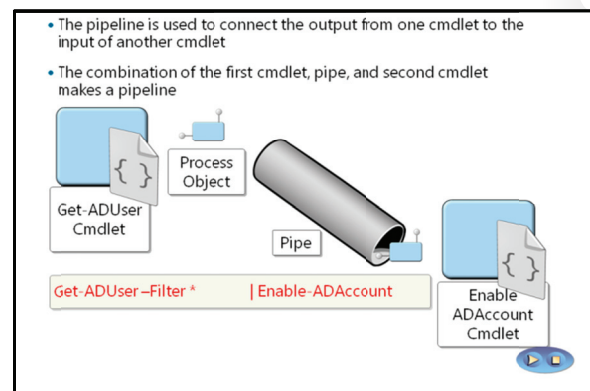
**Question:** How do you declare variables and assign values to them?

## The Windows PowerShell Pipeline

Windows PowerShell is an object-based environment. This means that the input and outputs of the cmdlets are objects that can be manipulated. In some instances, you may want to take the output of one cmdlet and pass it to another cmdlet for additional actions. For example, when you have to enable all disabled AD DS accounts in the domain, you could manually list each user by using the **Get-ADUser** cmdlet. Then by using Windows PowerShell, you can use the **Enable-ADAccount** cmdlet for each locked user account. To make this easier, you can

directly pass the output data from one cmdlet into another cmdlet, which is called piping. Piping is performed by putting the pipe (|) character between cmdlets. Each cmdlet is executed from the left to the right, each passing its output to the next cmdlet in line. For example, you can get a list of all users in the domain and then pipe the list to the **Enable-ADAccount** cmdlet, by running the following command:

```
Get-ADUser -Filter * | Enable-ADAccount
```



Piping can be used extensively in Windows PowerShell as it is in other shells. Windows PowerShell differs from typical shells because the data in the pipeline is an object instead of just simple text. Having an object in the pipeline enables you to easily persist all the properties of the returned data. The data in the pipeline is assigned to a special variable named `$_` which only exists while the pipeline is executing. For example, if you want to enable accounts that are disabled, you can use the **Where-Object** cmdlet to return only accounts are disabled. To do this, run the following command:

```
Get-ADUser | Where-Object {$_.Enabled -eq $false} | Enable-ADAccount
```

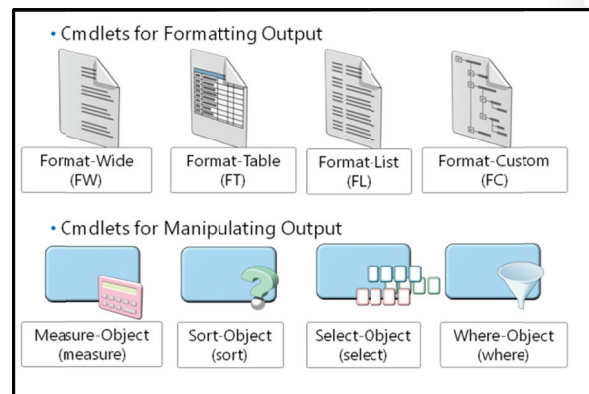
By piping an object with a list of all the users, you can use the **Where-Object** cmdlet to filter the accounts that are disabled based on the Enabled property of the account.



**Note:** This example is for teaching purposes only. It enables all the disabled accounts in the domain and should not be performed in a production environment because this may enable accounts that should remain disabled.

## Options for Formatting Windows PowerShell Output

When you work with AD DS data, you may have to retrieve lists of users, computers, or groups and have to visualize the data by using a tool such as Microsoft Office Excel® or you may have to view only the specific properties on screen. Windows PowerShell enables both such scenarios. First formatting data for viewing on screen. There are several default cmdlets available to control how data is formatted. These cmdlets are described in the following table.



Cmdlet	Description
<b>Format-List</b>	This cmdlet outputs data in a list format with each property on its own line. You can specify the properties that you want displayed by using the <i>-Property</i> parameter. You can call this cmdlet by using the alias of FL. This cmdlet is useful when you view a small number of objects with a large number of properties.
<b>Format-Table</b>	This cmdlet outputs data in a table format with each property as its own column. You can specify the properties that you want displayed by using the <i>-Property</i> parameter. You can call this cmdlet by using the alias of FT. This cmdlet is useful when you view a large number of objects with a small number of properties.
<b>Format-Wide</b>	This cmdlet outputs data in a table format with only one property for each object. You can specify the property that you want displayed by using the <i>-Property</i> parameter and the number of columns to display the data by using the <i>-column</i> parameter. You can call this cmdlet by using the alias of FW. This cmdlet is useful when you view a large number of objects and you only need to see one property for each object such as the name.

Cmdlet	Description
<b>Format-Custom</b>	This cmdlet outputs data in a format previously defined by using a PS1XML file. The settings in this file can specify which properties to show and how to arrange and group them. You can call this cmdlet by using the alias of FC. This cmdlet is useful when you view data that you access frequently and have to customize which properties are shown.

Another set of cmdlets enable complex formatting and reporting. These are listed in the following table.

Cmdlet	Description
<b>Measure-Object</b>	This cmdlet takes the input object from the pipelines or variable and performs calculations on specified properties and on text in strings and files. Calculations include counting objects, determining the average, minimum, maximum, and sum of property values. It can also count the number or occurrences of words and characters in a file or string. It is used when you have to quickly calculate the number of users selected as part of a query or determining the memory a set of processes is using.
<b>Select-Object</b>	This cmdlet takes the input object from the pipeline or variable and outputs objects that have only the selected properties. It can also select a subset of items in each object by using the <i>-First</i> , <i>-Last</i> , <i>-Unique</i> , and <i>-Index</i> parameters, which is valuable when you work large datasets.
<b>Sort-Object</b>	This cmdlet takes the input object from the pipeline or variable and sorts the data based on the selected properties. This is helpful when you have to provide a sorted list of data.
<b>Where-Object</b>	This cmdlet takes the input object from the pipeline or variable and then applies a filter that is based on a specified query. The queries used for filtering are enclosed in braces and include a comparison. This is helpful when you have to select specific types of data.

You can use all these cmdlets together to create customized output to the screen. You can also use the Out-File to write the output to a text file, or Export-Csv to export the data as a comma separated values (CSV) file.

## Creating and Running Windows PowerShell Scripts

You can perform complicated multi-step tasks by using a pipeline and multiple cmdlets. There may be times where you have to run multiple functions, make choices, wait for tasks to complete, or run the same code repeatedly. In these cases, you can use a Windows PowerShell script to put all the steps together. A script is a text-based file that includes at least one Windows PowerShell command and saved with a .PS1 file name extension. Scripts can be created to take input from the command line letting you customize how the script executes.

- Execution policy restricts script execution, the execution policies include:
  - Restricted
  - AllSigned
  - RemoteSigned
  - Unrestricted
  - Bypass
- Scripts are text files with a .ps1 extension
- Scripts contain one or more commands that you want the shell to execute in order
- Scripts, when run, require a relative or full path to be specified:

```
.\Get-LatestLogon.ps1
E:\Mod03\Democode\Get-LatestLogon.ps1
```



## Execution Policy

By default, the execution policy does not enable Windows PowerShell scripts to be executed automatically. This safeguards the computer from enabling unattended scripts to run without the administrator from knowing. There are four execution policies that can be set and are as follows:

- **Restricted.** This is the default policy for Windows Server 2012 and does not enable configuration files to load, nor does it enable scripts to be run. The Restricted execution policy is perfect for any computer for which you do not run scripts or for which you run scripts only rarely. (Be Aware That you could always manually open the shell with a less-restrictive execution policy.)
- **AllSigned.** This policy requires that all scripts and configuration files be signed by a trusted publisher, including scripts created on your local computer. This execution policy is useful for environments where you do not want to accidentally run *any* script unless it has an intact, trusted digital signature. This policy is less convenient because it requires you to digitally sign every script that you write, and re-sign each script every time that you make any changes to it.
- **RemoteSigned.** This policy requires that all scripts and configuration files downloaded from the Internet be signed by a trusted publisher. This execution policy is useful because it assumes that local scripts are ones that you create yourself, and you trust them. It does not require those scripts to be signed. Scripts that are downloaded from the Internet or received through e-mail, however, are not trusted unless they carry an intact, trusted digital signature. You could definitely still run those scripts—by running the shell under a lesser execution policy, for example, or even by signing the script yourself—but those are additional steps that you have to take, so it is unlikely that you would be able to run such a script accidentally or unknowingly.
- **Unrestricted.** This policy loads all configuration files and runs all scripts. If you run a script that was downloaded from the Internet, you are warned about potential dangers and must grant permission for the script to run. The Unrestricted execution policy is not usually appropriate for production environments because it provides little protection against accidentally or unknowingly running untrusted scripts.
- **Bypass.** This policy loads all configuration files and runs all scripts. If you run a script that was downloaded from the Internet, the script will run without any warnings. This execution policy is not usually appropriate for production environments because it provides no protection against accidentally or unknowingly running untrusted scripts.

You can view the execution policy for the computer by using the **Get-ExecutionPolicy** cmdlet. To configure the execution policy, you must open an elevated Windows PowerShell window and run the **Set-ExecutionPolicy** cmdlet. After the execution policy is configured, you can run a script by typing in the name of the script.

## Simple Scripts

Scripts are text files that have a .PS1 file name extension. These files contain one or more commands that you want the shell to execute in a particular order. You can edit scripts by using Notepad, but the Windows PowerShell ISE provides a better editing experience. In it, you can type commands interactively, obtain hints on the correct command syntax, and immediately see the results. You can then paste those results into a script for long-term use. Or you can type your commands directly into a script, highlight each command, and press F8 to execute only the highlighted command. If you are pleased with the results, you save the script and you are finished. Generally, there are very few differences between what you can do in a script and what you would do on the command line. Commands work in the same manner in a script. This means that a script can just be created by pasting commands that you have already tested at the command line. The following is a simple script in a text file that is named **Get-LatestLogon.ps1**.



```
# This script will return the last user who has logged on to the domain.
Get-ADUser -Filter * -Properties lastLogon | `
Sort-Object -Property lastLogon -Descending | `
Select-Object -first 5 | `
Format-Table name, `
@{Label="LastLogon";Expression={[datetime]::FromFileTime($_.lastLogon)}} `
-AutoSize
```

Although this script contains a single pipeline statement it is broken up by using the backtick (`) character. You can break up long lines of code by using the backtick character to make the script easier to read. Notice that the first line of this script starts with a hash mark (#). A line that begins with a hash mark will not be processed. Therefore, you can use start a line with a hash mark and write notes and comments about the script. To run a script, you must type either the full or the relative path of the script. For example, to run the **Get-LatestLogon.ps1** script, you can use either of the following options if the script is in your current directory or search path:

```
.\Get-LatestLogon.ps1
E:\ModXA\Democode\Get-LatestLogon.ps1
```

If the script name or path has spaces in it you have to enclose the name single or double quotation marks and echo the name to the console by using an ampersand (&) character. The following example shows how to do this by using both the relative and a full path.

```
& '.\Get Latest Logon.ps1'
& 'E:\ModXA\Democode\Get Latest Logon.ps1'
```

## Using Windows PowerShell Loops and Conditional Expressions

Advanced Windows PowerShell scripts may require repeating commands a certain number of times, until a specific condition is met, or only if a specific condition is met. These test conditions are defined by using comparison statements.

### Boolean Comparisons

Test, or comparison statements, are used as test conditions for loops and conditional constructs. These typically compare, either of two or more objects or two or more property values, and are designed to result in a True or False value. These comparisons are frequently known as *Boolean comparisons*, because they can only result in one of the two Boolean values, True or False. As part of designing a Windows PowerShell script using Boolean comparisons are common enough task: You might compare two computer names to see whether they are equal, or compare a performance counter value to a predetermined threshold value to see which of the two is greater. The comparison operators sit between the two items that you want to compare. You probably remember simple comparisons from grade school math with comparisons like  $10 > 4$ ,  $5 < 10$ , and  $15 = 15$ . Windows PowerShell performs comparisons the same way, although it has its own syntax. Some common comparison operators are as follows:

- -eq. Equal to
- -ne. Not equal to
- -le. Less than or equal to

```
foreach ($user in $group){
write-host $user " is in " $group}

if ($Today.DayOfWeek = "Monday") {
write-host "Today is Monday"}

while ($i -ne 25) {write-host $i "is not 25"}

for ($i=1; $i < 25; $i++) {
write-host $i "is not 25"}
```

- -ge. Greater than or equal to
- -gt. Greater than
- -lt. Less than

Windows PowerShell defines two special variables for comparisons, \$True, and \$False, which represent the Boolean values **true** and **false**. If a comparison is true, the expression is evaluated as \$True and if the comparison is not true, the expression is evaluated as \$False. For example, the comparison *4 is greater than 10* (*4 -gt 10*), will produce \$False as its result, whereas, *10 is equal to 10* (*10 -eq 10*) would produce \$True. Windows PowerShell enables you to execute comparisons right on the command line. Type your comparison and press Enter to see the result of the comparison. The real value of the Boolean comparisons are shown when they are used in loops and conditional expressions.

There are several Windows PowerShell constructs that make use Boolean comparisons to control the execution of code in a script. These constructs are **if**, **switch**, **for**, **while**, and **foreach**.

### The if Statement

The if statement can be used to execute a block of code if the specified criteria are met. The basic functionality of an if statement is shown in the following example:

```
if (Boolean comparison)
{
Code to complete if test expression is true
}
```

Another option available to allow for additional possibilities is using else and elseif statements. When you want to execute special code if a condition exists or execute other code if it does not exist, you can use the else. If there are additional conditions that you want to test for you could use the elseif statement consider the following example:

```
$Today = Get-Date
$Admin = Get-ADUser -Identity Administrator -Properties StreetAddress
Write-Host $Admin.Name "has an address of" $Admin.StreetAddress
if ($Today.DayOfWeek -eq "Monday")
{
Set-ADUser -Identity Administrator -StreetAddress "Headquarters"
}
elseif ($Today.DayOfWeek -eq "Thursday")
{
Set-ADUser -Identity Administrator -StreetAddress "London Office"
}
else
{
Set-ADUser -Identity Administrator -StreetAddress "Out of the Office"
}
# Confirm Settings were made
$Admin = Get-ADUser -Identity Administrator -Properties StreetAddress
Write-Host "Today is" $Today.DayOfWeek "and " $Admin.Name `
"is working from the" $Admin.StreetAddress
```

## The switch Statement

The switch statement is closely related to how ifelse statements work. The statement enables a single condition statement to have multiple options for execution. The switch statement has the following syntax:

```
switch (Value Testing)
{
Value 1 { Code run if value 1 condition exists}
Value 2 { Code run if value 2 condition exists}
Value 3 { Code run if value 3 condition exists}
default { Code run if no other condition exists}
}
```

Using the previous example, you can achieve the same functionality with less work as shown in this example:

```
$Today = Get-Date
$Admin = Get-ADUser -Identity Administrator -Properties StreetAddress
# Write current settings to console
Write-Host $Admin.Name "has an address of" $Admin.StreetAddress
switch ($Today.DayOfWeek)
{
"Monday" {Set-ADUser -Identity Administrator -StreetAddress "Headquarters"}
"Thursday" {Set-ADUser -Identity Administrator -StreetAddress `
    "London Office"}
default {Set-ADUser -Identity Administrator -StreetAddress `
    "Out of the office"}
}
# Confirm Settings were made
$Admin = Get-ADUser -Identity Administrator -Properties StreetAddress
Write-Host "Today is" $Today.DayOfWeek "and " $Admin.Name `
"is working from the" $Admin.StreetAddress
```

If a larger number of false statements are needed, the switch statement may be an easier option to use and debug.

## The for Loop

The for loop can be used to execute a block of code a specific number of times. This can be when multiple items have to be requested, or created. The for statement syntax is as follows:

```
for (setup loop variables ; Boolean comparison ; action after each loop)
{
Code to complete while Boolean comparison is true
}
```

The for loop begins with settings to configure variables, the Boolean comparison, and an action to complete after each loop. Consider the following example that creates five new computer accounts with unique names using a for statement:

```
# Create a variable named $i and assign it a value of 1
# Execute the for loop for as long as $i is less than 6
# After each loop add 1 to the value of $i
for ($i = 1 ; $i -lt 6 ; $i++)
{
# Create a variable with the name of the computer account
$ComputerAcct = "LON-SRV" + $i
New-ADComputer -Name $ComputerAcct
}
```

## The while Loop

The while loop can be used to execute a block of code while a specific condition exists and resembles the for loop, except that it does not have built in mechanisms to set up variables and actions to run after each loop. This enables the while statement to continue executing until a condition is met instead of a set number of times. The while statement syntax is as follows:

```
while (Boolean comparison)
{
Code to complete while Boolean expression is true
}
```

This script prints a random number on the screen until one of the random numbers is less than 50,000,000. The \$i variable's value must be set before the while loop so that the while loop executes as follows:

```
$i = 99999999999
while ($i -gt 50000000)
{
Write-Host "Random Value: " $i
$i = Get-Random
}
```

Also available is the do/while loop which works just as while loop however the Boolean expression is evaluated at the end of the loop instead of the beginning. This means that the code block in a do/while loop will always be executed at least one time. The value of \$i does not have to be set before the do/while loop because it is evaluated at the end of the loop. The following example shows a do/while loop:

```
do {
Write-Host "Random Value: " $i
$i = Get-Random
} while ($i -gt 50000000)
```

## The foreach Statement

The foreach statement iterates through an array (collection), item by item, assigning a specifically named variable to the current item of the collection. Then it runs the code block for that element.

```
foreach (item in collection)
{
Code to complete for each item in the collection.
}
```

Using the foreach statement can make batch modifications easier. Consider, for example, setting a description for all users who are members of a specific group, as shown in the following example:

```
# Get a list of the members of the Domain Admins group
$DAdmins = Get-ADGroupMember "Domain Admins"
# Go through each member and set the Description
foreach ($user in $DAdmins)
{
Set-ADUser $user -Description "In the Domain Admins Group"
}
```

## Demonstration: Managing AD DS by Using Windows PowerShell

In this demonstration, you will review how to manage users and group in Windows PowerShell.

### Demonstration Steps

1. Start and log on to LON-DC1. Log on as the domain administrator.
2. Open Windows PowerShell ISE as an administrator.
3. Refer to the demonstration script in virtual machine LON-DC1 at E:\ModXA\Democode\Managing Users and Groups.ps1.

## Active Directory Administrative Center Integration with Windows PowerShell

Active Directory Administrative Center is built on Windows PowerShell technology. It provides administrators the ability to perform enhanced data management by using a GUI. Using Active Directory Administrative Center, you can perform the following tasks:

- Manage user and computer accounts
- Manage groups
- Manage organizational units (OUs)
- Use build queries to filter Active Directory information

#### Active Directory Administrative Center:

- Allows management of user accounts, computer accounts, groups, and organizational units
- Provides a Windows PowerShell history of all commands used
- Is a Windows PowerShell learning tool

Because Active Directory Administrative Center is built on Windows PowerShell, it can expose the Windows PowerShell commands that are used to interact with the GUI. These commands can be used to learn Windows PowerShell, build Active Directory management scripts, and keep track of changes that are made within the GUI.

## Lesson 3

## Managing Servers by Using Windows PowerShell 3.0

As you become familiar with Windows PowerShell, you can perform administrative and management tasks with more ease. There are advanced features in Windows PowerShell 3.0 which let you manage a single server from a local console and to manage many servers from a remote location. The advanced features include Windows PowerShell Web Access, Windows PowerShell jobs, and Windows PowerShell workflow.

This lesson introduces some more advanced features of Windows PowerShell 3.0 and discusses how you might use the features to manage servers in your environment.

### Lesson Objectives

After completing this lesson, students will be able to:

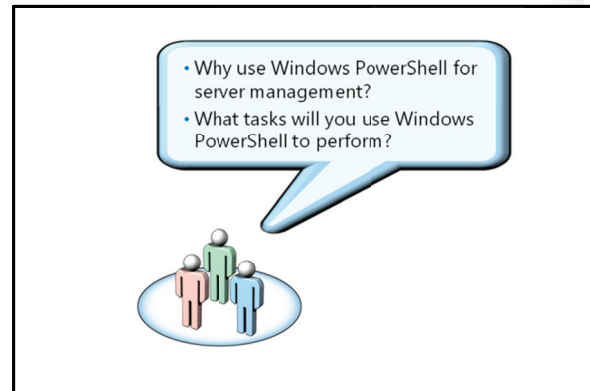
- Describe the need to use Windows PowerShell for managing servers.
- Describe how to configure and use Windows PowerShell Web Access.
- Describe Windows PowerShell jobs.
- Describe Windows PowerShell workflows and how they can be used.
- Manage a server by using Windows PowerShell 3.0.

### Discussion: The Need for Windows PowerShell for Server Management

Windows PowerShell has many features that make it useful in both large and small environments. Frequently the most difficult part of using Windows PowerShell is the starting point. Using Windows PowerShell to perform tasks that you perform every day will help you become more comfortable and more proficient in using it. Consider the following questions:

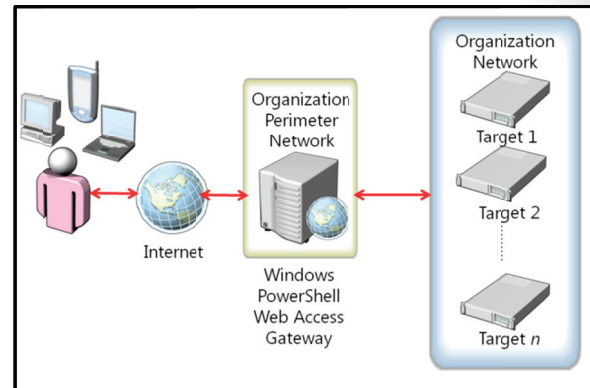
**Question:** Why use Windows PowerShell for server management?

**Question:** What tasks will you use Windows PowerShell to perform?



## What Is Windows PowerShell Web Access?

Windows PowerShell Web Access is a new feature in Windows Server 2012 that provides a web-based gateway to Windows PowerShell. This enables authorized users to administer a server without having management tools directly installed on their client computer, or having to use Remote Desktop to connect to the server. The administrator only has to configure a Windows PowerShell Web Access gateway, and use a web browser to connect.



Windows PowerShell Web Access gateway requires the Web Server Internet Information Services (IIS) role, and the .NET Framework 4.5 and Windows PowerShell 3.0 to be installed. Many client types are supported to access Windows PowerShell Web Access and still others are tested to work successfully. In order to work, the web browser must allow cookies, support connecting to the gateway by using Secure Sockets Layer (SSL), and also support JavaScript.

## Installing Windows PowerShell Web Access Gateway

To install Windows PowerShell Web Access gateway:

1. Install Windows PowerShell Web Access role.
2. Install a SSL certificate. An SSL certificate is required. A self-signed certificate can be created as part of the configuration process, however a trusted third-party certificate is recommended.
3. Create or configure an IIS site with the Windows PowerShell Web Access Gateway web application. This can be configured by using Internet Information Services Manager or by using the **Install-PswaWebApplication** cmdlet.
4. Configure Windows PowerShell Web Access authorization rules. By default, no one will be able to use Windows PowerShell Web Access until at least one authorization rule is created. An authorization rule defines which users and groups have access to specific cmdlets and which computers they can access from the gateway. Authorization rules are added by using the **Add-PswaAuthorizationRule** cmdlet. You can validate the functionality of the rules by using the **Test-PswaAuthorizationRule** cmdlet. Authorization rules are, by default, store in `%windir%\Web\PowerShellWebAccess\data\AuthorizationRules.xml`.
5. Configure destination computer authentication and authorization rules. You must configure the destination computer security settings to enable remote access from the gateway. As you assign administrative permission to the target computers, we recommend assigning only the minimally required permissions and setting the appropriate execution policy for your environment.
6. Configure additional security options. As in any environment, appropriate security best practices should be followed. One example is as installing and monitoring antivirus and anti-malware products on all the servers. Additionally, password expiration, lockout, and complexity policies should also be implemented.

## Using Windows PowerShell Web Access

To use Windows PowerShell Web Access, open a web browser and connect to the server by using `https://ServerName/pswa`. The logon page lets you connect directly to the gateway, to another server on the organization network, or to a custom URI. Using the optional connection settings on the logon page can specify one user account to log on to the gateway and specify another account to connect to the

server on the organization network. This is useful if the account authorized to connect to the gateway does not have permissions on the internal server.

After you have established a Windows PowerShell session by using Windows PowerShell Web Access, you can begin using Windows PowerShell cmdlets and executing scripts based on the execution policy settings. Although most of the functionality is the same as using Windows PowerShell remoting, there are some differences. For example, you cannot use some shortcut keys to interact with Windows PowerShell Web Access such as Ctrl+C to copy data, or any of the function keys used for things such as command history.



**Additional Reading:** Deploy Windows PowerShell Web Access  
<http://technet.microsoft.com/en-us/library/hh831611.aspx>

## What Are Windows PowerShell Jobs?

A Windows PowerShell background job runs a command or set of commands without interacting with the current Windows PowerShell session. You can start a background job by using the **Start-Job** cmdlet and then you can continue to work in the session. Using jobs can be useful when you perform tasks that can take an extended time to complete. You can also use jobs to perform the same task on several computers. The following example shows creating a new job on the local computer:

### Background Jobs:

- Enable extended tasks to be performed in the background
- Perform tasks on a number of remote servers

### Scheduled Jobs:

- Registered background jobs that can run on a schedule
- Triggers are created to define schedule

```
Start-Job -ScriptBlock {Get-ADUser -Filter *}
```

You can see the status of the job by using the **Get-Job** cmdlet and use the **Wait-Job** to be notified when the job is complete. If you have to remove a job that has not executed, you can do so with the **Remove-Job** cmdlet. These jobs are run in the background so they do not return results to your Windows PowerShell session. If you output data to the console in a background job, you can return those results by using the **Receive-Job** cmdlet.

Windows PowerShell 3.0 introduced an improvement to background jobs, which are known as scheduled jobs. These jobs can be triggered to start automatically or performed on a recurring schedule. When a scheduled job is created it is stored on disk and then registered in Task Scheduler. When a scheduled job is run, it creates an instance of the job that can then be managed by using the common job management cmdlets. The only difference between scheduled jobs and background jobs is that scheduled jobs save their results on disk.

Scheduled jobs are created by using the **Register-ScheduledJob** cmdlet. You can specify the *ScriptBlock* parameter to run a Windows PowerShell command, or you can specify a script by using the *FilePath* parameter. The following example shows how to register a scheduled job to run the **Get-LatestLogon.ps1** script.

```
Register-ScheduledJob -Name LastLogonJob -FilePath \\LON-SVR1\Scripts\Mod3\democode\Get-LastLogon.ps1
```



To enable the scheduled job to run, a schedule or trigger must be defined. Triggers are created by using the **New-JobTrigger** cmdlet. Using this cmdlet, you can use the **Add-JobTrigger** cmdlet to add the trigger to an already registered scheduled job or use it to assign a trigger when a new scheduled job is registered. Triggers can be scheduled once, daily, weekly, at server startup, when you log on. The following example shows creating a trigger that runs every Monday and Friday at 9:00 am and then registers the new scheduled job together with the trigger:

```
$Trigger = New-JobTrigger -Weekly -DaysOfWeek Monday, Friday -At 9:00AM
Register-ScheduledJob -Name ScheduledLastLogonJob -FilePath `
\\LON-SVR1\Scripts\Mod3\democode\Get-LastLogon.ps1 -Trigger $Trigger
```

You can also use the **Add-JobTrigger** cmdlet to modify an existing scheduled job as shown in the following example:

```
Add-JobTrigger -Name LastLogonJob -Trigger `
(New-JobTrigger -Daily -At 9:00AM)
```

Scheduled jobs can be used to automatically run task for: creating reports, verifying configuration settings, performing user and group maintenance, and many others.

## Introduction to Windows PowerShell Workflow

Windows PowerShell Workflow is a new feature in Windows PowerShell 3.0. It enables easy to use workflows, or task sequences within the familiar Windows PowerShell interface. A workflow can include individual Windows PowerShell commands or complete scripts. The difference between a workflow and perhaps an intricately designed script is that a workflow is designed to also be stopped, paused, and resumed. The workflow can wait until steps successfully complete to continue to the next workflow step. For example, you can create a workflow that makes changes to a multiple computers and waits for them all to restart before continuing to the next configuration step in the workflow.

- Windows PowerShell workflow:
  - Enables automating long-running and complex activities
  - Enables automating multiple server management and application provisioning
  - Enables processes to be resumed, paused, and restarted
  - Are created by using Windows PowerShell or Visual Studio Workflow Designer
- Windows Server 2012 includes over 60 predefined workflows

Windows PowerShell workflows can be created by using a Windows PowerShell console, the Windows PowerShell ISE, or by using Microsoft Visual Studio® Workflow Designer. Workflows created in Visual Studio Workflow Designer are saved as with a XAML file name extension. These workflows are imported by using the **Import-Module** cmdlet.

Workflows are run as Windows PowerShell jobs. Therefore, you can use the same cmdlets to manage running workflows as you do jobs. A workflow is created by using the following syntax:

```
Workflow WorkflowName { Commands to execute as part of the workflow }
```

After a workflow is created, it is executed as a cmdlet is executed. Each workflow can be executed with the parameters that are listed in the following table.

Parameter	Description
<i>-PSComputerName</i>	A list of target computers for the workflow to execute on
<i>-PSRunningTimeoutSec</i>	Length of time to allow for the workflow to execute
<i>-PSConnectionRetryCount</i>	Enable the workflow to retry connections several times
<i>-PSPersist</i>	Toggles the workflow to checkpoint data and state after each activity

In a workflow, commands can be performed in a parallel or sequential manner. Commands that can be run in parallel are identified by using the parallel keyword. Commands that must be performed sequentially are identified by using the sequence keyword. The following example shows a workflow with both keywords being used:

```
Workflow Get-DomainServerStats
{
# The following are executed in any order
Parallel
{
    Get-Process
    Get-ADUser -Filter *
# The following are executed sequentially
Sequence
{
    Set-AdUser Administrator -Description "Updated content"
    Get-ADUser Administrator -Properties Description
}
}
}
```

Windows has number of built in workflows to enable configuration of multi-server deployments of Remote Desktop Services, retrieve information about installed Windows roles, and restarting servers. To view defined workflows use the following command:

```
Get-Command -CommandCapability workflow
```

## Demonstration: Managing a Server by Using Windows PowerShell 3.0

In this demonstration, you will review how to use Windows PowerShell Web Access and Windows PowerShell jobs.

### Demonstration Steps

1. Start virtual machines LON-DC1, LON-SVR1, and LON-SVR2, and then log on to LON-DC1 as the domain administrator.
2. Open Windows PowerShell Web Access at **http://LON-DC1/pswa** by using the following information:
  - o User name: **Administrator**
  - o Password: **Pa\$\$w0rd**
  - o Computer: **LON-DC1**

3. Start a new job to list all Active Directory users, by using the **Start-Job** cmdlet.
4. Obtain the status of the job by running **Get-Job**.
5. Create a new scheduled job by running the following commands each followed by Enter:

```
$Trigger = New-JobTrigger -Weekly -DaysOfWeek Monday, Friday -At 9:00AM  
Register-ScheduledJob -Name ScheduledJob1 -ScriptBlock {Get-ADUser -Filter * } -  
Trigger $Trigger
```

6. Run the scheduled job immediately by using the **Start-Job** cmdlet.

## Lab: Managing Servers Running Windows Server 2012 by Using Windows PowerShell 3.0

### Scenario

As the A. Datum network grows in size and complexity, it is becoming increasingly apparent that some IT management processes have to be streamlined. The number of users in the organization is increasing quickly with users distributed in many locations. Servers are also being deployed in multiple data centers and in private and public clouds. A. Datum is deploying most new servers as virtual servers in Hyper-V. A. Datum has to ensure that both the host computers and virtual machines are managed consistently.

To address these server and AD DS management issues, you have to gain familiarity with Windows PowerShell. You have to understand how to run simple and complex commands and how to create scripts that will automate many of the regular management tasks.

### Objectives

After completing this lab, you will be able to:

- Explore Windows PowerShell commands and tools.
- Manage AD DS by using Windows PowerShell.
- Manage local and remote servers by using Windows PowerShell.

### Lab Setup

Estimated time: **30-60 minutes**

Virtual Machine(s)	20417-LON-DC1 20417-LON-SVR1 20417-LON-SVR2
User Name	Adatum\Administrator
Password	Pa\$\$w0rd

### Lab Setup

For this lab, you will use the available virtual machine environment. Before you begin the lab, you must complete the following steps:

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, click **20417A-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on using the following credentials:
  - a. User name: **Adatum\Administrator**
  - b. Password: **Pa\$\$w0rd**
5. Repeat steps 2-4 for **20417A-LON-SVR1** and **20417A-LON-SVR2**.

## Exercise 1: Introduction to Windows PowerShell 3.0

### Scenario

As a part of becoming familiar with the Windows PowerShell interface, you will explore interface and browse through available cmdlets.

The main tasks for this exercise are as follows:

1. Use Windows PowerShell ISE to retrieve basic information about LON-DC1.
2. Use Windows PowerShell ISE to retrieve a list of stopped services on LON-DC1.
3. Use a Remote Windows PowerShell session to install XPS Viewer on LON-SVR1.

#### ► Task 1: Use Windows PowerShell ISE to retrieve basic information about LON-DC1

1. Start the following virtual machines: LON-DC1, LON-SVR1, and LON-SVR2.
1. On LON-DC1, open Windows PowerShell ISE as an administrator.
2. Retrieve a list of installed Windows features by using **Get-WindowsFeature**.
3. List the contents of the E:\ModX\Democode directory by running **Get-ChildItem E:\ModXA\Democode**.
4. List the contents of C:\Windows, by running **dir C:\Windows**.
5. Use tab completion to find the correct cmdlet that begins with **Get-Ex** to see the execution policy setting on LON-DC1.

#### ► Task 2: Use Windows PowerShell ISE to retrieve a list of stopped services on LON-DC1

1. If it is necessary, open Windows PowerShell ISE as an administrator.
2. Retrieve a list of services by running **Get-Service**.
3. Assign the results of **Get-Service** to the \$Services variable.
4. Use the **Get-Help** cmdlet to view the examples of how to use **Where-Object**.
5. Use a pipeline to pipe the \$Services variable to the **Where-Object** cmdlet to show only services that have a status of stopped.

#### ► Task 3: Use a Remote Windows PowerShell session to install XPS Viewer on LON-SVR1

1. If it is necessary, open Windows PowerShell ISE as an administrator and open a new remote PowerShell tab.
2. Establish a Remote PowerShell session with LON-SVR1.
3. Retrieve a list of all installed Windows Features on LON-SVR1 by using **Get-WindowsFeature**.
4. Install XPS Viewer on LON-SVR by using **Add-WindowsFeature**.
5. Use command history to run **Get-WindowsFeature** and verify that XPS Viewer is installed.
6. Close the Remote PowerShell session.

**Results:** After this exercise, you will have explored the Windows PowerShell ISE interface and used cmdlets, variables, and pipelining.

## Exercise 2: Managing AD DS by Using Windows PowerShell 3.0

### Scenario

After you explore Windows PowerShell interface and cmdlets, you want to explore options and available cmdlets in the Active Directory module for Windows PowerShell and begin to use it for basic tasks such as formatting Windows PowerShell output, using variables and loops, and creating scripts.

The main tasks for this exercise are as follows:

1. Import the Active Directory PowerShell module and view the available cmdlets.
2. View options on how to create a report of users in the Active Directory domain.
3. Use a script to create new users in the domain by using a CSV-based file.
4. Create a script to modify the address of a user based on the day of the week.

#### ► Task 1: Import the Active Directory PowerShell module and view the available cmdlets

1. If it is necessary, open Windows PowerShell ISE as an administrator.
2. Import the Active Directory module by using the **Import-Module** cmdlet.
3. Use the **Get-Command** cmdlet to view the cmdlets available in the Active Directory module.

#### ► Task 2: View options on how to create a report of users in the Active Directory domain

1. If it is necessary, open Windows PowerShell ISE as an administrator and import the Active Directory module.
2. Use the **Get-Command** cmdlet to view the cmdlets available in the ActiveDirectory module.
3. Use Windows PowerShell to view a list of all Users in the domain. Review how **Format-List** modifies formatting by running the following commands by using:

```
Get-ADUser -Filter * | Format-List
Get-ADUser -Filter * |
Format-List -Property GivenName, Surname
Get-ADUser -Filter * -Properties * | Format-List *
```

4. Use Windows PowerShell to view a list of all Users in the domain. Review how **Format-Table** modifies the formatting by running the following commands by using:

```
Get-ADUser -Filter * | Format-Table
Get-ADUser -Filter * |
Format-Table -Property GivenName, Surname
Get-ADUser -Filter * -Properties * | Format-Table
```

5. Use Windows PowerShell to view a list of all OUs in the domain. Review how **Format-Wide** modifies the formatting by running the following commands:

```
Get-ADOrganizationalUnit -Filter * | Format-Wide
Get- ADOrganizationalUnit -Filter * |
Format-Wide -column 3
```

6. Use Windows PowerShell to adjust the formatting of the users report. Review how the **Sort-Object** cmdlet modified the output, by running the following:

```
Get-ADUser -Filter * | Sort-Object | Format-Wide
Get-ADUser -Filter * | Sort-Object -Property ObjectGUID | Format-Wide -Property
ObjectGUID
```

7. Run the following commands to see how to use the **Measure-Object** cmdlet:

```
Get-ADUser -Filter * | Measure-Object
```

► **Task 3: Use a script to create new users in the domain by using a CSV-based file**

1. On LON-DC1, browse to the Start screen and then type **Notepad.exe**. Press Enter.
2. Use Notepad.exe to view **E:\ModXA\Democode\LabUsers.csv**. You will need to change the file type to all files.
3. Use Windows PowerShell ISE to open the script that is located at **E:\ModXA\Democode\LabUsers.ps1**
4. On line 13, modify the **\$OU** variable to read: **\$OU = "ou=sales, dc=adatum, dc=com"**
5. Run the **LabUsers.ps1** script.
6. Use **Get-ADUser -Filter \* -SearchBase "OU=Sales, DC=Adatum, DC=com"** to confirm Luka Abrus, Marcel Truempy, Andy Brauning, and Cynthia Cary were created.

► **Task 4: Create a script to modify the address of a user based on the day of the week**

1. If it is necessary, open Windows PowerShell ISE as an administrator and import the Active Directory module.
2. Use Windows Powershell ISE to open the script that is located at **E:\ModXA\Democode\Using If Statements.ps1**
3. Verify that line 9 reads:  
**\$Admin = Get-ADUser -identity Administrator -Properties StreetAddress**
4. Review each section of the script and then run the script. Run the script a second time to view the changes.

**Results:** After completing this lab, you will have explored the Active Directory Windows PowerShell module, experienced formatting output in Windows PowerShell, used a Windows PowerShell script to create users, and used Windows PowerShell conditional loops to modify Active Directory properties.

## Exercise 3: Managing Servers by Using Windows PowerShell 3.0

### Scenario

Because of plans for remote server management, you want to explore possibilities to use Windows PowerShell for remote management. You want to test remote connections in Windows PowerShell and Windows PowerShell Web Access.

The main tasks for this exercise are as follows:

1. 1. Install and configure Windows PowerShell Web Access.
2. 2. Verify Windows PowerShell Web Access configuration.

### ► Task 1: Install and configure Windows PowerShell Web Access

1. Install Windows PowerShell Web Access on LON-DC1 by using the following command:

```
Install-WindowsFeature -Name WindowsPowerShellWebAccess -ComputerName LON-DC1 -
IncludeManagementTools -Restart
```

2. Configure Windows PowerShell Web Access by running **Install-PswaWebApplication – UseTestCertificate**.
3. Create a Windows PowerShell Web Access Authorization Rule that only enables the administrator to access the gateway by using the **Add-PSWaAuthorizationRule**.

### ► Task 2: Verify Windows PowerShell Web Access configuration

1. Open Internet Explorer and navigate to https://LON-DC1/pswa.
2. Sign in to Windows PowerShell Web Access by using the following information:
  - User: **Administrator**
  - Password: **Pa\$\$w0rd**
  - Computer: **LON-DC1**
3. Verify that you can retrieve information from LON-SVR1 by retrieving the five newest System events. Run the following command:

```
Get-EventLog System -Newest 5
```

4. Obtain the same information from LON-SVR2 and LON-DC1 by running the following command:

```
Invoke-Command -ScriptBlock { Get-Eventlog Security -Newest 20 } -ComputerName LON-
DC1,LON-SVR2
```

**Results:** After this exercise, you will have performed one to many management of remote servers by using Windows PowerShell, installed and configured Windows PowerShell Web Access, and managed servers by using Windows PowerShell Web Access.

### ► To prepare for the next module

When you are finished the lab, revert the virtual machines to their initial state. To do this, perform the following steps:

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20417A-LON-SVR1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20417A-LON-SVR2** and **20417A-LON-DC1**.



## Module Review and Takeaways

### Review Questions

**Question:** Which cmdlet will display the content of a text file?

**Question:** Which cmdlet will move a file to another directory?

**Question:** Which cmdlet will rename a file?

**Question:** Which cmdlet will create a new directory?

**Question:** Which cmdlet do you think would retrieve information from the event log?

**Question:** Which cmdlet do you think would start a stopped VM?

### Best Practices

- Make a goal to spend time learning how to use Windows PowerShell for your common tasks. This will make you more comfortable with working with Windows PowerShell and will equip you for using it to resolve more difficult problems.
- Save the commands that you have used to resolve problems in a script file for later reference.
- Use Windows PowerShell ISE to help write scripts and ensure you have the correct syntax.

### Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Administrators cannot find the correct Windows PowerShell cmdlet for a task.	
Administrator cannot connect to a server by using remote Windows PowerShell.	
<b>Get-Help</b> does not provide any help for cmdlets.	
An administrator is new to Windows PowerShell and is uncomfortable with the command-line.	

## Tools

You can use the tools in the following table to work with Windows PowerShell.

Tool	Description
Windows PowerShell Integrated Script Editor (ISE)	Windows PowerShell ISE provides a simple, yet powerful interface to create and test scripts, and discover new cmdlets.
Microsoft Visual Studio Workflow Designer	This is a development tool that is used to create Windows PowerShell workflows.
Powershell.exe	This is the Windows PowerShell executable.
Active Directory Administrative Center	This tool enables you to perform common Active Directory management tasks such as creating and modifying user and computer accounts. All the changes that you made by using this management tool are logged in the Windows PowerShell History pane.

## Real-world Issues and Scenarios

Many common tools can be replaced with Windows PowerShell cmdlets. The following table gives some examples of common commands that can be replaced with Windows PowerShell cmdlets in Windows Server 2012.

Old Command	Windows PowerShell Equivalent
<b>ipconfig /a</b>	<b>Get-NetIPConfiguration</b>
<b>Shutdown.exe</b>	<b>Restart-Computer</b>
<b>Net Start</b>	<b>Start-Service (Restart-Service)</b>
<b>Net Stop</b>	<b>Stop-Service (Restart-Service)</b>
<b>Net Use</b>	<b>New-SmbMapping</b>
<b>Netstat</b>	<b>Get-NetTCPConnection</b>
<b>Netsh advfirewall add</b>	<b>New-NetFirewallRule</b>
<b>Route Print</b>	<b>Get-NetRoute</b>

# Module 4

## Managing Storage for Windows Server 2012

### Contents:

Module Overview	4-1
Lesson 1: New Features in Windows Server 2012 Storage	4-2
Lesson 2: Configuring iSCSI Storage	4-12
Lesson 3: Configuring Storage Spaces in Windows Server 2012	4-18
Lab A: Managing Storage for Servers Based on Windows Server 2012	4-23
Lesson 4: Configuring BranchCache in Windows Server 2012	4-25
Lab B: Implementing BranchCache	4-36
Module Review and Takeaways	4-40

## Module Overview

Storage space requirements have been increasing ever since the invention of server-based file shares. The Windows Server® 2012 and Windows® 8 operating systems include two new features to reduce the disk space that is required and to effectively manage physical disks: data deduplication and storage spaces. This module provides an overview of these features and explains the steps required to configure them.

Another concern in storage is the connection between the storage and the remote disks. Internet small computer system interface (iSCSI) storage in Windows Server 2012 is a cost-effective feature that helps create a connection between the servers and the storage. To implement iSCSI storage in Windows Server 2012, you must be familiar with the iSCSI architecture and components. In addition, you must be familiar with the tools that are provided in Windows Server to implement an iSCSI-based storage. Also, in organizations that have branch offices, you have to consider slow links and how to use these links efficiently when data is sent between your offices. The BranchCache feature in Windows Server 2012 helps address the problem of slow connectivity. This module explains the BranchCache feature and the steps to configure BranchCache.

### Objectives

After completing this module, you will be able to:

- Describe the new features in Windows Server 2012 storage.
- Configure iSCSI storage.
- Configure storage spaces.
- Configure BranchCache.

## Lesson 1

# New Features in Windows Server 2012 Storage

The storage demand on servers is ever-increasing, and storage comprises a larger part of an IT department's budget. Larger volumes are required on flexible disks that can be added or removed dynamically. Windows Server 2012 includes changes to the storage area that will help administrators to ease the management of physical disks and provide technologies to reduce disk space consumption.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe the File and Storage Services in Windows Server 2012.
- Describe the data deduplication process.
- Configure data deduplication.
- Describe the capabilities of thin provisioning and trim storage.
- Describe the new features in File Server Resource Manager.
- Describe basic and dynamic disks.
- Describe Resilient File System (ReFS) and its advantages.
- Describe removed and deprecated features.

## File and Storage Services in Windows Server 2012

File and Storage Services includes technologies that help you set up and manage one or more file servers. File servers are servers that act as central locations on the network where you can store files and optionally, share them with users.

Windows Server 2012 offers the following new file and storage services features:

- *Multiterabyte volumes.* You can use this feature to deploy multiterabyte NTFS file system volumes, which support consolidation scenarios and maximizes storage use. The Chkdsk tool introduces a new approach that prioritizes volume availability and allows for the detection of corruption while the volume remains online with data available.
- *Data deduplication.* You can use this feature to save disk space by storing a single copy of identical data on the volume.
- *iSCSI target server.* You can use this feature to block storage to other servers and applications on the network by using the iSCSI standard.
- *Storage spaces and storage pools.* You can use this feature to virtualize storage by grouping industry-standard disks into storage pools, and then create storage spaces from the available capacity in the storage pools.

Windows Server 2012 provides the following new file and storage features:

- Multiterabyte volumes
- Data deduplication
- iSCSI target server
- Storage spaces and storage pools
- Unified remote management of File and Storage Services in Server Manager
- Windows PowerShell cmdlets for File and Storage Services

- *Unified remote management of File and Storage Services in Server Manager.* You can use this feature to remotely manage multiple file servers, including their role services and storage, all from a single window.
- *Windows PowerShell® cmdlets for File and Storage Services.* You can use the Windows PowerShell cmdlets for performing most administration tasks for file and storage servers.



**Additional Reading:** File and Storage Services overview  
[http://technet.microsoft.com/en-us/library/hh831487\(d=lightweight,v=ws.11\)](http://technet.microsoft.com/en-us/library/hh831487(d=lightweight,v=ws.11))

**Question:** Are you currently implementing volumes that are 10 terabytes or larger? What are the problems with volumes of that size?

## What Is Data Deduplication?

Data deduplication is a role service of Windows Server 2012. Data deduplication identifies and removes duplications within data without compromising its integrity to achieve the ultimate goal of storing more data while concurrently using less physical disk space.

Data integrity and recoverability are maintained in a process that involves evaluating checksum results and other algorithms. Data deduplication is highly scalable, resource efficient, and nonintrusive. It can run on dozens of large volumes of primary data concurrently without affecting other workloads on the server. Low impact on the server workloads is maintained by throttling the CPU and memory resources that are consumed. Using data deduplication jobs, you can schedule when data deduplication should run, specify the resources to deduplicate, and tune file selection.

When combined with BranchCache, the same optimization techniques are applied to data that is transferred over the wide area network (WAN) to a branch office. This results in faster file download times and reduced bandwidth consumption.

### Volume Requirements for Data Deduplication

After the feature is installed, you can enable data deduplication on a per volume basis. Each volume must meet the following requirements:

- Volumes must not be a system or boot volume. Deduplication is not supported on volumes where the operating system is installed.
- Volumes may be partitioned by using master boot record (MBR) or GUID partition table (GPT) format, and must be formatted by using the NTFS file system. The new Resilient File System (ReFS) file system is not supported for use on a data deduplication volume.
- Volumes must be exposed to Windows as non-removable drives, that is, no USB or floppy drives.
- Volumes can be on shared storage, such as a Fibre Channel or Serial Attached SCSI (SAS) array, or an iSCSI storage area network (SAN).
- Cluster Shared Volumes (CSV) volumes are not supported.

- Data deduplication identifies and removes duplications within data without compromising its integrity or fidelity with the ultimate goal to store more data on less space
- When you enable data deduplication on a volume, a background task runs with low-priority that:
  1. Segments data into small, variable sized chunks
  2. Identifies duplicate chunks
  3. Replaces redundant copies with a reference
  4. Compresses chunks
- You should consider using deduplication for the following areas:

File Shares

Software Deployment Shares

VHD Libraries

## The Data Deduplication Process

When you enable data deduplication on a volume, a background task runs with low-priority that processes the files on the volume. That is, the background task segments all file data on the volume into small, variable sized chunks (32 to 128 KB). Then, it identifies chunks that have one or more duplicates on the volume. All duplicate chunks are then replaced (erased from disk) with a reference to a single copy of that chunk. Finally, all remaining chunks are compressed so that even more disk space is saved.

## When to Use Data Deduplication

Data deduplication is designed to be installed on primary (and not logically extended) data volumes without adding any additional dedicated hardware. You can install and use the feature without affecting the primary workload on the server. The default settings are non-intrusive because only files older than 30 days are processed. The implementation is designed for low memory and CPU priority. However, if memory use becomes high, deduplication backs off and waits for available resources. You can schedule deduplication based on the type of data involved and the frequency and volume of changes that occur to the volume or particular file types.

You should consider using deduplication for the following areas:

- *File shares.* This includes group content publication or sharing, user home folders, and profile redirection (offline files). You may be able to save approximately 30–50 percent disk space.
- *Software deployment shares.* This includes software binaries, images, and updates. You may be able to save approximately 70–80 percent space.
- *Virtual hard disk (VHD) libraries.* This includes VHD file storage for provisioning to hypervisors. You may be able to save approximately 80–95 percent space.



**Note:** Use the deduplication evaluation tool (DDPEval.exe) to analyze a volume about expected savings that you would get when enabling deduplication. This utility is automatically installed to \\Windows\\System32\\ of the local computer when data deduplication is enabled.

When data deduplication is enabled, and the data is optimized, the volume contains the following:

- *Unoptimized files.* These are skipped files. For example, system state files, encrypted files, files with extended attributes, files smaller than 32KB, and reparse point files—previously optimized files that contain pointers to the respective chunks in the chunk store needed to build the file.
- *Optimized files.* These are stored as reference points to the chunk store.
- *Chunk store.* This is the optimized file data.



### Additional Reading:

Data Deduplication Overview

<http://technet.microsoft.com/en-us/library/hh831602>

Introduction to Data Deduplication in Windows Server 2012

<http://blogs.technet.com/b/filecab/archive/2012/05/21/introduction-to-data-deduplication-in-windows-server-2012.aspx>

**Question:** On which of your shares can you use data deduplication?

## Demonstration: Configuring Data Deduplication

In this demonstration, you will see how to add the data deduplication role service and enable data deduplication on drive E.

### Demonstration Steps

#### Add the Data Deduplication role service

1. Log on to LON-DC1 with a username of **Adatum\Administrator** and the password of **Pa\$\$w0rd**.
2. In Server Manager, start the **Add Roles and Features Wizard**, install the following roles and features to the local server and accept the default values:
  - **File And Storage Services (Installed)\File and iSCSI Services\Data Deduplication**

#### Enable Data Deduplication on E: Drive

1. On LON-DC1, in Server Manager, in the navigation pane, click **File and Storage Services**, and then click **Volumes**.
2. In the Volumes pane, right-click **E:**, and select **Configure Data Deduplication**.
3. Configure data deduplication with the following settings:
  - Enable data deduplication: **Enabled**
  - Deduplicate files older than (in days): **3**
  - Set Deduplication Schedule: **Enable throughput optimization**
  - Start time: **current time**

## What Are Thin Provisioning and Trim Storage?

Windows Server 2012 introduces two new storage concepts. They are:

- *Thin provisioning.* This is a functionality that you can use to allocate storage space on a just-in-time basis and is available with storage spaces or virtual disks. Using traditional disk provisioning methods, a volume would immediately consume all the disk space it was sized for. For example, a 2 GB volume would occupy 2 GB of disk space. Even if the data inside that volume is less than 2 GB, that entire storage amount is reserved on the disk.

Similar to a dynamically expanding VHD, a virtual disk configured as thin provisioning would only use the space from a storage pool on as-needed basis. The virtual disk is only allocated space on the volume as data is added. This also lets you create virtual disks that have a larger maximum size than the free space in the storage pool. For example, with thin provisioning, you can create a 1 terabyte virtual disk even though your storage pool only has 500 GB of free space available.

- Thin provisioning is the ability to allocate storage space just-in time
- Trim storage is the ability to reclaim storage that is no longer needed
- Thin provisioning and trim storage are available by default in Windows Server 2012; no feature or role needs to be installed
- Thin provisioning and trim storage include the following capabilities:
  - Identification
  - Notification
  - Optimization

- **Trim storage.** This is a functionality that you can use to reclaim storage that is no longer needed. The file system can inform an underlying physical storage device that the contents of specified sectors are no longer important. Therefore, these sectors can be used by another volume in a storage pool. Trim requests to a mounted VHD or inside Hyper-V® are now propagated to the underlying storage device.

Thin provisioning and trim storage are available by default in Windows Server 2012; no feature or role has to be installed.

Thin provisioning and trim storage in Windows Server 2012 provides the following capabilities:

- **Identification.** Windows Server 2012 uses a standardized method to detect and identify thinly-provisioned virtual disks, thereby enabling additional capabilities delivered by the storage stack. The storage stack is provided in the operating system and is available through storage management applications.
- **Notification.** When the configured physical storage use thresholds are reached, Windows Server 2012 notifies the administrator through events. This enables the administrator to take appropriate action as soon as possible. These events can also start automated actions from sophisticated management applications, such as Microsoft System Center.
- **Optimization.** Windows Server 2012 provides a new API that enables applications return storage when it is no longer needed. NTFS issues trim notifications in real time, when appropriate. Additionally, trim notifications are issued as part of storage consolidation (optimization), which is performed regularly on a scheduled basis.



**Additional Reading:** Thin Provisioning and Trim Storage Overview  
<http://technet.microsoft.com/en-us/library/hh831391.aspx>

## What's New in File Server Resource Manager?

You can use the File Server Resource Manager to manage and classify data that is stored on file servers. File Server Resource Manager includes the following features:

- **File classification infrastructure.** This feature automates the data classification process. You can dynamically apply access policies to files based on their classification. Example policies include Dynamic Access Control for restricting access to files, file encryption, and file expiration. You can classify files automatically by using file classification rules, or manually by modifying the properties of a selected file or folder.
- **File management tasks.** You can use this feature to apply a conditional policy or action to files, based on their classification. The conditions of a file management task include the file location, the classification properties, the date the file was created, the last modified date of the file, or the last time that the file was accessed. The actions that a file management task can take include the ability to expire files, encrypt files, or run a custom command.

You can use the File Server Resource Manager to manage and classify data that is stored on file servers

- File Server Resource Manager includes the following features:
  - File classification infrastructure
  - File management tasks
  - Quota management
  - File screening management
  - Storage reports

- The new features in File Server Resource Manager include:
  - Dynamic Access Control
  - Manual classification
  - Access-denied assistance
  - Filemanagement tasks
  - Automatic classification



- *Quota management.* You can use this feature to limit the space allowed for a volume or folder. Quotas can be automatically applied to new folders that are created on a volume. You can also define quota templates that you can apply to new volumes or folders.
- *File screening management.* You can use this feature to control the types of files that users can store on a file server. You can limit the extension that can be stored on your file shares. For example, you can create a file screen that does not enable files that have an MP3 extension to be stored in personal shared folders on a file server.
- *Storage reports.* You can use this feature to identify trends in disk usage and how your data is classified, and monitor attempts by a selected group of users to save unauthorized files.

You can configure and manage the File Server Resource Manager by using the File Server Resource Manager Microsoft Management Console (MMC) console or by using Windows PowerShell.

The following features of the File Server Resource Manager are new and are added in Windows Server 2012:

- *Dynamic Access Control.* Dynamic Access Control uses file classification infrastructure to help you centrally control and audit access to files on your file servers.
- *Manual classification.* Manual classification enables users to classify files and folders manually without the need to create automatic classification rules.
- *Access-denied assistance.* You can use access-denied assistance to customize the access denied error message that users see in Windows 8 Consumer Preview when they do not have access to a file or a folder.
- *File management tasks.* The updates to file management tasks include Active Directory® Rights Management Services (AD RMS) file management tasks, continuous file management tasks, and dynamic namespace for file management tasks.
- *Automatic classification.* The updates to automatic classification enable you to get more precise control on how data is classified on your file servers, including continuous classification, using Windows PowerShell for custom classification, updates to the existing content classifier, and dynamic namespace for classification rules.



**Additional Reading:** What's new in File Server Resource Manager  
<http://technet.microsoft.com/en-us/library/hh831746.aspx>

**Question:** Are you currently using the File Server Resource Manager in Windows Server 2008 R2? If yes, what areas do you use it for?

## What Are Basic and Dynamic Disks?

Windows Server 2012 continues to support basic disks and dynamic disks.

### Basic Disk

Basic storage uses typical partition tables supported by MS-DOS, and all versions of the Windows operating system. A disk initialized for basic storage is called a basic disk. A basic disk contains basic partitions, such as primary partitions and an extended partition. An extended partition can be subdivided into logical drives.

By default, when you initialize a disk in Windows, the disk is configured as a basic disk. Basic disks can easily be converted to dynamic disks without any loss of data. However, when you convert a dynamic disk to basic disk, all data on the disk will be lost.

Some applications such as the storage spaces feature in Windows Server 2012 cannot use dynamic disks. In addition, there is no performance gain by converting basic disks to dynamic disks. For these reasons, most administrators do not convert basic disks to dynamic disks unless they have to use some additional volume configuration options available with dynamic disks.

### Dynamic Disk

Dynamic storage is supported in all Windows operating systems including the Windows XP operating systems and the Microsoft® Windows NT Server 4.0 operating system. A disk initialized for dynamic storage is called a dynamic disk. A dynamic disk contains dynamic volumes. With dynamic storage, you can perform disk and volume management without the need to restart Windows.

When you configure dynamic disks, you create volumes instead of partitions. A volume is a storage unit made from free space on one or more disks. It can be formatted with a file system and can be assigned a drive letter or configured with a mount point.

The dynamic volumes include:

- *Simple volumes.* A simple volume uses free space from a single disk. It can be a single region on a disk or consist of multiple, concatenated regions. A simple volume can be extended within the same disk or onto additional disks. If a simple volume is extended across multiple disks, it becomes a spanned volume.
- *Spanned volumes.* A spanned volume is created from free disk space that is linked from multiple disks. You can extend a spanned volume onto a maximum of 32 disks. A spanned volume cannot be mirrored and is not fault-tolerant. Therefore if you lose one disk, you lose all the spanned volume.
- *Striped volumes.* A striped volume is a volume whose data is spread across two or more physical disks. The data on this type of volume is allocated alternately and evenly to each of the physical disks. A striped volume cannot be mirrored or extended and is not fault-tolerant, again meaning the loss of one disk will cause the loss of data immediately. Striping is also known as redundant array of independent disks (RAID)-0.
- *Mirrored volumes.* A mirrored volume is a fault-tolerant volume whose data is duplicated on two physical disks. All the data on one volume is copied to another disk to provide data redundancy. If one of the disks fails, the data can still be accessed from the remaining disk. A mirrored volume cannot be extended. Mirroring is also known as RAID-1.

- **Basic disks:**
  - Are disks initialized for basic storage
  - Are the default storage for Windows
- **Dynamic disks:**
  - Can be modified without restarting Windows
  - Provide several options for configuring volumes
- **Disk volume requirements include:**
  - System volume for hardware specific files required to start the server
  - Boot volume for the operating system files

- **RAID-5 volumes.** A RAID-5 volume is a fault-tolerant volume whose data is striped across a minimum of three or more disks. Parity (a calculated value that can be used to reconstruct data after a failure) is also striped across the disk array. If a physical disk fails, the portion of the RAID-5 volume that was on that failed disk can be re-created from the remaining data and the parity. A RAID-5 volume cannot be mirrored or extended.

### Required Disk Volumes

Regardless of which type of disk that you use, you must configure a system volume and a boot volume on one of the hard disks in the server:

- **System volumes.** The system volume contains the hardware-specific files that are needed to load Windows (for example, Bootmgr, BOOTSECT.bak, and BCD). The system volume can be, but does not have to be, the same as the boot volume.
- **Boot volumes.** The boot volume contains the Windows operating system files that are located in the %Systemroot% and %Systemroot%\System32 folders. The boot volume can be, but does not have to be, the same as the system volume.



**Note:** When you install the Windows 8 operating system or Windows Server 2012 in a clean installation, a separate system volume is created to enable encrypting the boot volume by using BitLocker®.



#### Additional Reading:

How Basic Disks and Volumes Work

<http://go.microsoft.com/fwlink/?LinkID=199648>

Dynamic Disks and Volumes

<http://go.microsoft.com/fwlink/?LinkID=199649>

## What Is the Resilient File System?

Resilient File System (ReFS) is a new file system provided in Windows Server 2012. ReFS is based on the NTFS file system and provides the following advantages:

- Metadata integrity with checksums
- Integrity streams providing optional user data integrity
- Allocation on write transactional model for robust disk updates (also known as copy on write)
- Large volume, file, and directory sizes
- Storage pooling and virtualization making file system creation and management easy
- Data striping for performance (bandwidth can be managed) and redundancy for fault tolerance
- Disk scrubbing for protection against latent disk errors
- Resiliency to corruptions with salvage for maximum volume availability in every case
- Shared storage pools across computers for additional failure tolerance and load balancing

Resilient File System (ReFS) is a new file system provided in Windows Server 2012

- ReFS provides the following advantages:
  - Metadata integrity with checksums
  - Integrity streams providing optional user data integrity
  - Allocation on write transactional model
  - Large volume, file, and directory sizes (2<sup>78</sup> bytes with 16-KB cluster size)
  - Storage pooling and virtualization
  - Data striping for performance and redundancy
  - Disk scrubbing for protection against latent disk errors
  - Resiliency to corruptions with salvage
  - Shared storage pools across machines

ReFS inherits the features from NTFS including BitLocker encryption, access-control lists for security, Update Sequence Number (USN) journal, change notifications, symbolic links, junction points, mount points, reparse points, volume snapshots, file IDs, and oplocks.

Because ReFS uses a subset of features from NTFS, it is designed to maintain backward compatibility with its older counterpart. Therefore, Windows 8 clients or earlier can read and write to ReFS hard-drive partitions and shares on a server, just as they can with those running NTFS. But, as implied in its name, the new file system offers more resiliency, meaning better data verification, error correction, and scalability.

Beyond its greater resiliency, ReFS also surpasses NTFS by offering larger maximum sizes for individual files, directories, disk volumes, and other items, as listed in the following table.

Attribute	Limit
Maximum size of a single file	$2^{64}-1$ bytes (18.446.744.073.709.551.616 bytes)
Maximum size of a single volume	$2^{78}$ bytes with 16KB cluster size ( $2^{64} * 16 * 2^{10}$ ) Windows stack addressing allows $2^{64}$ bytes
Maximum number of files in a directory	$2^{64}$
Maximum number of directories in a volume	$2^{64}$
Maximum file name length	32K unicode characters
Maximum path length	32K
Maximum size of any storage pool	4 petabyte
Maximum number of storage pools in a system	No limit
Maximum number of spaces in a storage pool	No limit

## Removed and Deprecated Features

The following storage-related features are removed and deprecated from Windows Server 2012:

- The Storage Manager for SANs snap-in for MMC is removed. Instead, you can manage storage with Windows PowerShell cmdlets and Server Manager.
- The Storage Explorer snap-in for MMC is removed.
- The SCSIport host-bus adapter driver is removed. Instead, you can either use a Storport driver or a different host-bus adapter.
- The File Server Resource Manager command-line tools such as dirquota.exe, filesnrm.exe, and storrep.exe are removed. This functionality is available in Windows PowerShell.
- The File Replication Service (FRS) is replaced by DFS Replication.

- The following features are removed and deprecated in Windows Server 2012:
  - Storage Manager for Storage Area Networks (SANs) snap-in
  - Storage Explorer snap-in
  - SCSIport host-bus adapter driver
  - File Server Resource Manager command-line tools
  - FRS
  - Share and Storage Management snap-in
  - Shared Folders snap-in
  - VDS provider

- The Share and Storage Management snap-in is replaced by the File and Storage Services role in Server Manager.
- The Shared Folders snap-in is replaced by the File and Storage Services role in Server Manager.
- The Virtual Disk Service (VDS) provider is replaced by the Storage Management APIs and storage provider or the Storage Management Initiative – Specification (SMI-S) standard and a compliant storage provider.

## Lesson 2

# Configuring iSCSI Storage

In this lesson, you will learn how to create a connection between servers and iSCSI storage. You will perform these tasks by using IP-based iSCSI storage. iSCSI storage is an inexpensive and simple way to configure a connection to remote disks. Many application requirements dictate that remote storage connections must be redundant in nature for fault tolerance or high availability. For this purpose, you will also learn how to create both single and redundant connections to an iSCSI target. You will do so by using the iSCSI initiator software that is available in Windows Server 2012.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe iSCSI and its components.
- Describe the iSCSI target server and the iSCSI initiator.
- Describe how to configure high-availability and locate iSCSI storage.
- Configure iSCSI target.
- Connect to the iSCSI storage.

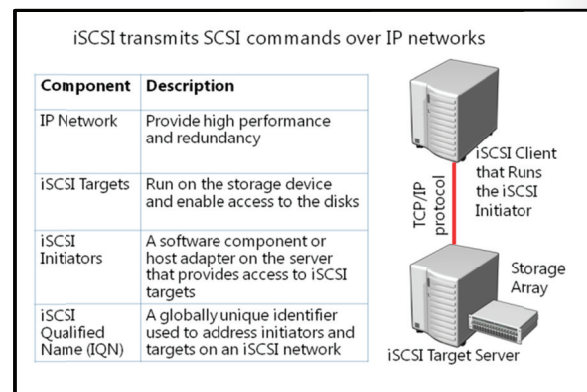
### What Is iSCSI?

iSCSI is a protocol that supports access to remote, SCSI-based storage devices over a TCP/IP network. iSCSI carries standard SCSI commands over IP networks to facilitate data transfers over intranets and to manage storage over long distances. You can use iSCSI to transmit data over LANs, WANs, or even over the larger Internet.

iSCSI relies on standard Ethernet networking architecture, and use of specialized hardware such as a host bus adapter (HBA) or network switches is optional. iSCSI uses TCP/IP (typically, TCP port 3260). This means that, iSCSI simply enables two hosts to negotiate (session establishment, flow control, and packet size, for example) and then exchange SCSI commands by using an existing Ethernet network. By doing this, iSCSI takes a popular, high performance, local storage bus subsystem architecture and emulates it over LANs and WANs, creating a SAN. Unlike some SAN protocols, iSCSI requires no specialized cabling; it can be run over existing switching and IP infrastructure. However, the performance of an iSCSI SAN deployment can be severely decreased if not operated on a dedicated network or subnet, as best practices recommend.



**Note:** While you can use a standard Ethernet network adapter to connect the server to the iSCSI storage device, you can also use dedicated HBAs.







## iSCSI Target Server

The iSCSI target server role service provides for software-based and hardware-independent iSCSI disk subsystem. You can use the iSCSI target server to create iSCSI targets and iSCSI virtual disks. You can then use the Server Manager to manage these iSCSI targets and virtual disks.

The iSCSI target server included in Windows Server 2012 provides the following functionality:

- *Network/diskless boot.* By using boot-capable network adapters or a software loader, you can use iSCSI targets to deploy diskless servers quickly. By using differencing virtual disks, you can save up to 90 percent of the storage space for the operating system images. This is ideal for large deployments of identical operating system images, such as a Hyper-V server farm or High Performance Computing (HPC) clusters.
- *Server application storage.* Some applications such as for example, Hyper-V and Exchange Server require block storage. The iSCSI target server can provide these applications with continuously available block storage. Because the storage is remotely accessible, it can also combine block storage for central or branch office locations.
- *Heterogeneous storage.* iSCSI target server supports iSCSI initiators that are not based on Windows, so you can share storage on Windows Servers in mixed environments.
- *Lab environments.* The iSCSI target server role enables your Windows Server 2012 computers to be a network-accessible block storage device. This is useful in situations such as when you want to test applications before deployment on SAN storage.

Enabling iSCSI target server to provide block storage takes advantage of your existing Ethernet network. No additional hardware is needed. If high availability is an important criterion, consider setting up a high availability cluster. With a high availability cluster, you will need shared storage for the cluster—either hardware Fibre Channel storage or a serial attached SCSI (SAS) storage array. iSCSI target server is directly integrated into the failover cluster feature as a cluster role.

## iSCSI Initiator

The iSCSI Initiator is included in Windows Server 2012 and Windows 8 as a service and installed by default. To connect your computer to an iSCSI target, you just have to start the service and configure it.



**Additional Reading:** Introduction of iSCSI Target in Windows Server 2012  
<http://blogs.technet.com/b/filecab/archive/2012/05/21/introduction-of-iscsi-target-in-windows-server-2012.aspx>

**Question:** When would you consider implementing diskless booting from iSCSI targets?



## Advanced iSCSI Configuration Options

In addition to configuring the basic iSCSI target server and iSCSI initiator settings, you can integrate these services into more advanced configurations.

### Locating iSCSI Storage

There are two common approaches for locating storage that is exposed to a network by an iSCSI Target.

The first approach is through the use of the iSCSI SendTargets command. This functionality is available within the iSCSI Initiator wizard of Windows Server. Using SendTargets in the iSCSI Initiator retrieves a list of available targets from a target device. To use this command, you must know both the IP address of the storage device that is hosting the targets, and whether the device is suitable for your storage needs. The iSCSI SendTargets command is only workable in smaller iSCSI environments because as the number of iSCSI targets increases in your company, the more complex this approach is.

The second approach is for large networks. On large networks, locating storage can be more difficult. One solution that can help you is the Internet Storage Name Service (iSNS), which is a Windows Server 2012 feature similar to Domain Name System (DNS) and lets you locate a target on several target devices.

iSNS contains three distinct services:

- *Name Registration Service.* This service enables initiators and targets to register and query the iSNS server directory for information about initiator and target IDs and addresses.
- *Network Zoning and Logon Control Service.* You can use this service to restrict iSNS initiators to zones so that iSCSI initiators do not discover any target devices outside their own zone or discovery domains. This prevents initiators from accessing storage devices that are not intended for their use. Logon control enables targets to determine which initiators can access them.
- *State Change Notification Service.* This service enables iSNS to notify clients of changes in the network, such as the addition or removal of targets, or changes in zoning membership. Only initiators that you register to receive notifications will get these packets, which reduces random broadcast traffic on the network.

### Configuring iSCSI for High Availability

Creating a single connection to iSCSI storage makes that storage available. However, it does not make that storage highly available. Losing the connection results in the server losing access to its storage. Therefore, most iSCSI storage connections are made redundant through one of two high-availability technologies: Multiple Connections per Session (MCS) and Multipath I/O (MPIO).

Although similar in the result they achieve, these two technologies use different approaches to achieve high availability for iSCSI storage connections.

MCS is a feature of the iSCSI protocol that:

- Enables multiple TCP/IP connections from the initiator to the target for the same iSCSI session.
- Supports automatic failover. If a failure were to occur, all outstanding iSCSI commands are reassigned to another connection automatically.
- Requires explicit support by iSCSI SAN devices, although the iSCSI target server role supports it.

#### Locating iSCSI storage. There are two approaches:

- iSCSI SendTargets. Use this command to receive a list of targets from an iSCSI target server
- iSNS. Use this for larger networks; similar to DNS, iSNS stores iSCSI targets

Configuring iSCSI for high availability There are two technologies:

- MCS. In the event of a failure, all outstanding iSCSI commands are reassigned to another connection automatically
- MPIO. If you have multiple network interface cards (NICs) in your iSCSI initiator and iSCSI target server, you can use MPIO to provide failover redundancy during network outages

MPIO is a different way to provide redundancy that:

- Requires a device specific module (DSM) if you want to connect to a third SAN device such as HP's EVA SAN connected to the iSCSI initiator. Windows includes a default MPIO DSM, installed as the Multipath I/O feature within Server Manager.
- Is widely supported. Many SANs can use the default DSM without any additional software, while others require a specialized DSM from the manufacturer.
- Is more complex to configure and not as fully automated during failover as MCS.

## Demonstration: Configuring iSCSI Target

In this demonstration, you will add an iSCSI target server role service and create an iSCSI virtual disk and iSCSI target on LON-DC1.

### Demonstration Steps

#### Add the iSCSI Target Server role service

1. On LON-DC1, in Server Manager, click the Dashboard button.
2. In the Add Roles and Features Wizard, install the following roles and features to the local server and accept the default values:
  - **File And Storage Services (Installed)\File and iSCSI Services\iSCSI Target Server**

#### Create two iSCSI virtual disks and an iSCSI target on LON-DC1

1. On LON-DC1, in Server Manager, in the navigation pane, click **File and Storage Services**, and then click **iSCSI**.
2. In the iSCSI VIRTUAL DISKS pane, click **TASKS**, and then in the TASKS drop-down list, click **New iSCSI Virtual Disk**. Create a virtual disk that has the following settings:
  - Name: **iSCSIDisk1**
  - Disk size: **5 GB**
  - iSCSI target: **New**
  - Target name: **LON-SVR2**
  - Access servers: **172.16.0.22**
3. On the **View results** page, wait until the creation is completed, and then close the **View Results** page.
4. In the iSCSI VIRTUAL DISKS pane, click **TASKS**, and then in the **TASKS** drop-down list, click **New iSCSI Virtual Disk**. Create a virtual disk that has these settings:
  - Name: **iSCSIDisk2**
  - Disk size: **5 GB**
  - iSCSI target: **LON-SVR2**
5. On the **View Results** page, wait until the creation is completed, and then close the **View Results** page.

## Demonstration: Connecting to the iSCSI Storage

In this demonstration, you will connect LON-SVR2 to the iSCSI target and verify the presence of the iSCSI drive.

### Demonstration Steps

#### Connect LON-SVR2 to the iSCSI target

1. Log on to LON-SVR2 with username of **Adatum\Administrator** and password of **Pa\$\$w0rd**.
2. In Server Manager on the **Tools** menu, open **iSCSI Initiator**.
3. In the **iSCSI Initiator Properties** dialog box, configure the following:
  - Quick Connect: **LON-DC1**
  - Discover targets: **iqn.1991-05.com.microsoft:lon-dc1-lon-svr2-target**

#### Verify the presence of the iSCSI drive

1. In Server Manager, on the **Tools** menu, open **Computer Management**.
2. In the Computer Management console, under Storage, access Disk Management.  
Notice that the new disks are added. They all are currently offline and not formatted.

## Lesson 3

# Configuring Storage Spaces in Windows Server 2012

Managing physical disks attached directly to a server proved to be a tedious task for the administrators. To overcome this problem, many organizations used SANs that basically grouped physically disks together.

However, SANs require special configuration and sometimes special hardware and are therefore expensive. To overcome these issues, storage spaces in Windows Server 2012 is a feature that pools disks together and presents them to the operating system as a single disk. This lesson explains how to configure and implement storage spaces.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe the use of storage spaces.
- Describe the features of storage spaces.
- Configure a storage space.
- Implement redundant storage spaces.

### What Are Storage Spaces?

A storage space is a storage virtualization capability built into Windows Server 2012 and Windows 8. You can use storage spaces to add physical disks of any type and size to a storage pool and create highly-available virtual disks from it. The primary advantage of storage spaces is that you do not manage single disks any longer, but manage them as one unit.

To create a highly-available virtual disk, you must have the following:


- *Disk drive*. This is a volume that you can access from your OS. For example, using a drive letter.
- *Virtual disk (or storage space)*. This resembles a physical disk from the perspective of users and applications. However, virtual disks are more flexible because it includes thin provisioning or just-in-time allocations and resiliency to physical disk failures with built-in functionality such as mirroring.
- *Storage pool*. A storage pool is a collection of one or more physical disks that you can use to create virtual disks. You can add to a storage pool any available physical disk that is not formatted or attached to another storage pool.
- *Physical disk*. These are connected physical disks such as SAS disks attached to your server. If you want to add them to a storage pool, they have to satisfy the following requirements:
  - One physical drive is required to create a storage pool; a minimum of two physical drives is required to create a resilient mirror virtual disk.
  - A minimum of three physical drives are required to create a virtual disk with resiliency through parity.

You can use storage spaces to add physical disks of any type and size to a storage pool and create highly-available virtual disks from it

To create a virtual disk, you need the following:

- One or more physical disks
- Storage pool that includes the disks
- Virtual drives (or storage spaces) that are created with disks from the storage pool
- Disk drives that are based on virtual drives

Virtual drives are not VHDs; they should be considered as a drive in Disk Manager



- Three-way mirroring requires at least five physical drives.
- Drives must be blank and unformatted, no volume must exist on them.
- Drives can be attached using different bus interfaces including iSCSI, SAS, Serial Advanced Technology Attachment (SATA), SCSI, and USB. You cannot use SATA, USB or SCSI disks in a failover cluster.

A storage space is a feature available for both NTFS and ReFS volumes that can provide redundancy and pooled storage for many internal and external drives of different sizes and interfaces.

## Storage Spaces Features

To configure storage spaces as per your requirements, you must have to consider the features described in the following table before you implement virtual disks.

To optimally use storage spaces in your environment, you should consider the following features

Feature	Options
Storage layout	Simple Two-way or three-way mirrors Parity
Disk sector size	512 or 512e
Drive allocation	Data-store Manual Hot-Spare
Provisioning schemes	Thin provisioning space Fixed provisioning space

To support failover clustering, all assigned drives must support a multi-initiator protocol, such as Serial-Attached SCSI (SAS)

Feature	Description
Storage layout	<p>This defines the number of disks from the storage pool that are allocated. Valid options are:</p> <ul style="list-style-type: none"> <li>• <i>Simple</i>. A simple space has data striping but no redundancy. In data striping, logically sequential data is segmented across all disks in a way that access to these sequential segments can be made to different physical storage drives. Striping makes it possible to access multiple segments of data at the same time. Do not host important data on a simple volume, because it provides no failover capabilities when the disk where the data is stored on fails.</li> <li>• <i>Two-way and three-way mirrors</i>. Mirror spaces maintain two or three copies of the data they host (two data copies for two-way mirrors and three data copies for three-way mirrors). Duplication happens with every write to ensure all data copies are always current. Mirror spaces also stripe the data across multiple physical drives. Mirror spaces are attractive because of their greater data throughput and lower access latency. They also do not introduce a risk of corrupting at-rest data and do not require the additional journaling stage when writing data.</li> <li>• <i>Parity</i>. A parity space resembles a simple space. Data, along with parity information, is striped across multiple physical drives. Parity enables storage spaces to continue to service read and write requests even when a drive has failed. Parity is always rotated across available disks to enable IO optimization. A storage space requires a minimum of three physical drives for parity spaces. Parity spaces have increased resiliency through journaling.</li> </ul>

Feature	Description
Disk sector size	A storage pool's sector size is set the moment it is created. If the list of drives being used contains only 512 and 512e drives, the pool is defaulted to 512e. However, if the list contains at least one 4-KB drive, the pool sector size is defaulted to 4 KB. Optionally, an administrator can explicitly define the sector size that all contained spaces in the pool will inherit. After an administrator defines this, Windows will only enable addition of drives that have a compliant sector size, that is: 512 or 512e for a 512e storage pool and 512, 512e, or 4 KB for a 4-KB pool.
Cluster disk requirement	Failover clustering prevents interruption to workloads or data if there is a computer failure. For a pool to support failover, clustering all assigned drives must support a multi-initiator protocol, such as SAS.
Drive allocation	<p>This defines how the drive is allocated to the pool. Options are:</p> <ul style="list-style-type: none"> <li>• <i>Data-store</i>. This is the default allocation when any drive is added to a pool. Storage spaces can automatically select available capacity on data-store drives for both storage space creation and just-in-time allocation.</li> <li>• <i>Manual</i>. Administrators can choose to specify manual as the usage type for drives added to a pool. A manual drive is not automatically used as part of a storage space unless it is specifically selected at the creation of that storage space. This usage property lets administrators specify particular types of drives for use by only certain storage spaces.</li> <li>• <i>Hot-Spare</i>. Drives added as "Hot-Spares" to a pool are reserve drives that are not used in the creation of a storage space. If a failure occurs on a drive that is hosting columns of a storage space, a reserve drive is called on to replace the failed drive.</li> </ul>
Provisioning schemes	<p>You can provision a virtual disk by using two schemes:</p> <ul style="list-style-type: none"> <li>• <i>Thin provisioning space</i>. Thin provisioning is a mechanism that enables storage to be easily allocated on a just-enough and just-in-time basis. Storage capacity in the pool is organized into provisioning slabs that are not allocated until the point in time when datasets grow to actually require the storage. Instead of the traditional fixed storage allocation method, where large pools of storage capacity are allocated but may remain unused, thin provisioning optimizes use of available storage. Organizations are also able to save on operating costs such as electricity and floor space associated with keeping unused drives spinning.</li> <li>• <i>Fixed provisioning space</i>. In storage spaces, fixed provisioned spaces also use the flexible provisioning slabs. The difference here is that the storage capacity is allocated up front, at the time that the space is created.</li> </ul>



**Note:** Storage spaces allows for the creation of both thin and fixed provisioning virtual disks within the same storage pool. Having both provisioned types in the same storage pool is very convenient especially when they are related to the same workload. For example, you can choose to have a thin provisioning space to host a database and a fixed provisioning space to host its log.

## Demonstration: Configuring a Storage Space

In this demonstration, you will create a storage pool and create a simple virtual disk and a volume.

### Demonstration Steps

#### Create a storage pool

1. On LON-SVR2, in Server Manager, navigate to **File and Storage Services**, and **Storage Pools**.
2. In the STORAGE POOLS pane, create a **New Storage Pool** named **StoragePool1**, and then add all available disks.

#### Create a simple virtual disk and a volume

1. In the VIRTUAL DISKS pane, create a **New Virtual Disk** with these settings:
  - Storage pool: **StoragePool1**
  - Disk name: **Simple vDisk**
  - Storage layout: **Simple**
  - Provisioning type: **Thin**
  - Size: **2 GB**
2. On the **View results** page, wait until the creation is completed, make sure **Create a volume when this wizard closes** is checked, and then click **Close**.
3. In the New Volume Wizard, create a volume with these settings:
  - Virtual disk: **Simple vDisk**
  - File system: **ReFS**
  - Volume label: **Simple Volume**

## Demonstration: Implementing Redundant Storage Spaces

In this demonstration, you will create a redundant virtual disk and a volume, simulate a drive failure, and test volume access.

### Demonstration Steps

#### Create a redundant virtual disk and a volume

1. On LON-SVR2, in Server Manager, in the VIRTUAL DISKS pane, click **TASKS**, and then in the **TASKS** drop-down list, select **New Virtual Disk** and create a virtual disk with these settings:
  - Storage pool: **StoragePool1**
  - Disk name: **Mirrored vDisk**
  - Storage layout: **Mirror**
  - Provisioning type: **Thin**
  - Size: **5 GB**
2. On the **View results** page, wait until the creation is completed, make sure **Create a volume when this wizard closes** is checked, and then click **Close**.

3. In the New Volume Wizard, create a volume with these settings:
  - o Virtual disk: **Mirrored vDisk**
  - o File system: **ReFS**
  - o Volume label: **Mirrored Volume**
4. On the **Completion** page, wait until the creation is completed, and then click **Close**.
5. On the Start screen, type **command prompt** and then press Enter.
6. At the command prompt, type the following command and then press Enter:

```
Copy C:\windows\system32\write.exe F:\
```

7. In Server Manager, on the menu bar, click **Tools** and then in the **Tools** drop-down list, select **Computer Management**.
8. In the Computer Management console, under **Storage**, click **Disk Management**.  
Notice that the two volumes E: and F: are available.

#### **Simulate a drive failure and test volume access**

1. On LON-DC1, in Server Manager, in the left pane, click **File and Storage Services**.
2. In the File and Storage Services pane, click **iSCSI**.
3. In the iSCSI VIRTUAL DISKS pane, in the **LON-DC1** list, right-click **iSCSIDisk1.vhd**, and then click **Disable iSCSI Virtual Disk**.
4. Switch to LON-SVR2.
5. In the Computer Management console, under **Storage**, right-click **Disk Management**, and then in drop-down list, select **Rescan Disks**.  
Notice that the Simple Volume (E:) is not available and the Mirrored Volume (F:) is available.
6. On the taskbar, open Windows Explorer and then click **Mirrored Volume (F:)**. You should now see **write.exe** in the file list.
7. In Server Manager, in the STORAGE POOLS pane, on the menu bar, click the **Refresh "Storage Pools"** button. Notice the warning that appears right next to **Mirrored vDisk**.
8. In the VIRTUAL DISKS pane, in the drop-down list, right-click **Simple vDisk**, and then select **Properties**.
9. In the **Simple vDisk Properties** dialog box, in the navigation pane, click **Health**.  
Notice the **Health Status** that should indicate **Unknown**. The **Operational Status** should indicate **Detached**. This means that the disk is not available on this computer any longer.
10. In the VIRTUAL DISKS pane, right-click **Mirrored vDisk**, and then in the drop-down list, select **Properties**.
11. In the **Mirrored vDisk Properties** window, in the navigation pane, click **Health**.  
Notice the **Health Status** should indicate a **Warning**. The Operational Status should indicate **Incomplete** or **Degraded**.



## Lab A: Managing Storage for Servers Based on Windows Server 2012

### Scenario

With the growth in A. Datum, the requirements for managing storage and shared file access has also expanded. Although the cost of storage has decreased significantly over the last few years, the data produced by the A. Datum business groups has increased even more. The organization is considering alternative ways to reduce the cost of storing data on the network in addition to the options for optimizing data access for both physical and virtual servers. Also, to meet some requirements for high availability, the organization is exploring options for making storage highly available.

As one of the senior network administrators at A. Datum, you are responsible for implementing some new file storage technologies for the organization. You will implement iSCSI storage to provide a less complex option for deploying large amounts of storage in the organization. You will also implement the storage spaces on the Windows Server 2012 servers to simplify storage access and to provide redundancy at the storage level.

### Objectives

After completing this lab, you will be able to:

- Configure iSCSI storage for Windows Server 2012 servers.
- Configure a redundant storage space.

### Lab Setup

Estimated time: **40 minutes**

Virtual Machine(s)	20417A-LON-DC1 20417A-LON-SVR2
User Name	Adatum\Administrator
Password	Pa\$\$w0rd

### Lab Setup

For this lab, you will use the available virtual machine environment. Before you begin the lab, you must complete the following steps:

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, click **20417A-LON-DC1**, and in the **Actions** pane, click **Start**.
3. In the **Actions** pane, click **Connect**. Wait until the virtual machine starts.
4. Log on using the following credentials:
  - a. User name: **Adatum\Administrator**
  - b. Password: **Pa\$\$w0rd**
5. Repeat steps 2 to 4 for **20417A-LON-SVR2**.

For this lab, on **20417A-LON-SVR2**, disable Routing and Remote Access. In Server Manager, click **Tools**, and then click **Routing and Remote Access**. In the **Routing and Remote Access** console, right-click **LON-SVR2** and then click **Disable Routing and Remote Access**.

## Exercise 1: Configuring iSCSI Storage

### Scenario

In order to reduce the cost and complexity of configuring centralized storage, A. Datum is exploring the option of using iSCSI to provide storage. To get started, you will install and configure the iSCSI targets, and configure access to the targets by configuring the iSCSI initiators.

The main tasks for this exercise are as follows:

1. Install the iSCSI Target feature.
2. Configure the iSCSI targets.
3. Configure MPIO.
4. Connect to and configure the iSCSI targets.

#### ► Task 1: Install the iSCSI Target feature

1. Log on to LON-DC1 with username of **Adatum\Administrator** and the password of **Pa\$\$w0rd**.
2. In Server Manager, start the **Add Roles and Features Wizard**, install the following roles and features to the local server and accept the default values:
  - **File And Storage Services (Installed)\File and iSCSI Services\iSCSI Target Server**

#### ► Task 2: Configure the iSCSI targets

1. On LON-DC1, in Server Manager, in the navigation pane, click **File and Storage Services**, and then click **iSCSI**.
2. Create a virtual disk with these settings:
  - Storage location: **C:**
  - Disk name: **iSCSIDisk1**
  - Size: **5 GB**
  - iSCSI target: **New**
  - Target name: **lon-svr2**
  - Access servers: **172.16.0.22** and **131.107.0.2**
3. On the **View results** page, wait until the creation is completed, and then click **Close**.
4. Create a New iSCSI Virtual Disk with these settings:
  - Storage location: **C:**
  - Disk name: **iSCSIDisk2**
  - Size: **5 GB**
  - iSCSI target: **lon-svr2**
5. Create a New iSCSI Virtual Disk with these settings:
  - Storage location: **C:**
  - Disk name: **iSCSIDisk3**
  - Size: **5 GB**
  - iSCSI target: **lon-svr2**

6. Create a New iSCSI Virtual Disk with these settings:

- Storage location: **C:**
- Disk name: **iSCSIDisk4**
- Size: **5 GB**
- iSCSI target: **lon-svr2**

7. Create a New iSCSI Virtual Disk with these settings:

- Storage location: **C:**
- Disk name: **iSCSIDisk5**
- Size: **5 GB**
- iSCSI target: **lon-svr2**

► **Task 3: Configure MPIO**

1. Log on to LON-SVR2.
2. In Server Manager, start the Add Roles and Features Wizard and install the **Multipath I/O** feature.
3. In Server Manager, on the **Tools** menu, open **iSCSI Initiator**, and configure the following:
  - Enable the **iSCSI Initiator** service
  - **Quick Connect** to target: **LON-DC1**
4. In Server Manager, on the **Tools** menu, open **MPIO**, and configure the following:
  - Enable **Add support for iSCSI devices** on Discover Multi-paths
5. After the computer restarts, log on to LON-SVR2, on the **Tools** menu in Server Manager, open **MPIO** and verify that **Device Hardware ID MSFT2005iSCSIBusType\_0x9** is added to the list.

► **Task 4: Connect to and configure the iSCSI targets**

1. On LON-SVR2, in Server Manager, on the **Tools** menu, open **iSCSI Initiator**.
2. In the **iSCSI Initiator Properties** dialog box, perform the following steps:
  - a. **Disconnect** all **Targets**.
  - b. **Connect** and **Enable multi-path**.
  - c. Set **Advanced** options as follows:
    - Local Adapter: **Microsoft iSCSI Initiator**
    - Initiator IP: **172.16.0.22**
    - Target Portal IP: **172.16.0.10 / 3260**
  - d. **Connect** to another target, **enable multi-path**, and configure the following **Advanced** settings:
    - Local Adapter: Microsoft iSCSI Initiator
    - Initiator IP: 131.107.0.2
    - Target Portal IP: 131.107.0.1 / 3260
3. In the **Targets** list, open **Devices** for **iqn.1991-05.com.microsoft:lon-dc1-lon-svr2-target**, access the MPIO information, and then verify that in **Load balance policy**, **Round Robin** is selected. Verify that two paths are listed by looking at the IP addresses of both network adapters.

**Results:** After completing this exercise, you will have configured and connected to iSCSI targets.

## Exercise 2: Configuring a Redundant Storage Space

### Scenario

After you have configured the iSCSI components, you want to take advantage of the storage pools to simplify the configuration of storage on the Windows Server 2012 servers. To meet some requirements for high availability, you decided to evaluate redundancy features in storage spaces. Also, you want to test provisioning of new disks to the storage pool.

The main tasks for this exercise are as follows:

1. Create a storage pool by using the iSCSI disks attached to the server.
2. Create a 3-way mirrored disk.
3. Copy a file to the volume and verify visibility in Windows Explorer.
4. Disconnect an iSCSI disk.
5. Verify that the file is still accessible and check the health of the virtual disk.
6. Add a new iSCSI virtual disk.
7. Add the new disk to the storage pool and extend the virtual disk.

#### ► Task 1: Create a storage pool by using the iSCSI disks attached to the server

1. On LON-SVR2, open Server Manager by clicking the icon on the taskbar.
2. In the navigation pane, click **File and Storage Services**, and then in the Servers pane, click **Storage Pools**.
3. Create a storage pool with the following settings:
  - o Name: **StoragePool1**
4. On the **View results** page, wait until the creation is completed, then click **Close**.

#### ► Task 2: Create a 3-way mirrored disk

1. On LON-SVR2, in Server Manager, in the VIRTUAL DISKS pane, create a virtual disk with these settings:
  - o Storage pool: **StoragePool1**
  - o Name: **Mirrored vDisk**
  - o Storage Layout: **Mirror**
  - o Resiliency settings: **Three-way mirror**
  - o Provisioning type: **Thin**
  - o Virtual disk size: **10 GB**
2. On the **View results** page, wait until the creation is completed, make sure **Create a volume when this wizard closes** is checked, and then click **Close**.
3. In the New Volume Wizard, create a volume with these settings:
  - o Virtual disk: **Mirrored vDisk**
  - o Drive letter: **E**
  - o File system: **ReFS**
  - o Volume label: **Mirrored Volume**
4. On the **Completion** page, wait until the creation is completed, and then click **Close**.

► **Task 3: Copy a file to the volume and verify visibility in Windows Explorer**

1. On the Start screen, type **command prompt** and then press ENTER.
2. Type the following command:

```
Copy C:\windows\system32\write.exe E:\
```

3. Use Windows Explorer and access **Mirrored Volume (E:)**. You should now see write.exe in the file list.

► **Task 4: Disconnect an iSCSI disk**

1. Switch to LON-DC1.
2. In the iSCSI VIRTUAL DISKS pane, in the LON-DC1 list, disable the iSCSI Virtual Disk named **iSCSIDisk1.vhd**.

► **Task 5: Verify that the file is still accessible and check the health of the virtual disk**

1. Switch to LON-SVR2.
2. Use Windows Explorer and open **E:\write.exe** to make sure access to the volume is still available.
3. In Server Manager, in the STORAGE POOLS pane, on the menu bar, click the **Refresh "Storage Pools"** button. Notice the warning that appears right next to Mirrored vDisk.
4. In the VIRTUAL DISK pane, right-click **Mirrored vDisk**, in the drop-down list, select **Properties**.
5. In Mirrored vDisk Properties window, in the **Health** pane, notice that the Health Status indicates a Warning. The Operational Status should indicate Degraded.

► **Task 6: Add a new iSCSI virtual disk**

1. Switch to LON-DC1.
2. In Server Manager, in the iSCSI Virtual VIRTUAL DISKS pane, click **TASKS**, and then in the **TASKS** drop-down list, select **New iSCSI Virtual Disk**.
3. Create a NEW iSCSI Virtual Disk with these settings:
  - Storage location: **C:**
  - Disk name: **iSCSIDisk6**
  - Size: **5 GB**
  - iSCSI target: **lon-svr2**

► **Task 7: Add the new disk to the storage pool and extend the virtual disk**

1. Switch to LON-SVR2.
2. In Server Manager, in the STORAGE POOLS pane, on the menu bar, click the **Refresh "Storage Pools"** button.
3. In the STORAGE POOLS pane, right-click **StoragePool1**, and then in the drop-down list, select **Add Physical Disk**, and add **PhysicalDisk1 (LON-SVR2)**.
4. In the VIRTUAL DISKS pane, right-click **Mirrored vDisk**, and then in the drop-down list, select **Extend Virtual Disk** and extend the disk to 15 GB.

**Results:** After completing this exercise, you will have created a storage pool and added a new disk to the storage pool and extended the disk.

► **To prepare for the next lab**

When you finish the lab, revert the virtual machines to their initial state. To do this, perform the following steps:

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20417A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20417A-LON-SVR2**.

## Lesson 4

# Configuring BranchCache in Windows Server 2012

Branch offices have unique management challenges. A branch office typically has slow connectivity to the enterprise network and limited infrastructure for securing servers. Therefore, the challenge is being able to provide efficient access to network resources for users in branch offices. The BranchCache feature helps you overcome these problems by caching files so they do not have to be transferred over the network again.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe how BranchCache works.
- Describe the BranchCache requirements.
- Configure the BranchCache server settings.
- Configure the BranchCache client settings.
- Configure BranchCache.
- Describe how to monitor BranchCache.

### How Does BranchCache Work?


The BranchCache feature introduced with Windows Server 2008 R2 and Windows 7 reduces the network use on WAN connections between branch offices and the headquarters by locally caching frequently used files on computers in the branch office.

BranchCache improves the performance of applications that use one of the following protocols:

- *HTTP or HTTPS protocols.* These protocols are used by web browsers and other applications.
- *Server message block (SMB), including signed SMB traffic protocol.* This protocol is used for accessing shared folders.
- *Background Intelligent Transfer Service (BITS).* A Windows component that distributes content from a server to clients by using only idle network bandwidth.

BranchCache retrieves data from a server when the client requests the data. Because BranchCache is a passive cache, it will not increase WAN use. BranchCache only caches the read requests and will not interfere when a user saves a file.

BranchCache improves the responsiveness of common network applications that access intranet servers across slow WAN links. Because BranchCache does not require additional infrastructure, you can improve the performance of remote networks by deploying Windows 7 or 8 to client computers and Windows Server 2012 to servers, and by enabling the BranchCache feature.



BranchCache works seamlessly with network security technologies, including Secure Sockets Layer (SSL), SMB Signing, and end-to-end Internet Protocol Security (IPsec). You can use BranchCache to reduce the network bandwidth use and improve application performance, even if the content is encrypted.

You can configure BranchCache to use Hosted Cache mode or Distributed Cache mode:

- *Hosted Cache.* This mode operates by deploying a computer that is running Windows Server 2008 R2 or later versions as a hosted cache server in the branch office. Client computers are configured with the fully qualified domain name (FQDN) of the host computer so that they can retrieve content from the Hosted Cache when available. If the content is not available in the Hosted Cache, the content is retrieved from the content server by using a WAN link and then provided to the Hosted Cache so that the successive client requests can get it from there.
- *Distributed Cache.* You can configure BranchCache in the Distributed Cache mode for small remote offices without requiring a server. In this mode, local client computers running Windows 7 or Windows 8 keep a copy of the content and make it available to other authorized clients that request the same data. This eliminates the need to have a server in the branch office. However, unlike the Hosted Cache mode, this configuration works across a single subnet only. In addition, clients who hibernate or disconnect from the network cannot provide content to other requesting clients.

BranchCache in Windows Server 2012 is improved in the following ways:

- More than one hosted cache servers per location to allow for scale.
- New underlying database that uses the Extensible Storage Engine (ESE) database technology from Microsoft Exchange Server. This enables a hosted cache server to store significantly more data (in the order of terabytes).
- The deployment is made much simpler such that you do not require a Group Policy Object (GPO) for each location. A single GPO that contains the settings is all that is required to deploy BranchCache.

### **How Client Computer Retrieves Data by Using BranchCache**

When BranchCache is enabled on the client computer and the server, the client computer performs the following process to retrieve data when using the HTTP, HTTPS, or SMB protocol:

1. The client computer that is running Windows 7 connects to a content server that is running Windows Server 2008 R2 in the head office and requests content similar to the way it would retrieve content without using BranchCache.
2. The content server in the head office authenticates the user and verifies that the user is authorized to access the data.
3. The content server in the head office returns identifiers or hashes of the requested content to the client computer instead of sending the content itself. The content server sends that data over the same connection that the content would have typically been sent.
4. Using retrieved identifiers, the client computer does the following:
  - If you configure it to use Distributed Cache, the client computer multicasts on the local subnet to find other client computers that have already downloaded the content.
  - If you configure it to use Hosted Cache, the client computer searches for the content on the configured Hosted Cache.
5. If the content is available in the branch office, either on one or more clients or on the Hosted Cache, the client computer retrieves the data from the branch office and ensures that the data is updated and has not been tampered with or corrupted.



6. If the content is not available in the remote office, the client computer retrieves the content directly from the server across the WAN link. The client computer then either makes it available on the local network to other requesting client computers (Distributed Cache mode) or sends it to the Hosted Cache, where it is made available to other client computers.

## BranchCache Requirements

BranchCache optimizes traffic flow between head office and branch offices. Windows Server 2008 R2, Windows Server 2012, and clients based on client computers running Windows 7 or Windows 8 Enterprise Edition can only benefit from BranchCache. The earlier versions of Windows operating systems do not benefit from this feature. You can cache only the content that is stored on file servers or web servers running Windows Server 2008 R2 or Windows Server 2012 by using BranchCache.

Requirements for using BranchCache	Requirements for Distributed and Hosted Cache modes
<ul style="list-style-type: none"> <li>• Install the BranchCache feature or the BranchCache for Network Files role service on the server that is hosting the content</li> <li>• Configure client computers, either by using Group Policy or the <code>netsh branchcache set service</code> command</li> </ul>	<ul style="list-style-type: none"> <li>• In the Distributed Cache mode, no server is required in the branch office; just Windows 7, Windows 8 or Windows Server 2008 R2 or later as client computers are required</li> <li>• In the Hosted Cache mode, Windows Server 2012 server must be configured for BranchCache host in the branch office</li> <li>• The BranchCache host server must have a digital certificate</li> </ul>

### Requirements for Using BranchCache

To use BranchCache, you must perform the following tasks:

- Install the BranchCache feature or the BranchCache for Network Files role service on the server running Windows Server 2012 that is hosting the data.
- Configure client computers either by using Group Policy or the **netsh branchcache set service** command.

If you want to use BranchCache for caching content from the web server, you must install the BranchCache feature on the web server. Additional configurations are not needed. If you want to use BranchCache to cache content from the file server, you must install the BranchCache for the Network Files role service on the file server, configure hash publication for BranchCache, and create BranchCache-enabled file shares.

BranchCache is supported on Full Installation of Windows Server 2012 and on Server Core.

### Requirements for Distributed Cache and Hosted Cache Modes

In the Distributed Cache mode, BranchCache works across a single subnet only. If client computers are configured to use the Distributed Cache mode, any client computer can search locally for the computer that has already downloaded and cached the content by using a multicast protocol called WS-Discovery. In the Distributed Cache mode, content servers across the WAN link must run Windows Server 2008 R2 or later versions, and the clients in the branch must run at least Windows 7 or Windows Server 2008 R2. You should configure the client firewall to enable incoming traffic, HTTP, and WS-Discovery.

In the Hosted Cache mode, the client computers are configured with the FQDN of the host server to retrieve content from the Hosted Cache. Therefore, the BranchCache host server must have a digital certificate, which is used to encrypt communication with client computers. In the Hosted Cache mode, content servers across the WAN link must run Windows Server 2008 R2 or later versions. Hosted Cache in the branch must run Windows Server 2008 R2 or later versions and the client in the branch must run at least Windows 7. You must configure a firewall to enable incoming HTTP traffic from the Hosted Cache server. In both cache modes, BranchCache uses the HTTP protocol for data transfer between client computers and the computer that is hosting the cached data.



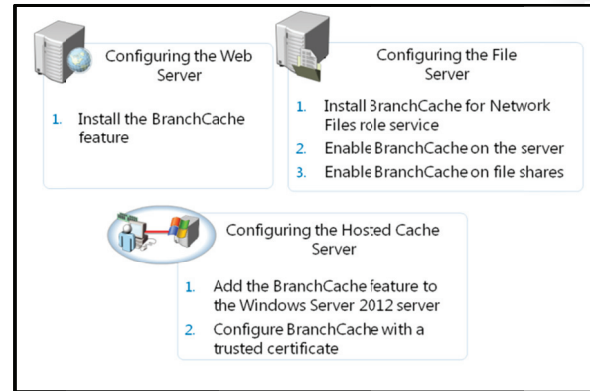
**Additional Reading:** Windows Server 2008 R2

<http://go.microsoft.com/fwlink/?LinkID=214828&clcid=0x409>

## Configuring BranchCache Server Settings

You can use BranchCache to cache web content, which is delivered by HTTP or HTTPS. You can also use BranchCache to cache shared folder content, which is delivered by the SMB protocol. By default, BranchCache is not installed on Windows Server 2012.

The following table lists the servers that you can configure for BranchCache.



Server	Description
Web server or Background Intelligent Transfer Service (BITS) server	To configure a Windows Server 2012 web server or an application server that uses the BITS protocol, install the BranchCache feature. Ensure that the BranchCache service has started. Then, configure clients who will use the BranchCache feature; no additional configuration of the web server is needed.
File server	The BranchCache for the Network Files role service of the File Services server role has to be installed before you can enable BranchCache for any file shares. After you install the BranchCache for the Network Files role service, use Group Policy to enable BranchCache on the server. Finally, you must configure each file share to enable BranchCache. You also have to configure clients who will use the BranchCache feature.
Hosted Cache server	<p>For the Hosted Cache mode, you must add the BranchCache feature to the Windows Server 2012 server that you are configuring as a Hosted Cache server.</p> <p>To help secure communication, client computers use Transport Layer Security (TLS) when communicating with the Hosted Cache server. To support authentication, the Hosted Cache server must be provisioned with a certificate that is trusted by clients and is suitable for server authentication.</p> <p>By default, BranchCache allocates five percent of disk space on the active partition for hosting cache data. However, you can change this value by using Group Policy or the netsh tool.</p>

```
netsh branchcache set service mode=hostedclient location=<Hosted Cache server>
```

- Setting the cache size
- Setting the location of the Hosted Cache server
- Clearing the cache
- Creating and replicating a shared key for using in a server cluster

## Configuring the Client Firewall To Enable BranchCache Protocols

In the Distributed Cache mode, BranchCache clients use the HTTP protocol for data transfer between client computers and the WS-Discovery protocol (WSD) for cached content discovery. You should configure the client firewall to enable the following incoming rules:

- BranchCache–Content Retrieval (Uses HTTP)
- BranchCache–Peer Discovery (Uses WSD)

In the Hosted Cache mode, BranchCache clients use the HTTP protocol for data transfer between client computers, but it does not use the WS-Discovery protocol. In the Hosted Cache mode, you should configure the client firewall to enable the incoming rule, BranchCache–Content Retrieval (Uses HTTP).

## Additional Configuration Tasks for BranchCache

After you configure BranchCache, clients can access the cached data in BranchCache-enabled content servers, available locally in the branch office, and not across a slow WAN link. You can modify BranchCache settings and perform additional configuration tasks, such as:

- Setting the cache size
- Setting the location of the Hosted Cache server
- Clearing the cache
- Creating and replicating a shared key for using in a server cluster

## Demonstration: How to Configure BranchCache

In this demonstration, you will add BranchCache for the Network Files role service, configure BranchCache in Local Group Policy Editor, and enable BranchCache for a file share.

### Demonstration Steps

#### Add BranchCache for the Network Files role service

1. Log on to LON-DC1 and open Server Manager.
2. In the **Add Roles and Features Wizard**, install the following roles and features to the local server:
  - **File And Storage Services (Installed)\File and iSCSI Services\BranchCache for Network Files**

#### Enable BranchCache for the server

1. On the Start screen, type **gpedit.msc**, and press **ENTER**.
2. Browse to **Computer Configuration\Administrative Templates\Network\Lanman Server** and do the following:
  - Enable **Hash Publication for BranchCache**
  - Select **Allow hash publication only for shared folder on which BranchCache is enabled**

#### Enable BranchCache for a file share

1. Open Windows Explorer and create a folder named **Share** on **C:\**.
2. Configure the **Share** folder properties as follows:
  - Enable **Share this folder**
  - Check **Enable BranchCache** in **Offline Settings**

## Monitoring BranchCache


After the initial configuration, you might want to verify that BranchCache is configured correctly and functioning correctly. You can use the **netsh branchcache show status all** command to display the BranchCache service status. On client and Hosted Cache servers, additional information such as the location of the local cache, the size of the local cache, and the status of the firewall rules for HTTP and WS-Discovery protocols that BranchCache uses is shown.

You can also use the following tools to monitor BranchCache:

- *Event Viewer.* You can use this tool to monitor BranchCache events in Event Viewer.
- *Performance counters.* You can use this tool to monitor BranchCache work and performance by using the BranchCache performance monitor counters. BranchCache performance monitor counters are useful debugging tools for monitoring BranchCache effectiveness and health. You can also use BranchCache performance monitor for determining the bandwidth savings in the Distributed Cache mode or in the Hosted Cache mode. If you have System Center Operations Manager 2007 SP2 or later versions implemented in the environment, you can use Windows BranchCache Management Pack for Operations Manager 2007

The BranchCache monitoring tools include:

- The netsh branchcache show status all command
- Event Viewer
- Performance counters



## Lab B: Implementing BranchCache

### Scenario

A. Datum has deployed a new branch office. This office has a single server. To support branch staff requirements, you must configure BranchCache. Data is centralized at the head office. To reduce WAN use out to the branch office, you must configure BranchCache for these data.

### Objectives

After completing this lab, you will be able to:

- Perform initial configuration tasks for BranchCache.
- Configure BranchCache clients.
- Configure BranchCache on the branch server.

### Lab Setup

Estimated time: **40 minutes**

Virtual Machine(s)	20417A-LON-DC1 20417A-LON-SVR1 20417A-LON-CL1 20417A-LON-CL2
User Name	Adatum\Administrator
Password	Pa\$\$w0rd

### Lab Setup

For this lab, you will use the available virtual machine environment. Before you begin the lab, you must complete the following steps:

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, click **20417A-LON-DC1**, and in the **Actions** pane, click **Start**.
3. In the **Actions** pane, click **Connect**. Wait until the virtual machine starts.
4. Log on using the following credentials:
  - a. User name: **Adatum\Administrator**
  - b. Password: **Pa\$\$w0rd**
5. Do not start **20417A-LON-SVR1**, **20417A-LON-CL1** and **20417A-LON-CL2** until directed to do so.

## Exercise 1: Performing Initial Configuration Tasks for BranchCache

### Scenario

Before you can configure the BranchCache feature for your branch offices, you must configure the network components.

The main tasks for this exercise are as follows:

1. Configure LON-DC1 to use BranchCache.
2. Simulate slow link to the branch office.
3. Enable a file share for BranchCache.
4. Configure client firewall rules for BranchCache.

#### ► Task 1: Configure LON-DC1 to use BranchCache

1. Switch to LON-DC1.
2. Open Server Manager and install the **BranchCache for network files** role service.
3. Open the Local Group Policy Editor (**gpedit.msc**).
4. Navigate to and open **Computer Configuration/Administrative Templates/Network /Lanman Server/Hash Publication for BranchCache**. Enable this setting and then select **Allow hash publication only for shared folders on which BranchCache is enabled**.

#### ► Task 2: Simulate slow Link to the branch office

1. Navigate to **Computer Configuration\Windows Settings\Policy-based QoS**.
2. Create a new policy with the following settings:
  - Name: **Limit to 100Kbps**
  - Specify Outbound Throttle Rate: **100**



**Note:** This task is required to simulate a slow network connection in a test environment where all the computers are connected by a fast network connection.

#### ► Task 3: Enable a file share for BranchCache

1. In Windows Explorer, create a new folder named **C:\Share**.
2. Share this folder with the following properties:
  - Sharename: **Share**
  - Permissions: default
  - Caching: **Enable BranchCache**
3. Copy **C:\Windows\System32\mspaint.exe** to the **C:\Share** folder.

#### ► Task 4: Configure client firewall rules for BranchCache

1. On LON-DC1, open **Group Policy Management**.
2. Navigate to **Forest: Adatum.com\Domains\Adatum.com\Default Domain Policy**. Open the policy for editing.
3. Navigate to **Computer Configuration\Policies\Windows Settings\Security Settings\Windows Firewall with Advanced Security\Windows Firewall with Advanced Security\Inbound Rules**.

4. Create a new inbound firewall rule with the following properties:
  - Rule type: **predefined**
  - Use **BranchCache – Content Retrieval (Uses HTTP)**
  - Action: **Allow**
5. Create a new inbound firewall rule with the following properties:
  - Rule type: **predefined**
  - Use **BranchCache – Peer Discovery (Uses WSD)**
  - Action: **Allow**

**Results:** At the end of this exercise, you will have deployed BranchCache, configured a slow link, and enabled BranchCache on a file share.

## Exercise 2: Configuring BranchCache Client Computers

### Scenario

After you have configured the network components, you must now make sure the client computers are configured correctly. This is a preparatory task to be able to use BranchCache.

The main task for this exercise is to configure client computers to use BranchCache in the Hosted Cache mode.

### ► Task: Configure client computers to use BranchCache in the Hosted Cache mode

1. On LON-DC1, in Group Policy Management Editor, and configure the following at **Computer Configuration\Policies\Administrative Templates\Network\BranchCache**:
  - Turn on BranchCache: Enable
  - Set BranchCache Hosted Cache mode: Enable
  - Type the name of the hosted Cache server: **LON-SVR1.adatum.com**
  - Configure BranchCache for network files: Enable
  - Type the maximum round trip network latency value (milliseconds) after which caching begins: **0**
2. Start the 20417A-LON-CL1, open a Command Prompt window, and refresh the Group Policy settings (**gpupdate /force**).
3. At the command prompt, type **netsh branchcache show status all**, and then press Enter.
4. Start the 20417A-LON-CL2, open the Command Prompt window, and refresh the Group Policy settings (**gpupdate /force**).
5. At the command prompt, type **netsh branchcache show status all**, and then press Enter.



**Note:** To test BranchCache in a test lab, you should deploy two client computers. This enables you to request a file from one of the client computers, and then verify that the file is retrieved from the local cache on the second client computer.

**Results:** At the end of this exercise, you will have configured the client computers for BranchCache.



## Exercise 3: Configuring BranchCache on the Branch Server

### Scenario

The next step you must perform is to configure a file server for the BranchCache feature. You will install the BranchCache feature and configure it as BranchCache Host Server.

The main tasks for this exercise are as follows:

1. Install the BranchCache Feature on LON-SVR1.
2. Start the BranchCache Host Server.

#### ► Task 1: Install the BranchCache feature on LON-SVR1

1. Start **20417A-LON-SVR1**. After the computer starts, log on as **Adatum\Administrator** with the password of **Pa\$\$w0rd**.
2. Open Server Manager and add the **BranchCache for Network Files** role service.
3. Add the **BranchCache** feature.

#### ► Task 2: Start the BranchCache host server

1. On, LON-DC1, open Active Directory Users and Computers. Create a new OU called **BranchCacheHost** and move LON-SVR1 into this OU.
2. Open Group Policy Management and block GPO inheritance on the BranchCacheHost OU.
3. Switch to LON-SVR1 and restart the computer. Log on as **Adatum\Administrator** with the password of **Pa\$\$w0rd**
4. Open Windows PowerShell by clicking the icon on the taskbar and run the following cmdlets:

```
Enable-BCHostedServer -RegisterSCP
Get-BCStatus
```



**Note:** BranchCache is only available on Windows 8 Enterprise edition. This edition was not available when this course was created, so the BranchCache verification steps are not included in this lab.

**Results:** At the end of this exercise, you will have enabled the BranchCache server in the branch office.

#### ► To prepare for the next module

When you finish the lab, revert the virtual machines to their initial state. To do this, perform the following steps:

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20417A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20417A-LON-SVR1**, **20417A-LON-CL1**, and **20417A-LON-CL2**.

## Module Review and Takeaways

**Question:** How does BranchCache differ from DFS?

**Question:** Why would you want to implement BranchCache in Hosted Cache mode instead of the Distributed Cache mode?

**Question:** Is the storage spaces feature also available on Windows 8?

**Question:** Can you configure data deduplication on a boot volume?

### Tools

Tool	Use	Where to find it
iSCSI target server	Configure iSCSI targets	In Server Manager, under File and Storage Servers
iSCSI initiator	Configure a client to connect to an iSCSI target virtual disk	In Server Manager, in the Tools drop-down list
Deduplication Evaluation tool (DDPEval.exe)	Analyze a volume on the potential saving when enabling data deduplication	C:\windows\system32

# Module 5

## Implementing Network Services

### Contents:

Module Overview	5-1
Lesson 1: Implementing DNS and DHCP Enhancements	5-2
Lesson 2: Implementing IP Address Management	5-10
Lesson 3: NAP Overview	5-14
Lesson 4: Implementing NAP	5-20
Lab: Implementing Network Services	5-25
Module Review and Takeaways	5-31

## Module Overview

As seasoned administrators are aware, network services such as Domain Name System (DNS) provide critical support for name resolution of network and Internet resources. With Dynamic Host Configuration Protocol (DHCP) you can manage and distribute IP addresses to client computers. DHCP is essential in managing IP-based networks. DHCP failover can prevent client computers from losing access to the network if there is a DHCP server failure. IP Address Management provides a unified means of controlling IP addressing. With Network Access Protection (NAP), administrators can control which computers have access to corporate networks based on the computer's adherence to corporate security policies.

This module introduces DNS and DHCP improvements, what is new in IP address management, and describes how to implement these features. It also provides an overview and implementation guidance for NAP.

### Objectives

After completing this module, you will be able to:

- Implement DHCP and DNS enhancements.
- Implement IP address management.
- Describe NAP.
- Implement NAP.

## Lesson 1

# Implementing DNS and DHCP Enhancements

In TCP/IP networks of any size, certain services are required. DNS is one of the most important network services. Many other applications and services, including Active Directory® Domain Services (AD DS), rely on DNS to resolve resource names to IP addresses. Without DNS availability user authentications can fail, and network based resources and applications can become inaccessible. To prevent this, DNS has to be protected. Windows Server® 2012 implements DNS Security Extensions (DNSSEC) to protect the authenticity of DNS responses.

DHCP has long been used to ease the distribution of IP addresses to network client computers. Windows Server 2012 improves the functionality of DHCP by providing failover capabilities.

### Lesson Objectives

After completing this lesson, you will be able to:

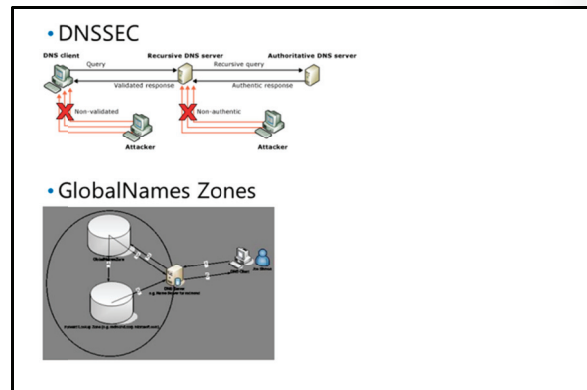
- Describe the new DNS features in Windows Server 2012.
- Configure DNSSEC.
- Describe the new DHCP features in Windows Server 2012.
- Configure failover for DHCP.

### What's New in DNS in Windows Server 2012

DNSSEC and Global Name Zones are two features that continue to be available in Windows Server 2012. However, the DNSSEC implementation has been simplified in Windows Server 2012.

#### DNSSEC

Intercepting and tampering with an organization's DNS query response is a common attack method. If an attacker can alter the response from a DNS server, or send a spoofed response to point client computers to their own servers, they can gain access to sensitive information. This is known as a man-in-the-middle attack. Any service that relies on DNS for the initial connection, such as e-commerce web servers and email servers are vulnerable. DNSSEC is intended to protect clients that are making DNS queries from accepting false DNS responses.



#### New Resource Records

Validation of DNS responses is achieved by associating a private/public key pair (generated by the administrator) with a DNS zone and defining additional DNS resource records to sign and publish keys. Resource records distribute the public key while the private key remains on the server. When the client requests validation, DNSSEC adds data to the response that enables the client to authenticate the response.

Windows Server 2012 defines the new resource records in the following table.

Resource Record	Purpose
DNSKEY	This record publishes the public key for the zone. It checks the authority of a response against the private key held by the DNS server. These keys require periodic replacement. This is known as key rollovers. Windows Server 2012 supports automated key rollovers.
DS	This is a delegation record that contains the hash of the public key of a child zone. This record is signed by the parent zone's private key. If a child zone of a signed parent is also signed, the DS records from the child must be manually added to the parent so a chain of trust can be created.
RRSIG	This record holds a signature for a set of DNS records. It is used to check the authority of a response.
NSEC	When the DNS response has no data to provide to the client this record authenticates that the host does not exist.

### Trust Anchors

A trust anchor is an authoritative entity represented by a public key. The TrustAnchors zone stores preconfigured public keys that are associated with a specific zone. In DNS the trust anchor is the DNSKEY or DS resource record. Client computers use these records to build trust chains. A trust anchor from the zone must be configured on every domain DNS server in order to validate responses from that signed zone. If the DNS server is a domain controller then Active Directory integrated zones can distribute the trust anchors.

### Name Resolution Policy Table (NRPT)

The NRPT contains rules that control the DNS client behavior for sending DNS queries and processing the responses from those queries. For example, a DNSSEC rule prompts the client computer to check for validation of the response for a particular DNS domain suffix. Group policy is the preferred method of configuring the NRPT. If there is no NRPT present the client computer does not validate responses.

### Considerations when implementing DNSSEC

Consider the following before you implement DNSSEC:

- The zone replication scope or type cannot be changed while a zone is signed.
- DNS response messages are larger.
- DNS traffic increases are caused by queries for DNSKEY records.
- Zone files are larger.
- The client computer has to spend more time authenticating responses.
- There is an added level of administration to maintain.

### GlobalNames Zones

GlobalNames zones address a problem in multiple DNS domain environments. GlobalName zones are used when you must maintain a list of DNS search suffixes on client computers to resolve names among these multiple DNS domains. For example, if an organization supports two DNS domains, such as Widgets.com and Corp.com, users in the Widgets.com DNS domain have to use the fully qualified domain name (FQDN) to locate the servers in corp or the domain administrator has to add a DNS search suffix for Corp.com on all the systems in the Widgets.com domain. In other words, if users in the Widgets.com

domain want to locate a server named Data in the Corp.com domain, they would have to search for the FQDN of Data.Corp.com to locate that server. If they just search for the server name Data, then the search would fail.

Global names are based on creating Canonical Name (CNAME) records (or aliases) in a special forward lookup zone that use single names to point to FQDNs. GlobalNames zones enables clients in any DNS domain to use a single label name, such as Data, to locate a server whose FQDN is Data.corp.com without having to use the FQDN.

### Creating GlobalNames Zones

To create GlobalNames zones:

- Use the Dnscmd utility to enable GlobalNames zones functionality.
- Create a new forward lookup zone named GlobalNames (not case-sensitive). Do not enable dynamic updates for this zone.
- Manually create CNAME records that point to records that already exist in the other zones hosted on your DNS servers.

For example, you could create a CNAME record in the GlobalNames zone for Data that points to Data.corp.com. This enables clients from any DNS domain in the organization to find this server by the single label name of Data.

### How to Configure DNSSEC

Although DNSSEC was supported in Windows Server 2008 R2, most of the configurations and administration were performed manually, and zones were signed when they were offline. Windows Server 2012 includes a DNSSEC wizard to simplify the configuration and signing process, and enables online signing.

- DNSSEC is simpler to deploy in Windows Server 2012 than in previous versions of Windows Server.
- To Deploy DNSSEC:
  - Assign the DNS server role
  - Sign the zones
  - Configure trust anchor distribution points
  - Configure NRPT on clients

### Deploying DNSSEC

To deploy DNSSEC:

1. Install Windows Server 2012 in the environment and assign the server the DNS role. Typically a domain controller also acts as the DNS server. However, that is not a requirement.
2. Sign the DNS zone by using the DNSSEC configuration wizard in the DNS Manager console.
3. Configure trust anchor distribution points.
4. Configure the NRPT on the client computers.

### Assign the DNS Server Role

To add the DNS server role, from the Server Manager Dashboard, use the Add Roles and Features Wizard. You can also add this role can when you add the AD DS role. Configure the primary zones on the DNS server. After a zone is signed, any new DNS servers on Windows Server 2012 automatically receives the DNSSEC parameters.

## Sign the Zone

To access the DNSSEC zone signing wizard, right-click the primary zone. You can sign zones on any Windows Server 2012 that hosts a primary DNS zone. You cannot configure DNSSEC on secondary zones. The wizard guides you through all the configuration steps required to sign the zone.

The following signing options are available:

- The **Configure the zone signing parameters** option guides you through the steps and enables you to set all values for the Key Signing Key (KSK) and the Zone Signing Key (ZSK).
- The **Sign the zone with parameters of an existing zone** option enables you to keep the same values and options as another signed zone.
- The **Use recommended settings** option signs the zone by using the default values.



**Note:** Zones can also be unsigned by using the DNSSEC management user interface.

## Configure Trust Anchor Distribution Points

If the zone is Active Directory Integrated, you should select to distribute the trust anchors to all the servers in the forest. If trust anchors are required on computers that are not joined to the domain, for example, a DNS server in the perimeter network (also known as DMZ, demilitarized zone, and screened subnet), then you should enable automated key rollover.

## Configure NRPT on Client Computers

The DNS client computer only performs DNSSEC validation on domain names where it is configured to do so by the NRPT. A client computer running Windows® 7 is DNSSEC aware, but does not perform validation. It relies on the security aware DNS server to perform validation on its behalf.

## Demonstration: Configuring DNSSEC

In this demo you will see how to use the wizard in the DNS management console to configure DNSSEC.

### Demonstration Steps

1. Log on to LON-DC1 as Adatum\Administrator.
2. Start the DNS Management console.
3. Use the DNSSEC zone signing wizard to sign the Adatum.com zone. Accept all the default settings.
4. Verify the DNSKEY resource records were created in the Trust Points zone.
5. Use the Group Policy Management Console to configure NRPT. Create a rule that enables DNSSEC for the Adatum.com suffix and requires DNS client computers to check that the name and address data is validated.
6. Close all open Windows.

## What's New in DHCP in Windows Server 2012

DHCP failover is a new feature for Windows Server 2012. It addresses the issue of client computers losing connectivity to the network and all its resources if there is DHCP server failure.

Another new feature in Windows Server 2012 is DHCP name protection. Names that are registered in DNS by DHCP on behalf of systems must be protected from being overwritten by non-Microsoft systems that have the same name. For example, a Unix based system named Client1 could potentially overwrite the DNS address that was assigned and registered by DHCP on behalf of a Windows-based system also named Client1. DHCP name protection addresses this issue.

- DNCP name protection can be configured in properties at the IP level or scope level

DHCP Limitations	WS 2012 solution
Failure of DHCP will result in loss of network connectivity for clients	DHCP failover
Windows systems can have their DNS name registrations overwritten by non-Microsoft systems bearing the same system name	DHCP name protection

### DHCP Failover

DHCP client computers renew their lease on their IP address at regular, configurable intervals. If the DHCP server service fails, then leases time-out, and eventually client computers no longer have IP addresses. In the past, DHCP failover was not possible because DHCP servers were independent and unaware of one another. Configuring two separate DHCP servers to distribute IP addresses within the same pool could lead to duplicate address assignment if the administrator incorrectly configured overlapping ranges. The DHCP server failover feature enables an alternative DHCP server to distribute IP addresses and associated option configuration to the same subnet or scope. Lease information is replicated between the two DHCP servers. If one of the DHCP servers fails, then the other DHCP server services the client computers for the whole subnet. In Windows Server 2012 you can configure one alternative DHCP server for failover. Additionally, only IPv4 scopes and subnets are supported because IPv6 uses a different IP address assignment scheme.



**Note:** For more information about DHCP options in IPv6, see: <http://technet.microsoft.com/en-us/library/cc753493>.

### DHCP Name Protection

"Name squatting" describes the problem where a DHCP client computer registers a name with DNS, but that name is actively being used by another computer. The original computer then becomes inaccessible. This problem typically occurs between non-Windows systems that have duplicate names of Windows systems. DHCP Name Protection uses a resource record known as a DHCID to keep track of which computer originally requested the name. This record is provided by the DHCP server and stored in DNS. When the DHCP server receives a request to update a host record that is currently associated with a different computer, the DHCP server can verify the DHCID in DNS to check whether the requester is the original owner of the name. If it is not the same computer, the record in DNS is not updated. To resolve this issue, either the current host name owner must release the IP address, or the requester must use a different host name. You can implement name protection for both IPv4 and IPv6. Configuration is set in the properties page at the IP address level or the scope level.



## How to Configure Failover for DHCP

To configure failover of DHCP you must establish a failover relationship between the two servers. You must give this relationship a unique name. This name is exchanged with the failover partner during the configuration. This enables a single DHCP server to have multiple failover relationships with other DHCP servers, as long as they all have unique names. Failover is configured through a wizard that you can start on the shortcut menu of the IP node or the scope node.

- Failover relationships must have unique names
- The MCLT determines when a failover partner takes control of the subnet or scope
- Failover supports two modes:
  - Hot Standby Mode
  - Load Sharing Mode
- Auto State Switchover Interval determines when a failover partner is considered to be down
- Message authentication can validate the failover messages
- Firewall rules are auto-configured during DHCP installation



**Note:** DHCP failover is time sensitive. Time must be kept synchronized between the partners in the relationship. If the time difference is greater than one minute the failover process will stop with a critical error.

### Configure Maximum Client Lead Time

The administrator configures the Maximum Client Lead Time (MCLT) parameter to determine the time that a DHCP server waits if the partner is unavailable before assuming control of the whole address range. This value cannot be zero and the default is one hour.

### Configure Failover Mode

Failover can be configured in one of two modes:

Mode	Characteristics
Hot Standby Mode	In this mode one server is the primary server and the other is a secondary. The primary server actively distributes IP configurations for the scope or subnet. The other DHCP server will only take over this role if the primary server becomes unavailable. A DHCP server can act as the primary for one scope or subnet while it is the secondary for another. Administrators must configure a percentage of the scope addresses to be assigned to the standby server. These addresses are distributed during the MCLT interval if the primary server is down. The default value is 5 percent of the scope. The secondary takes control of the whole range after the MCLT has passed. Hot Standby mode is best suited to deployments where a data recovery (DR) site is located at a different location. Then, the DHCP server does not service client computers unless there is an outage of the main server.
Load Sharing Mode	This is the default mode. In this mode both servers concurrently distribute IP configuration to client computers. Which server responds to IP configuration requests depends on how the administrator configures the load distribution ratio. The default ratio is 50:50.

### Configure Auto State Switchover Interval

When a server loses contact with its partner it goes into a communication interrupted state. Because the server cannot determine what is causing the communication loss, it stays in this state until the administrator manually changes it to a partner down state. The administrator can also enable automatic transition to partner down state by configuring the auto state switchover interval. The default value for this interval is 10 minutes.

### Configure Message Authentication

Windows Server 2012 enables you to authenticate the failover message traffic between the replication partners. The administrator can establish a shared secret, much like a password, in the configuration wizard for DHCP failover. This validates that the failover message comes from the failover partner.

### Firewall Considerations

DHCP uses TCP port 647 to listen for failover traffic. The DHCP installation creates the following incoming and outgoing firewall rules:

- Microsoft-Windows-DHCP-Failover-TCP-In
- Microsoft-Windows-DHCP-Failover-TCP-Out

### Configure DHCP Failover

The Configuration Failover Wizard steps you through the process of creating a failover relationship. The wizard prompts you to enter the following information:

- Name of the relationship
- Which scopes are selected for failover
- Name of the partner server
- The MCLT
- The Mode
- The Load Balance Percentage
- The Auto State Switchover Interval
- Message Authentication setting
- A shared secret

The failover relationship can then be modified as required through the Failover tab in the properties of IPv4.

## Demonstration: Configuring Failover for DHCP

In the demonstration you will see how to use the DHCP console to configure DHCP failover in load sharing mode.

### Demonstration Steps

1. Log on to LON-SVR1 as the Adatum\administrator.
2. Start the DHCP console and view the current state of DHCP. Note the server is authorized but no scopes are configured.
3. Switch to LON-DC1.
4. Open the DHCP Management console and start the Configure Failover Wizard.
5. Configure failover replication with the following settings:
  - Partner server = 172.16.0.21
  - Relationship Name = Adatum
  - Maximum Client Lead Time = 15 minutes
  - Mode = Load balance
  - Load Balance Percentage = 50%
  - State Switchover Interval = 60 minutes
  - Message authentication shared secret: Pa\$\$w0rd
6. Complete the wizard.

## Lesson 2

# Implementing IP Address Management

With the development of IPv6 and more and more devices requiring IP addresses, networks have become very complex and difficult to manage. Windows Server 2012 has implemented IP Address Management (IPAM) as a tool to manage IP addresses.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe IPAM.
- Describe the IPAM architecture.
- Describe the requirements for IPAM.

### What is IP Address Management?

IP management is difficult in large networks because tracking IP address usage is largely a manual operation. IPAM is a framework for discovering, utilization monitoring, auditing, and managing the IP address space in a network. IPAM enables the administration and monitoring of DHCP and DNS. IPAM provides a comprehensive view of where IP addresses are used. IPAM collects information from domain controllers and Network Policy Servers (NPS) and stores that information in the Windows Internal Database.

IPAM assists in the areas of IP administration shown in the following table.

- IPAM assists in the following areas of IP address management:
  - Planning
  - Managing
  - Tracking
  - Auditing
- IPAM provides multiple benefits for IP administrators

IP Administration Area	IPAM Capabilities
Planning	Provides a tool set that can reduce the time and expense of the planning process when changes occur in the network.
Managing	Provides a single point of management and assists in optimizing utilization and capacity planning for DHCP and DNS.
Tracking	Enables tracking and forecasting of IP address utilization.
Auditing	Assists with compliance requirements, such as HIPAA and Sarbanes-Oxley, and provides reporting for forensics and change management.

### Benefits of IPAM

IPAM benefits include:

- IPv4 and IPv6 address space planning and allocation.
- IP address space utilization statistics and trend monitoring.
- Static IP inventory management, lifetime management and DHCP and DNS record creation and deletion.

- Service and zone monitoring of DNS services.
- IP address lease and logon event tracking.
- Role-based access control.
- Remote administration support through Remote Server Administration Tools (RSAT).



**Note:** IPAM does not support management and configuration of non-Microsoft network elements.

## IPAM Architecture

IPAM consists of four main modules, as shown in the following table:

- IPAM has four main modules:
  - IPAM discovery
  - IP address space management
  - Multi-server management and monitoring
  - Operational auditing and IP address tracking
- IPAM can be deployed in three topologies:
  - Distributed
  - Centralized
  - Hybrid
- IPAM has two components:
  - IPAM Server
  - IPAM Client

Module	Description
IPAM discovery	You use Active Directory to discover servers running Windows Server 2008 and later versions that have DNS, DHCP, or AD DS installed. Administrators can define the scope of discovery to a subset of domains in the forest. They can also manually add servers.
IP address space management (ASM)	You can use this module to view, monitor and manage the IP address space. You can dynamically issue or statically assign addresses. You can also track address utilization and detect overlapping DHCP scopes.
Multi-server management and monitoring	You can manage and monitor multiple DHCP servers. This enables tasks to be executed across multiple servers. For example, you can configure and edit DHCP properties and scopes and track the status of DHCP and scope utilization. You can also monitor Multiple DNS servers, and monitor the health and status of DNS zones across authoritative DNS servers.
Operational auditing and IP address tracking	You can track use the auditing tools to track potential configuration problems. You can also collect, manage, and view details of configuration changes from managed DHCP servers. You can also collect address lease tracking from DHCP lease logs, and collect logon event information from Network Policy Servers (NPS) and domain controllers.

The IPAM server can only manage one Active Directory forest. IPAM is deployed in one of three topologies:

- **Distributed** – An IPAM server is deployed to every site in the forest.
- **Centralized** – Only one IPAM server is deployed in the forest.
- **Hybrid** – A central IPAM server is deployed together with a dedicated IPAM server in each site.



**Note:** IPAM servers do not communicate with one another or share database information. If you deploy multiple IPAM servers, you must customize the discovery scope of each server.

IPAM has two main components:

- **IPAM Server** – performs the data collection from the managed servers. It also manages the Windows Internal Database and provides role based access control.
- **IPAM Client** – provides the client computer user interface and interacts with the IPAM server and invokes PowerShell to perform DHCP configuration tasks, DNS monitoring and remote management.

## Requirements for IPAM Implementation

You must meet several prerequisites to ensure a successful IPAM deployment:

- The IPAM server must be a domain member, but cannot be a domain controller.
- The IPAM server should be a single purpose server. Do not install other network roles such as DHCP or DNS on the same server.
- To manage the IPv6 address, space IPv6 must be enabled on the IPAM server.
- Log on to the IPAM server with a domain account, not a local account.
- You must be a member of the correct IPAM local security group on the IPAM server.
- Ensure that logging of account logon events is enabled on DC and NPS servers for the IP Address Tracking and auditing feature of IPAM.

Hardware and software requirements:

- Dual core processor of 2.0 GHZ or higher
- Windows Server 2012 operating system
- 4 GB of RAM or more
- 80 GB of free hard disk space

- IPAM requirements:
  - IPAM server must belong to the domain
  - IPAM server cannot be a domain controller
  - IPv6 must be enabled to manage IPv6
  - Log on with a domain account
  - You must be in the correct IP security group
  - Logging account logon events must be enabled for IP address tracking and auditing
- Hardware and software:
  - CPU – dual core 2.0 GHZ or higher
  - Windows Server 2012 Operating system
  - 4 GB of RAM
  - 80 GB free disk space

## Demonstration: Implementing IPAM

In this demonstration you will see how to install IPAM. You will also see how to create the related GPOs and begin server discovery.

### Demonstration Steps

1. Log on to LON-SVR1 as Adatum\Administrator.
2. In Server Manager add the IPAM feature and all required supporting features.
3. From the IPAM Overview pane provision the IPAM server by using Group Policy.
4. Enter IPAM as the GPO name prefix and provision IPAM.
5. From the IPAM Overview pane configure server discovery for the Adatum domain.
6. From the IPAM Overview pane start the server discovery process.

## Lesson 3

# NAP Overview

NAP is a policy-enforcement platform that is built into the Windows XP with Service Pack 3 (SP3) and later operating systems, and into Windows Server 2008 and later operating systems. NAP enables you to protect network assets by enforcing compliance with system-health requirements. NAP provides the necessary software components to help ensure that computers that are connected or connecting to the network remain manageable so that they do not become a security risk to the network and other attached computers.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe NAP.
- Describe NAP architecture.
- Describe scenarios for using NAP.
- Describe the considerations for using NAP.

### What is NAP?

NAP enforces client computer health before it enables client computers to access the network. Client health can be based on characteristics such as antivirus software status, Windows Firewall status, or the installation of security updates. The monitored characteristics are based on which system health agents are installed.

NAP enables you to create solutions for validating computers that connect to your networks, in addition to providing needed updates or access to needed health update resources, and limiting the access or communication of noncompliant computers.

You can integrate NAP's enforcement features with software from other vendors or with custom programs. You can customize the health-maintenance solution that developers within your organization might develop and deploy, whether for monitoring the computers accessing the network for health policy compliance, automatically updating computers with software updates to meet health policy requirements, or limiting the access to a restricted network of computers that do not meet health policy requirements.

NAP does not protect a network from malicious users. Instead, it enables you maintain the health of your organization's networked computers automatically, which in turn helps maintain the network's overall integrity. For example, if a computer has all the software and configuration settings that the health policy requires, the computer is compliant and has unlimited network access. NAP does not prevent an authorized user who has a compliant computer from uploading a malicious program to the network or engaging in other unsuitable behavior.

Also, unless configured specifically, NAP cannot determine whether a client computer is free of viruses, trojans, rootkits or malware. Default behavior is to check for compliance in having current antivirus software and configurations.

- Network Access Protection can:
  - Enforce health-requirement policies on client computers
  - Ensure client computers are compliant with policies
  - Offer remediation support for computers that do not meet health requirements
- Network Access Protection cannot:
  - Protect the network from malicious users
  - Guarantee that a client computer is not infected



## Features of NAP

NAP has three important and distinct features:

- **Health state validation:** When a client computer tries to connect to the network, NAP validates the computer's health state against the health-requirement policies that the administrator defines. You can also define what to do if a computer is not compliant. In a monitoring-only environment, all computers have their health state evaluated and the compliance state of each computer is logged for analysis. In a limited access environment, computers that comply with the health-requirement policies have unlimited network access. Computers that do not comply with health-requirement policies could find their access limited to a restricted network.
- **Health policy compliance:** You can help ensure compliance with health-requirement policies by choosing to update noncompliant computers automatically with missing software updates or configuration changes through management software, such as Microsoft System Center Configuration Manager. In a monitoring-only environment, computers have network access before they are updated with required updates or configuration changes. In a limited access environment, noncompliant computers have limited access until the updates and configuration changes are completed. In both environments, computers that are compatible with NAP can become compliant automatically and you can define exceptions for computers that are not NAP compatible.
- **Limited Access:** You can protect your networks by limiting the access of noncompliant computers. You can base limited network access on a specific time, or on the resources that the noncompliant computer can access. In the latter case, you define a restricted network that contains health update resources, and the limited access lasts until the noncompliant computer comes into compliance. You can also configure exceptions so that computers that are incompatible with NAP do not have limited network access.

## What's New for NAP in Windows Server 2012

### Support for Windows PowerShell

You can now use Windows PowerShell® to automate the installation of the Network Policy and Access Services server role. You can also use Windows PowerShell to deploy and configure some aspects of Network Policy Server.

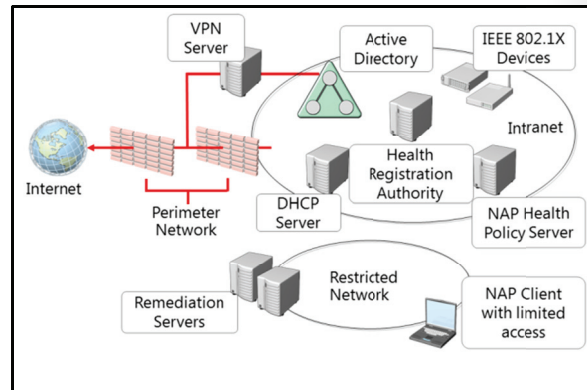
### Removed Functionality

In Windows Server 2008 R2 and Windows Server 2008, Network Policy and Access Services included the Routing and Remote Access Service role service. In Windows Server 2012, RRAS is now a role service in the Remote Access server role

- Support for Windows PowerShell
- RRAS is now a role service in the Remote Access server role

## NAP Architecture

The following table describes the NAP components.



Components	Description
NAP Clients	<p>Computers that support the NAP platform for system health-validated network access or communication. Client architecture consists of:</p> <ul style="list-style-type: none"> <li>• <b>NAP enforcement client (EC):</b> ECs monitor attempts to connect to the network. Different EC components exist for different types of network access.</li> <li>• <b>System health agents (SHA):</b> SHAs report on one or more elements of system health. For example, there might be an SHA for checking antivirus definitions and another for checking Windows updates. The SHA returns a statement of health (SoH) to the NAP agent which passes that to the NAP health policy server for evaluation.</li> <li>• <b>NAP agent:</b> Collects and stores SoHs from the SHAs and supplies it to the ECs when requested.</li> </ul>
NAP enforcement points	<p>NAP enforcement points are computers or network-access devices that use NAP to evaluate a NAP client computer's health state. NAP enforcement points rely on policies from a Network Policy Server (NPS) to perform that evaluation and determine whether network access or communication is enabled, and the set of remediation actions that a noncompliant NAP client computer must perform.</p> <p>NAP enforcement points can include:</p> <ul style="list-style-type: none"> <li>• <b>Health Registration Authority (HRA)</b> is a server running Windows Server 2012 with Internet Information Services (IIS) installed that obtains health certificates from a certification authority (CA) for compliant computers.</li> <li>• <b>VPN server</b> is a Windows 2012 server that runs Routing and Remote Access, and that enables remote access VPN intranet connections through remote access.</li> <li>• <b>DHCP server</b> is a Windows 2012 server that runs the DHCP Server service.</li> <li>• <b>Network access devices</b> are Ethernet switches or wireless access points that support IEEE 802.1X authentication.</li> </ul>

Components	Description
NAP health policy servers	<p>Windows 2012 servers run the NPS service and store health-requirement policies and provide health-state validation for NAP. NPS replaces the Internet Authentication Service (IAS), and the Remote Authentication Dial-In User Service (RADIUS) server and proxy that Windows Server 2003 provides. The NAP health policy server has the following components:</p> <ul style="list-style-type: none"> <li>• <b>NPS service:</b> Receives RADIUS requests and extracts the System State of Health (SSoH) and passes it to the NAP administration server component.</li> <li>• <b>NAP Administration Server:</b> Makes Communication Easier between the NPS service and the SHVs.</li> <li>• <b>System Health Validators (SHV):</b> You define SHVs for system health elements and match them to an SHA. An example of these would be a SHV for an antivirus software that tracks the latest version of the antivirus definition file.</li> </ul> <p>NPS also acts as an authentication, authorization, and accounting (AAA) server for network access. When acting as an AAA server or NAP health policy server, NPS typically runs on a separate server for centralized configuration of network access and health-requirement policies. The NPS service also runs on Windows Server 2012-based NAP enforcement points that do not have a built-in RADIUS client computer, such as an HRA or DHCP server. However, in these configurations, the NPS service acts as a RADIUS proxy to exchange RADIUS messages with a NAP health policy server.</p>
AD DS	AD DS stores account credentials and properties, and stores Group Policy settings. Although not required for health-state validation, Active Directory is required for IPSec-protected communications, 802.1X-authenticated connections, and remote access VPN connections.
Restricted network	<p>This is a separate logical or physical network that has the following components:</p> <ul style="list-style-type: none"> <li>• Remediation servers that contain health update resources, such as antivirus definition distribution points and Windows software update servers, which NAP client computers can access to remedy their noncompliant state.</li> <li>• NAP client computers that have limited access are added on the restricted network when they do not comply with health-requirement policies.</li> </ul>

## Scenarios for Using NAP

NAP provides a solution for the common scenarios described in this section. Depending on your needs, you can configure a solution to address any of these scenarios for your network.

### Roaming Portable computers

Portability and flexibility are two primary portable computer advantages, but these features also present a system health threat. Users frequently connect their portable computers to other networks. When users are away from your organization, their portable computers might not receive the most recent software updates or

- Roaming laptops
- Desktop computers
- Visiting laptops
- Unmanaged home computers

configuration changes. Additionally, exposure to unprotected networks, such as the Internet, could introduce security-related threats to the portable computers. NAP lets you check any portable computer's health state when it reconnects to the organization's network, whether through a VPN, DirectAccess connection, or the workplace network connection.

### Desktop Computers

Although desktop computers are usually not taken out of the company building, they still can present a threat to the network. To minimize this threat, you must maintain these computers with the most recent updates and required software. Otherwise, these computers are at risk of infection from websites, email, files from shared folders, and other publicly available resources. NAP enables you to automate health state checks to verify each desktop computer's compliance with health-requirement policies. You can check log files to determine which computers do not comply. Additionally, by using management software enables you to generate automatic reports and automatically update noncompliant computers. When you change health-requirement policies, computers can be provisioned automatically with the most recent updates.

### Visiting Portable Computers

Organizations frequently have to enable consultants, business partners, and guests to connect to their private networks. The portable computers that these visitors bring into your organization might not meet system health requirements and can present health risks. NAP enables you to determine which visiting portable computers are noncompliant and limit their access to restricted networks. Typically, you would not require or provide any updates or configuration changes for visiting portable computers. You can configure Internet access for visiting portable computers, but not for other organizational computers that have limited access.

### Unmanaged Home Computers

Unmanaged home computers that are not a member of the company's Active Directory domain can connect to a managed company network through VPN. Unmanaged home computers provide an additional challenge because you cannot physically access these computers. Lack of physical access makes enforcing compliance with health requirements—such as the use of antivirus software—more difficult.

However, NAP enables you to verify the health state of a home computer every time that it makes a VPN connection to the company network, and to limit its access to a restricted network until it meets system health requirements.

## Considerations for NAP

Before you implement NAP, you must consider the following points.

### Considerations for NAP Client Computer Deployment

Before you can use NAP on client computers, you must configure the NAP settings. Although you can use the Netsh commands to configure all aspects of the NAP client computer, Group Policy is the preferred method of deploying client computer settings. The NAP Client Configuration console and NAP client computer configuration settings in the Group Policy Management Console provide a graphical user interface for configuring NAP client computer settings.

- Use group policy to deploy client settings
- Plan the enforcement type you wish to enforce
- Plan for a remediation network
- Ensure you can provide the administrative support for the solution

### Consideration for a NAP Enforcement Type

Deciding on the best enforcement type for your organization is very important.

NAP provides four mechanisms:

- **VPN:** The VPN server relays the policy from the Network Policy Server (NPS) to the requesting client computer and performs the validation. This method requires a computer certificate to perform PEAP-based user or computer authentication.
- **DHCP:** The DHCP server interacts with the policies from the NPS to determine the client computer's compliance.
- **IPsec:** enforces the policy and configures the systems out of compliance with a limited access local IP security policy for remediation. This method requires a computer certificate to perform PEAP-based user or computer authentication.
- **802.1X:** authenticates over an 802.1X authenticated network and is the best solution when integrating hardware from other vendors.

### Considerations for a Remediation Network

You can provide a remediation network as a location for client computers that are out of compliance to resolve issues and then gain access to the network. It is important to make the remediation network a place where client computers can gain the required updates or definitions without help desk intervention.

### Administrative Effort and Support

NAP is not a simple solution to implement and requires a good level of understanding and ongoing support.

## Lesson 4

# Implementing NAP

There are different NAP procedures, depending on the type of enforcement you are implementing. This lesson describes the main requirements for each of the NAP enforcement methods.

### Lesson Objectives

After completing this lesson you will be able to:

- Describe the requirements for implementing NAP.
- Describe the requirements for NAP with VPN.
- Describe the requirements for NAP with IPsec.
- Describe the requirements for NAP DHCP.
- Describe the requirements for NAP with 802.1X.

### Requirements for Implementing NAP

All NAP enforcement methods require that the NAP Agent service is running on the client computer and that at least one enforcement client computer is enabled. Depending on the desired enforcement method there may be other services and settings required.

A Network Policy Server (NPS) is required to create and enforce organization-wide network access policies for client computer health, connection request authentication and authorization. The NPS can also act as a RADIUS server. The NPS evaluates the statements of health (SoH) sent by NAP client computers.

- All enforcement methods require the NAP agent to run on the client
- A Network Policy Server (NPS) is required to create and enforce policies
- SHVs are required to determine what will be evaluated on the client
- System health policies are required to determine client compliance or noncompliance
- Certificates are required to validate computer identities for PEAP authentication
- Remediation networks can provide a way for clients to become compliant and gain access to the network

System Health Validators (SHVs) are required to determine what the system health policy checks for. SHVs can check for Windows Firewall settings, antivirus and spyware protection, Windows Updates, and so on.

Health policies compare the state of a client computer's health according to SHVs that are defined by corporate requirements and determine whether the client computer is compliant or noncompliant with the corporate policy. A health policy can be defined to check one of the following:

- Client passes all SHV checks
- Client fails all SHV checks
- Client passes one or more SHV checks
- Client fails one or more SHV checks

Network policies are required to determine what happens if the client computer requesting network access is compliant or noncompliant. These policies determine what level of access, if any, the client computer will receive to the network.

A certification authority (CA) is required to issue computer certificates to validate computer identity if Protected EAP (PEAP) is used for authentication. This may be an enterprise CA or a third-party CA.

Remediation networks are not an absolute requirement, but can provide a means for a client computer to become compliant. For example, a network policy can direct a noncompliant client computer to a network segment that contains a Web site from which the client computer can obtain current virus definitions or Windows Updates.

## NAP with VPN

NAP enforcement for VPN method works by using a set of remote access IP packet filters to limit the traffic of a noncompliant VPN client computer so that it can only reach the resources on the restricted network. Compliant client computers will be granted full access. VPN servers can enforce the health policy for computers that are considered to be noncompliant by applying the filters.



**Note:** Site-to-site VPN connections do not support NAP health evaluation.

- The VPN server uses the NPS server as primary RADIUS
- VPN servers are configured as RADIUS clients in NPS
- Connection request policy has the VPN server as source
- Configure SHVs to test for health conditions
- Health policies pass compliant clients and fail noncompliant clients
- Network policy grants full access to compliant clients and limited access to noncompliant clients
- Group policy or local policy can enable the ECs on client computers
- NAP agent service must be enabled on clients
- Computer certificates are required for PEAP authentication

To deploy NAP with VPN you must:

- Install RRAS as a VPN server and configure the NPS as the primary RADIUS server.
- Configure the VPN servers as RADIUS client computers in the NPS.
- Configure a connection request policy with the source set to the VPN server.
- Configure SHVs to test for health conditions.
- Create compliant health policies to pass selected SHVs and a noncompliant health policy to fail selected SHVs.
- Configure a network policy with the source set to the VPN server. Full access will be granted to compliant computers and limited access to noncompliant computers.
- Enable the NAP Remote Access and EAP enforcement clients on client computers. You can do this by using Group Policy or local policy settings.
- Enable the NAP agent service on client computers.
- Issue computer certificates to use PEAP authentication.



## NAP with IPsec

NAP IP security (IPsec) enforcement provides the strongest and most flexible method for maintaining client computer compliance with network health requirements.

To implement NAP with IPsec you must:

- Configure a certification authority (CA) to issue health certificates: the System Health Authentication template must be issued and the HRA must be granted permission to enroll the certificate.
- Install Health Registration Authority (HRA): the HRA is a component of NAP that is central to IPsec enforcement. The HRA obtains health certificates on behalf of NAP client computers when they are compliant with network health requirements. These health certificates authenticate NAP client computers for IPsec-protected communications with other NAP client computers on an intranet. If a NAP client computer does not have a health certificate, the IPsec peer authentication fails.
- Select authentication requirements: the HRA can provide health certificate to authenticated domain users only, or optionally provide health certificates to anonymous users.
- Configure the NPS server with the required health policies.
- Configure NAP client computers for IPsec NAP enforcement: NAP agent must be running and the NAP IPsec EC must be running. You can do this through Group Policy or local policy or Netsh commands.
- Use IPsec policies to create logical networks: IPsec enforcement divides a physical network into three logical networks. A computer is a member of only one logical network at any time. The logical networks are:
  - Secure network - Computers on the secure network have health certificates and require that incoming communication is authenticated by using these certificates.
  - Boundary network - Computers on the boundary network have health certificates, but do not require IPsec authentication of incoming communication attempts.
  - Restricted network - Computers on the restricted network do not have health certificates.

NAP with IPsec requires:

- A CA to issue health certificates
- An HRA to authenticate and obtain health certificate on behalf of clients
- Authentication requirements: domain only or anonymous
- An NPS server
- Clients configured for IPsec enforcement
- IPsec policies to create logical networks

## NAP with DHCP

NAP enforcement can be integrated with DHCP so that NAP policies can be enforced when a client computer tries to lease or renew its DHCP address. The NPS server uses health policies and SHVs to evaluate client computer health. Based on the evaluation the NPS tells the DHCP server to provide full access to compliant computers and to restrict access to noncompliant computers.

- NAP enforcement can be integrated with DHCP
- NPS server uses health policies and SHVs to evaluate client health
- NPS tells the DHCP server to provide full access to compliant computers and to restrict access to noncompliant computers



The components listed in the following table must be defined on the NPS.

Component	Description
Radius client computers	If DHCP is installed on a separate computer, the NAP DHCP server must be configured as a RADIUS client computer in NPS. You must also select RADIUS client computer is NAP-capable.
Network policy	Source must be set to DHCP server. Both compliant and noncompliant policies are set to grant access.
Connection request policy	Source is set to DHCP server. The policy authenticates requests on this server.
Health policies	Must be configured to pass SHVs in the compliant policy and fail SHVs in the noncompliant policy.
SHVs	Health checks are configured on the NPS server.
NAP agent	Must be running on the client computer.
IP address configuration	Must be configured to use DHCP. Clients that have static IP address cannot be evaluated.

## Demonstration: Implementing NAP with DHCP

Because you are configuring NPS on the DHCP server you do not have to designate the DHCP server as a RADIUS client computer.

You will configure the policy for all scopes.

### Demonstration Steps

1. Install **Network Policy and Access Services** on LON-DC1.
2. Use the Configure NAP Wizard to create a DHCP enforcement policy.
3. Configure DHCP to enable Network Access Protection for all scopes.

## Network Access Protection with 802.1X

You can provide NAP enforcement to an IEEE 802.1X-capable device, such as a wireless access point, authenticating switch, or other network device. NAP enforcement occurs when client computers try to access the network through these devices.

NAP with 802.1x has the following characteristics:

- Radius client computers must be added in the NPS console and are identified by host name or IP address.
- A shared secret must be configured in the NPS server and the device to identify the radius client computer.

NAP with 802.1x characteristics:

- Radius clients are identified by host name or IP address
- Shared secrets must be configured
- Server certificates must be installed
- Network authentication must use EAP methods
- VLANs may be configured
- Type of network access server should be set to unspecified
- Connection requests must use PEAP

- Server certificates must be installed and client computers must trust these certificates.
- Network authentication must use EAP authentication methods – secure passwords, smart cards or other certificates.
- If your access points support VLANs, you can configure that information for NPS. For example, the restricted network may be a VLAN.
- When you create network policies and connection request policies, the type of network access server should be set to Unspecified.
- Connection request policies must be configured to use PEAP authentication in the policy.

## Lab: Implementing Network Services

### Scenario

A. Datum has grown quickly over the last few years in several ways. The company has deployed several new branch offices, it has significantly increased the number of users in the organization, and it has expanded the number of partner organizations and customers who are accessing A. Datum websites and applications. This expansion has resulted in increasing complexity of the network infrastructure at A. Datum, and has also meant that the organization has to be much more aware of network level security.

IT management and the security group at A. Datum are also concerned with the level of compliance for all client computers on the network. A. Datum plans to implement NAP for all client computers and all client computer connections, but is starting with a pilot program to enable NAP for VPN users.

As one of the senior network administrators at A. Datum, you are responsible for implementing the new features in the Windows Server 2012 environment. You will implement some new DHCP and DNS features, and then implement IPAM to simplify the process for managing the IP infrastructure. You will also implement NAP for external VPN users.

### Objectives

- Configure new features in DNS and DHCP.
- Configure IP Address Management.
- Configure NAP for VPN client computers.
- Verify the NAP deployment.

### Lab Setup

Estimated time: **75 minutes**

Virtual Machines	20417A-LON-DC1 20417A-LON-SVR1 20417A-LON-SVR2 20417A-LON-CL1
User Name	Adatum\Administrator
Password	Pa\$\$w0rd

For this lab, you will use the available virtual machine environment. Before you begin the lab, you must complete the following steps:

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, click **20417A-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on using the following credentials:
  - a. User name: **Adatum\Administrator**
  - b. Password: **Pa\$\$w0rd**
5. Repeat steps 2 - 4 for **20417A-LON-SVR1**, **20417A-LON-SVR2** and **20417A-LON-CL1**.

## Exercise 1: Configure new features in DNS and DHCP

### Scenario

To increase security in your network, you want to implement new security features in DNS and DHCP. Also, you want to achieve high availability for IP addressing system. Therefore, you decided to implement DHCP Failover.

The main tasks for this exercise are as follows:

1. Configure DNSSEC.
2. Configure DHCP Name Protection.
3. Configure DHCP Failover.

#### ► Task 1: Configure DNSSEC

1. On LON-DC1, start the DNS Management console.
2. Use the DNSSEC zone signing wizard to sign the Adatum.com zone. Accept all the default settings.
3. Verify the DNSKEY resource records were created in the Trust Points zone.
4. Close the DNS Management console.
5. Use the Group Policy Management Console to configure NRPT. Create a rule that enables DNSSEC for the Adatum.com suffix and requires DNS client computers to check that the name and address data is validated.
6. Close the Group Policy Management Editor and Group Policy Management console.

#### ► Task 2: Configure DHCP Name Protection

1. Start the DHCP Management console.
2. Configure Name Protection for the IPv4 node.

#### ► Task 3: Configure DHCP Failover

1. On LON-SVR1, start the DHCP console and view the current state of DHCP. Note the server is authorized but no scopes are configured.
2. On LON-DC1, in the DHCP Management console, start the failover wizard.
3. Configure failover replication with the following settings:
  - Partner server = 172.16.0.21
  - Relationship Name = Adatum
  - Maximum Client Lead Time = 15 minutes
  - Mode = Load balance
  - Load Balance Percentage = 50%
  - State Switchover Interval = 60 minutes
  - Message authentication shared secret is Pa\$\$w0rd
  - Complete the wizard
4. Switch to LON-SVR1 and notice that the IPv4 node is active and the Adatum scope is configured.
5. Close the DHCP console on both LON-DC1 and LON-SVR1.

**Results:** After completing this exercise you will be able to configure DNSSEC, configure DHCP name protection, and configure and verify DHCP failover.

## Exercise 2: Configuring IP Address Management

### Scenario

A. Datum is evaluating solutions for simplifying IP management. Because you implemented Windows Server 2012, you decide to implement IPAM.

The main tasks for this exercise are as follows:

1. Install the IPAM Feature.
2. Configure IPAM Related GPOs.
3. Configure IP Management Server Discovery.
4. Configure Managed Servers.
5. Configure and Verify a New DHCP Scope with IPAM.

#### ► Task 1: Install the IPAM Feature

- On LON-SVR2, in Server Manager, add the IPAM feature and all required supporting features.

#### ► Task 2: Configure IPAM Related GPOs

1. On LON-SVR2, in Server Manager, click **IPAM**.
2. From the IPAM Overview pane provision the IPAM server.
3. Enter **IPAM** as the GPO name prefix.

#### ► Task 3: Configure IP Management Server Discovery

1. From the IPAM Overview pane, configure server discovery for the Adatum domain.
2. From the IPAM Overview pane, start the server discovery process.
3. In the yellow banner, click the **More** link to determine the discovery status.

#### ► Task 4: Configure Managed Servers

1. From the IPAM Overview pane, add the servers to manage. Verify that IPAM access is currently blocked for LON-DC1.
2. Start Windows PowerShell and grant the IPAM server permission. Use the following command:

```
Invoke-IPamGpoProvisioning -Domain Adatum.com -GpoPrefixName IPAM -IpamServerFqdn  
LON-SVR2.adatum.com
```

3. In the IPAM console, for LON-SVR1 and LON-DC1, set the manageability status to be **Managed**.
4. Switch to LON-DC1 and refresh Group Policy.
5. Switch to LON-SVR1, and refresh Group Policy.
6. Switch back to LON-SVR2 and refresh the IPAM console view.
7. Switch back to LON-SVR2, and in the IPAM console, configure LON-SVR1 to be **Managed**.
8. Refresh the Server Access Status and refresh the console view until LON-DC1 and LON-SVR1 shows an IPAM Access Status Unblocked. This may take 10-15 minutes to complete.
9. From the IPAM Overview pane retrieve data from the managed server.

### ► Task 5: Configure and Verify a New DHCP Scope with IPAM

1. Use IPAM to create a new DHCP scope called **TestScope** with the following parameters:
  - The scope start address will be **10.0.0.50**.
  - The scope end address will be **10.0.0.100**.
  - The subnet mask will be **255.0.0.0**.
  - The default gateway will be **10.0.0.1**.
2. On LON-DC1, verify the TestScope in the DHCP MMC.
3. Right-click the **TestScope** and then click **Deactivate**. Click **Yes**.
4. Close the DHCP console.
5. On LON-SVR2, close all open windows.

**Results:** After completing this exercise you will be able to install and configure the IPAM feature, configure IPAM related GPOs, configure IP Management server discovery, configure managed servers, and configure and verify a new DHCP scope with IPAM.

## Exercise 3: Configuring NAP

### Scenario

A. Datum has identified that remote client computers who connect through VPN have inconsistent security configuration. Because these client computers are accessing important data, it is important for all client computers to comply with company security policy. To increase security of your network and better manage client computers who establish remote connection, you decide to implement NAP for all VPN connections.

The main tasks for this exercise are as follows:

1. Configure Server and Client Certificate Requirements.
2. Install the Network Policy Server Role.
3. Configure Health Policies.
4. Configure Network Policies for Compliant and Noncompliant Computers.
5. Configure Connection Request Policies for VPN.

### ► Task 1: Configure Server and Client Certificate Requirements

1. On LON-SVR2, create a new management console for **Certificates** focused on the local computer.
2. Enroll a **Computer** certificate for LON-SVR2.
3. Switch to LON-CL1 and log on as **Adatum\administrator** with the password of **Pa\$\$w0rd**.
4. Create a new management console for **Certificates** focused on the local computer.
5. Enroll a **Computer** certificate for LON-CL1.

### ► Task 2: Install the Network Policy Server Role

- On LON-SVR2, add the Network Policy Server role service.

### ► Task 3: Configure Health Policies

1. On LON-SVR2, open the Network Policy Server console.
2. Configure the Windows Security Health Validator to only validate that the Windows Firewall is enabled.
3. Create two new Health Policies. One for compliant computers that pass all SHV checks and one for noncompliant computers that fail one or more SHV checks.

### ► Task 4: Configure Network Policies for Compliant and Noncompliant Computers

1. Configure a network policy for compliant computers in such a way that the health policy allows them full network access. Name the policy **Compliant Full-Access**.
2. Configure a network policy for noncompliant computers in such a way that the health policy enables them to exchange packets with LON-DC1 at 172.16.0.10 only. Name the policy **Noncompliant-Restricted**.

### ► Task 5: Configure Connection Request Policies for VPN

1. Disable the two default connection request policies.
2. Configure a new Connection Request Policy called **VPN connections**.
3. Add conditions for Point to Point Tunneling Protocol (PPTP), Secure Socket Tunneling Protocol (SSTP), and Layer 2 Tunneling Protocol (L2TP).
4. Ensure requests are authenticated on this server and will override network policy authentication.
5. Add Protected Extensible Authentication Protocol (PEAP) and edit it to enforce network access protection.

**Results:** After completing this exercise you will be able to configure server and client computer certificate requirements, install the NPS server role, configure health policies, configure network policies, and configure connection request policies for VPN.

## Exercise 4: Verifying the NAP Deployment

### Scenario

After you implemented NAP infrastructure and configured policies, you want to test NAP with VPN client computer.

The main tasks for this exercise are as follows:

1. Configure Security Center.
2. Enable a Client NAP Enforcement Method.
3. Allow Ping on LON-SVR2.
4. Move the Client to the Internet and Establish a VPN Connection.
5. To prepare for next module.

### ► Task 1: Configure Security Center

1. Log on to LON-CL1 as **Adatum\Administrator** with a password of **Pa\$\$w0rd**.
2. Use **gpedit.msc** to open Local Group Policy and turn on the Security Center.

► **Task 2: Enable a Client NAP Enforcement Method**

1. Use the NAP Client Configuration MMC to enable the EAP Quarantine Enforcement Client on LON-CL1.
2. Enable and start the NAP agent service.

► **Task 3: Allow Ping on LON-SVR2**

- On LON-SVR2, open Windows Firewall with Advanced Security.
- Configure a new inbound rule that allows ICMPv4 echo packets through the firewall.

► **Task 4: Move the Client to the Internet and Establish a VPN Connection**

1. Configure LON-CL1 with the following IP address settings:
    - o IP address: **131.107.0.20**
    - o Subnet Mask: **255.255.0.0**
  2. In Hyper-V Manager, right-click **20417A-LON-CL1** and then click **Settings**.
  3. Click **Legacy Network Adapter** and then under Network select Private Network 2, click **OK**.
  4. Verify that you can ping 131.107.0.1.
  5. Create a VPN on LON-CL1 with the following settings:
    - o Name: Adatum **VPN**
    - o Internet address: **131.107.0.2**
  6. Right-click the **Adatum VPN** connection, click **Properties**, and then click the **Security** tab.
  7. Under Authentication, click **Use Extensible Authentication Protocol (EAP)**.
  8. In the **Microsoft: Secured password (EAP-MSCHAP v2) (encryption enabled)** list, click **Microsoft: Protected EAP (PEAP) (encryption enabled)** and then click **Properties**.
  9. Ensure that the **Verify the server's identity by validating the certificate** check box is already selected. Clear the **Connect to these servers** check box, and then ensure that **Secured password (EAP-MSCHAP v2)** is already selected under Select Authentication Method. Clear the **Enable Fast Reconnect** check box and then select the **Enforce Network Access Protection** check box.
  10. Test the VPN connection.
- **To prepare for next module**
- Revert virtual machines to their initial state.

**Results:** After completing this exercise you will be able to configure Security Center, enable a client computer NAP enforcement method, allow Ping on LON-SVR2, and move the client computer to the Internet and establish a VPN connection.



## Module Review and Takeaways

### Best Practices

- Ensure that IPv6 is enabled on the IPAM server in order to manage IPv6 address spaces.
- Use Group Policy to configure NRPT tables for DNSSEC client computers.
- Disable authentication protocols that you are not using.
- Document the NPS configuration by using the NetshNps Show Config>Path\File.txt to save the configuration to a text file.

### Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Unable to connect to the IPAM server.	
Noncompliant NAP client computers are being denied network access instead of being sent to the restricted network	

### Review Question

**Question:** What is a major drawback of IPAM?

### Real-world Issues and Scenarios

**Scenario:** Tailspin Toys wants to implement IPsec NAP enforcement. What infrastructure components have to be in place to support this method?

**Scenario:** You have implemented DNSSEC, but now you have to disable DNSSEC. How will you disable DNSSEC?

### Tools

Tool	Use	Where to find it
DNS Management Console	Configure all aspects of DNS	In Server Manager under the Tools drop-down list.
DHCP Management Console	Configure all aspects of DHCP	In Server Manager under the Tools drop-down list.
Remote Access Management Console	Configure remote access such as VPN	In Server Manager under the Tools drop-down list.
NAP configuration wizard	Configure the NAP Enforcement Point	Open the NPS (Local) console. In Getting Started, under Standard Configuration, select Network Access Protection (NAP), and then click Configure NAP.

**MCT USE ONLY. STUDENT USE PROHIBITED**

# Module 6

## Implementing DirectAccess

### Contents:

Module Overview	6-1
Lesson 1: Overview of DirectAccess	6-2
Lesson 2: Installing and Configuring DirectAccess Components	6-14
Lab: Implementing DirectAccess	6-24
Module Review and Takeaways	6-33

## Module Overview

Introduced in Windows Server® 2008 R2, the DirectAccess feature is a technology that enables users to securely connect to data and resources in corporate networks without using traditional virtual private network (VPN) technology. In Windows Server 2012, DirectAccess is now one of three component technologies (DirectAccess, Routing, and Remote Access) that is integrated with a single, unified server role called Windows Server 2012 Remote Access. DirectAccess seamlessly integrates and coexists with what was formerly called Routing and Remote Access service (RRAS). Direct Access itself is expanded to add features such as integrated accounting, express setup for small and medium deployments, and multiple domain support.

In this module, you will learn how DirectAccess works for internal and external clients. You will also learn the new DirectAccess features introduced in Windows Server 2012 and Windows® 8. In addition, you will learn how to install and configure DirectAccess in different scenarios.

### Objectives

After completing this module, you will be able to:

- Describe the DirectAccess functionality in Windows Server 2012 and Windows 8.
- Install and configure DirectAccess in Windows Server 2012 and Windows 8.

## Lesson 1

# Overview of DirectAccess

DirectAccess enables remote users to securely access corporate resources, such as email servers, shared folders, or internal websites without connecting to a VPN. Also, DirectAccess provides increased productivity for a mobile workforce by offering the same connectivity experience both inside and outside the office. With the new unified management experience, you can configure DirectAccess and older VPN connections from one location. Other enhancements in DirectAccess include simplified deployment, and improved performance and scalability. This lesson provides an overview of the DirectAccess architecture and components.

### Lesson Objectives

After completing this lesson, you will be able to:

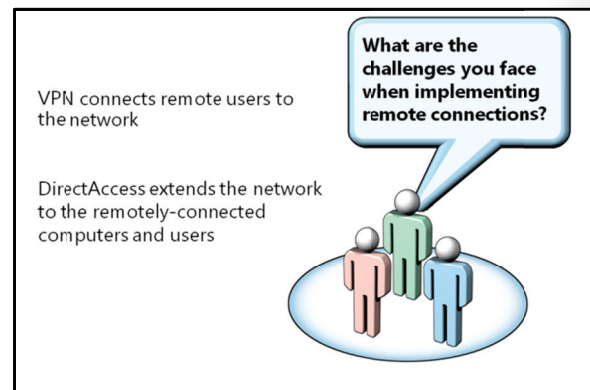
- Discuss the problems with remote connections.
- Describe the use of DirectAccess.
- Describe the new features of DirectAccess in Windows Server 2012.
- Describe the DirectAccess components.
- Describe the use of the Name Resolution Policy Table.
- Describe how DirectAccess works for internal clients.
- Describe how DirectAccess works for external clients.

### Problems with Remote Connections

Organizations often rely on traditional VPN connections to provide remote users with secure access to data and resources on the corporate network. VPN connections need to be configured most of the time manually. This sometimes present interoperability issues in situations when the users are using multiple different VPN clients. Additionally, VPN connections face the following problems:

- The user must initiate the VPN connection.
- The connection requires several steps and the connection process takes at least several seconds, or even more.
- The connection could require additional configuration on the corporate firewall. If not properly configured on the firewall, VPN connections usually enable remote access to the entire corporate network.
- Troubleshooting failed VPN connections can make up a significant portion of Help Desk calls for many organizations.

Moreover, organizations cannot effectively manage remote computers unless they are connected. VPN-based remote client computers present a challenge to IT professionals because these computers might not connect to the internal network for weeks at a time, preventing them from downloading Group Policy objects (GPOs) and software updates.



Also, if the organization does not require additional health checks in order to establish a network VPN connection, computers that are not updated and protected on a regular basis may contain malware. This malware could attempt to spread inside the corporate network through e-mail, shared folders, or automated network attacks.

### DirectAccess Extends the Network to the Remotely-Connected Computers and Users

To overcome these limitations in traditional VPN connections, organizations can implement DirectAccess to provide a seamless connection between the internal network and the remote computer on the Internet. With DirectAccess, organizations can effortlessly manage remote computers because they are always connected.

### What Is DirectAccess?

The DirectAccess feature in Windows Server 2012 enables seamless remote access to intranet resources without first establishing a user-initiated VPN connection. The DirectAccess feature also ensures seamless connectivity to the application infrastructure for internal users and remote users.

Unlike traditional VPNs that require user intervention to initiate a connection to an intranet, DirectAccess enables any IPv6-capable application on the client computer to have complete access to intranet resources.

DirectAccess also enables you to specify resources and client-side applications that are restricted for remote access.

Organizations benefit from DirectAccess because remote computers can be managed as if they are local computers. Using the same management and update servers, you can ensure they are always up-to-date and in compliance with security and system health policies. You can also define more detailed access control policies for remote access when compared with defining access control policies in VPN solutions.

DirectAccess offers the following features:

- Connects automatically to corporate intranet when connected to the Internet
- Uses various protocols, including HTTPS, to establish IPv6 connectivity—HTTPS is typically allowed through firewalls and proxy servers
- Supports selected server access and end-to-end Internet Protocol Security (IPsec) authentication with intranet network servers
- Supports end-to-end authentication and encryption with intranet network servers
- Supports management of remote client computers
- Allows remote users to connect directly to intranet servers

DirectAccess provides the following benefits:

- *Always-on connectivity.* Whenever the user connects the client computer to the Internet, the client computer is also connected to the intranet. This connectivity enables remote client computers to access and update applications more easily. It also makes intranet resources always available, and enables users to connect to the corporate intranet from anywhere and anytime, thereby improving their productivity and performance.

#### Features of DirectAccess

- Connects automatically to the corporate network over the public network
- Uses various protocols, including HTTPS, to establish IPv6 connectivity
- Supports selected server access and IPsec authentication
- Supports end-to-end authentication and encryption
- Supports management of remote client computers
- Allows remote users to connect directly to intranet servers

#### Benefits of DirectAccess

- Always-on connectivity
- Seamless connectivity
- Bidirectional access
- Manage-out Support
- Improved security
- Integrated solution



- *Seamless connectivity.* DirectAccess provides a consistent connectivity experience whether the client computer is local or remote. This allows users to focus more on productivity and less on connectivity options and process. This consistency can reduce training costs for users, with fewer support incidents.
- *Bidirectional access.* You can configure DirectAccess in a way that the DirectAccess clients have access to intranet resources and you can also have access from the intranet to those DirectAccess clients. Therefore, DirectAccess can be bidirectional. This ensures that the client computers are always updated with recent security updates, the domain Group Policy is enforced, and there is no difference whether the users are on the corporate intranet or on the public network. This bidirectional access also results in:
  - Decreased update time
  - Increased security
  - Decreased update miss rate
  - Improved compliance monitoring
- *Manage-out Support.* This feature is new in Windows Server 2012 and provides the ability to enable only remote management functionality in the DirectAccess client. This new sub-option of the DirectAccess client configuration wizard automates the deployment of policies that are used for managing the client computer. Manage-out support does not implement any policy options that allow users to connect to the network for file or application access. Manage-out support is unidirectional, incoming only access for administration purposes only.
- *Improved security.* Unlike traditional VPNs, DirectAccess offers many levels of access control to network resources. This tighter degree of control allows security architects to precisely control remote users who access specified resources. You can use a granular policy to specifically define which user can use DirectAccess, and the location from which the user can access it. IPsec encryption is used for protecting DirectAccess traffic so that users can ensure that their communication is safe.
- *Integrated solution.* DirectAccess fully integrates with Server and Domain Isolation and Network Access Protection (NAP) solutions, resulting in the seamless integration of security, access, and health requirement policies between the intranet and remote computers.

## What's New in DirectAccess in Windows Server 2012

In Windows Server 2012, DirectAccess has several enhancements, especially in regards to bypassing some common technology issues such as requirements for public key infrastructure (PKI) and public IP addresses.

### Improved DirectAccess Management

DirectAccess in Windows Server 2012 has been improved in the following ways:

- *DirectAccess and RRAS coexistence.* Windows Server 2012 DirectAccess and RRAS unified server role solve the problems of interoperability of Denial of Service Protection (DoSP) and Internet Key Exchange version 2 (IKEv2).

The new features of DirectAccess include:

#### Improved DirectAccess Management

- Rich monitoring of clients
- DirectAccess and RRAS coexistence
- Accounting and reporting
- Windows PowerShell and Server Core support
- Unified management wizard and tools



- *Rich monitoring of clients.* You can view the health of user computers and servers along with deployment monitoring and diagnostics in a single console in DirectAccess. Using the dashboard, you can have top-level information about Remote Access servers and client activity. User and client computer monitoring can provide you with information on which resources are accessed by the clients.
- *Integrated accounting and reporting.* Accounting and reporting is now integrated in the console and provides the ability to measure specific metrics. It also enables administrators to generate rich usage reports on various user and server statistics.
- *Windows PowerShell® and Server Core support.* Windows Server 2012 provides full Windows PowerShell support for the setup, configuration, management, monitoring, and troubleshooting of the Remote Access Server Role.
- *Unified management wizard and tools.* You can use a single wizard and console for DirectAccess configuration, management, and monitoring.
- *Works with existing infrastructure.* You do not need to upgrade your existing domain controllers to Windows Server 2012.
- *IPv6 for internal network is no longer required.* This is because transition technologies such as network address translation 64 (NAT64) and Domain Name System 64 (DNS64) allow access to internal resources that are run only on IPv4 computers. Previously, this functionality was only possible to achieve with deployments that included Microsoft Unified Access Gateway Server.
- *Single network adapter.* You can implement your DirectAccess server behind a NAT with a single network adapter.
- *Single IP address.* In certain deployment scenarios, you can even use a single IP address for the DirectAccess server. This makes deployment easier in comparison to the DirectAccess deployment in Windows Server 2008.

### **Simplified DirectAccess Deployment**

The DirectAccess deployment has been simplified. Windows Server 2012 provides Express Setup for small and medium deployment. Express Setup includes the following characteristics:

- PKI deployment is optional, because the wizard creates a self-signed certificate without the need for certificate revocation lists (CRL) lists. This functionality is achieved by the using the HTTPS-based Kerberos proxy (built into Windows Server 2012) which accepts client authentication requests and sends them to domain controllers on behalf of the client.
- Single IPsec tunnel configuration.
- Single factor authentication only; no support for smart card integration or using one-time password (OTP).
- Works only with client computers running Windows 8.

### **Performance and Scalability Improvements**

DirectAccess includes the following improved features in performance and scalability:

- *Support for high availability and external load balancers.* Windows Server 2012 supports network load balancing (NLB) to achieve high availability and scalability for both DirectAccess and RRAS. The setup process also provides integrated support for third party external hardware-based load balancer solutions.

- *Improved support for Receive Site Scaling (RSS).* DirectAccess provides support for RSS and supports running DirectAccess in virtual machines with increased density:
  - *IP-HTTPS interoperability and performance improvements.* Windows Server 2012 DirectAccess implementation removes double encryption when using IP-HTTPS. Also, it reduces the time for duplicate address detection, resulting in a significant performance improvement.
  - *Lower bandwidth utilization.* Windows Server 2012 reduces the overhead associated with establishing of connectivity methods, optimizes batched send behavior, and receives buffers, which result in overall lower bandwidth utilization. Additionally Windows Server 2012 DirectAccess receives site scaling with User Datagram Protocol (UDP).

## New Deployment Scenarios


The new DirectAccess deployment scenarios in Windows Server 2012 include:

- *Deploying multiple endpoints.* When you implement DirectAccess on multiple servers in different network locations, the Windows 8 device automatically chooses the closest endpoint. (For the Windows 7 operating system, you have to specify the endpoint manually). This also works for distributed file system (DFS) shares that are redirected to an appropriate Active Directory® site.
- *Multiple domain support.* This feature is integrated with Windows Server 2012.
- *Deploy a server behind a NAT.* You can deploy Windows Server 2012 DirectAccess behind a NAT device, with the support for a single or multiple interfaces, removing the prerequisite for a public address. In this configuration, only IP over HTTPS (IP-HTTPS) is deployed which allows secure IP tunnel to be established by using a secure HTTP connection.
- *Support for OTP and virtual smart cards.* This feature requires a PKI deployment. If the option is selected in the DirectAccess Setup Wizard, the Use computer certificates option is automatically selected. Also, DirectAccess can use the Trusted Platform Module (TPM)–based virtual smart card which use TPM of a client computer to act as a virtual smart card for two-factor authentication.
- *Offload network adapters with support for network teaming.* Network teaming in Windows Server 2012 is fully supported without the need for third-party drivers.
- *Off-premise provisioning.* With the new djoin tool, you can easily provision non-domain computer with an Active Directory blob, so that the computer can be joined in a domain without the need to be ever connected in your internal premises.

## DirectAccess Components

To deploy and configure DirectAccess, your organization must support the following infrastructure components:

- DirectAccess server
- DirectAccess clients
- Network location server
- Internal resources
- Active Directory domain
- Group Policy
- PKI (Optional for the internal network)
- DNS server
- NAP server





## DirectAccess Server

DirectAccess server can be any Windows Server 2012 joined in a domain, which accepts connections from DirectAccess clients and establishes communication with intranet resources. This server provides authentication services for DirectAccess clients and acts as an IPsec tunnel mode endpoint for external traffic. The new Remote Access server role allows centralized administration, configuration, and monitoring for both DirectAccess and VPN connectivity.

Compared with previous implementation in Windows Server 2008 R2, the new wizard-based setup simplifies DirectAccess management for small and medium organizations, by removing the need for full PKI deployment and removing the requirement for two consecutive public IPv4 addresses for the physical adapter that is connected to the Internet. In Windows Server 2012, the wizard detects the actual implementation state of the DirectAccess server, and automatically selects the best deployment; thereby, hiding from the administrator the complexity of configuring manually IPv6 transition technologies.

## DirectAccess Clients

DirectAccess clients can be any domain-joined computer running Windows 8, Windows 7 Enterprise Edition, or Windows 7 Ultimate Edition.



**Note:** With off-premise provisioning, you can join the client computer in a domain without connecting the client computer in your internal premises.

The DirectAccess client computer connects to the DirectAccess server by using IPv6 and IPsec. If a native IPv6 network is not available, the client establishes an IPv6-over-IPv4 tunnel by using 6to4 or Teredo. Note that the user does not have to be logged on to the computer for this step to complete.

If a firewall or proxy server prevents the client computer using 6to4 or Teredo from connecting to the DirectAccess server, the client computer automatically attempts to connect by using the IP-HTTPS protocol, which uses a Secure Sockets Layer (SSL) connection to ensure connectivity.

## Network Location Server

DirectAccess clients use the network location server (NLS) to determine their location. If the client computer can connect with HTTPS, then the client computer assumes it is on the intranet and disables DirectAccess components. If the NLS is not contactable, the client assumes it is on the Internet. The NLS server is installed with the web server role.



**Note:** The URL for the NLS is distributed by using GPO.

## Internal Resources

You can configure any IPv6-capable application which is running on internal servers or client computers to be available for DirectAccess clients. For older applications and servers not based on Windows and have no IPv6 support, Windows Server 2012 now includes native support for protocol translation (NAT64) and name resolution (DNS64) gateway to convert IPv6 communication from DirectAccess client to IPv4 for the internal servers.



**Note:** As done in the past, this functionality can also be achieved with Microsoft® Forefront® Unified Access Gateway Server. Likewise, as in past versions, these translation services do not support sessions initiated by internal devices; rather they support requests originating from ipv6 DirectAccess clients only.

## Active Directory Domain

You must deploy at least one Active Directory domain, running at a minimum Windows Server 2008 R2 domain functional level. Windows Server 2012 DirectAccess provides integrated multiple domain support which allows client computers from different domains to access resources that may be located in different trusted domains.

## Group Policy

Group Policy is required for the centralized administration and deployment of DirectAccess settings. The DirectAccess Setup Wizard creates a set of GPOs and settings for DirectAccess clients, the DirectAccess server, and selected servers.

## PKI

PKI deployment is optional for simplified configuration and management. Windows Server 2012 DirectAccess enables client authentication requests to be sent over a HTTPS based Kerberos proxy service running on the DirectAccess server. This eliminates the need for establishing a second IPsec tunnel between clients and domain controllers. The Kerberos proxy will send Kerberos requests to domain controllers on behalf of the client.

However, for a full DirectAccess configuration, that allows NAP integration, two-factor authentication, and force tunneling, you still need to implement certificates for authentication for every client that will participate in DirectAccess communication.

## DNS Server

When using Intra-Site Automatic Tunnel Addressing Protocol (ISATAP), you must use at least Windows Server 2008 R2, Windows Server 2008 with the Q958194 hotfix, Windows Server 2008 SP2 or later, or a third-party DNS server that supports DNS message exchanges over the ISATAP.

## NAP Servers

NAP is an optional component of the DirectAccess solution that allows you to provide compliance checking and enforce security policy for DirectAccess clients over the Internet. Windows Server 2012 DirectAccess provides the ability to configure NAP health check directly from the setup user interface instead of manual editing of GPO as it was in Windows Server 2008 R2 DirectAccess.



**Additional Reading:** The DNS server does not listen on the ISATAP interface on a Windows Server 2008-based computer  
<http://go.microsoft.com/fwlink/?LinkID=159951>

IPv6 - Technology Overview  
<http://technet.microsoft.com/en-us/library/hh831730.aspx>

## Name Resolution Policy Table

To separate Internet traffic from intranet traffic in DirectAccess, Windows Server 2012 and Windows 8 include the Name Resolution Policy Table (NRPT), a feature that allows DNS servers to be defined per DNS namespace, rather than per interface.

The NRPT stores a list of rules. Each rule defines a DNS namespace and configuration settings that describe the DNS client's behavior for that namespace.

When a DirectAccess client is on the Internet, each name query request is compared against the namespace rules stored in the NRPT:

- If a match is found, the request is processed according to the settings in the NRPT rule.
- If a name query request does not match a namespace listed in the NRPT, the request is sent to the DNS servers configured in the TCP/IP settings for the specified network interface.

DNS settings are configured depending on the client location:

- For a remote client computer, the DNS servers are typically the Internet DNS servers configured through the Internet Service Provider (ISP).
- For a DirectAccess client on the intranet, the DNS servers are typically the intranet DNS servers configured through Dynamic Host Configuration Protocol (DHCP).

Single-label names, for example, `http://internal`, typically have configured DNS search suffixes appended to the name before they are checked against the NRPT.

If no DNS search suffixes are configured, and the single-label name does not match any other single-label name entry in the NRPT, the request is sent to the DNS servers specified in the client's TCP/IP settings.

Namespaces, for example, `internal.adatum.com`, are entered into the NRPT, followed by the DNS servers to which requests matching that namespace should be directed. If an IP address is entered for the DNS server, all DNS requests are sent directly to the DNS server over the DirectAccess connection. You need not specify any additional security for such configurations. However, if a name is specified for the DNS server, such as `dns.adatum.com` in the NRPT, the name must be publicly resolvable when the client queries the DNS servers specified in its TCP/IP settings.

The NRPT allows DirectAccess clients to use intranet DNS servers for name resolution of internal resources and Internet DNS for name resolution of other resources. Dedicated DNS servers are not required for name resolution. DirectAccess is designed to prevent the exposure of your intranet namespace to the Internet.

Some names need to be treated differently with regards to name resolution; these names should not be resolved by using intranet DNS servers. To ensure that these names are resolved with the DNS servers specified in the client's TCP/IP settings, you must add them as NRPT exemptions.

NRPT is controlled through Group Policy. When the computer is configured to use NRPT, the name resolution mechanism uses the following in order:

- The local name cache
- The hosts file
- NRPT

NRPT is a table that defines DNS servers for different namespaces and corresponding security settings; NRPT is used before the adapter's DNS settings

Using NRPT:

- DNS servers can be defined for each DNS namespace rather than for each interface
- DNS queries for specific namespaces can be optionally secured by using IPSec



Then, the name resolution mechanism finally sends the query to the DNS servers specified in the TCP/IP settings.


**Question:** How can you benefit from NRPT?

**Question:** How can you benefit by using connection security rules for Direct Access?

## How DirectAccess Works for Internal Client Computers

An NLS is an internal network server that hosts an HTTPS-based URL. DirectAccess clients try to access a NLS URL to determine if they are located on the intranet or on a public network. The DirectAccess server can also be the NLS. In some organizations where DirectAccess is a business-critical service, the NLS should be highly available. Generally, the web server on the NLS does not have to be dedicated just for supporting DirectAccess clients.

It is critical that the NLS is available from each company location, because the behavior of the DirectAccess client depends on the response from the NLS. Branch locations may need a separate NLS at each branch location to ensure that the NLS remains accessible even when there is a link failure between branches.



## How DirectAccess Works for Internal Clients

The DirectAccess connection process happens automatically, without requiring user intervention. DirectAccess clients use the following process to connect to intranet resources:

1. The DirectAccess client tries to resolve the fully qualified domain name (FQDN) of the NLS URL.  
Because the FQDN of the NLS URL corresponds to an exemption rule in the NRPT, the DirectAccess client instead sends the DNS query to a locally-configured DNS server (an intranet-based DNS server). The intranet-based DNS server resolves the name.
2. The DirectAccess client accesses the HTTPS-based URL of the NLS, during which process it obtains the certificate of the NLS.
3. Based on the CRL distribution points field of the NLS's certificate, the DirectAccess client checks the CRL revocation files in the CRL distribution point to determine if the NLS's certificate has been revoked.
4. Based on an HTTP 200 Success of the NLS URL (successful access and certificate authentication and revocation check), the DirectAccess client switches to domain firewall profile and ignores the DirectAccess rules in the NRPT for the remainder of the session.
5. The DirectAccess client computer attempts to locate and log on to the Active Directory Domain Services (AD DS) domain by using its computer account.

Because the client no longer references any DirectAccess rules in the NRPT for the rest of the connected session, all DNS queries are sent through interface-configured DNS servers (intranet-based DNS servers).

With the combination of network location detection and computer domain logon, the DirectAccess client configures itself for normal intranet access.

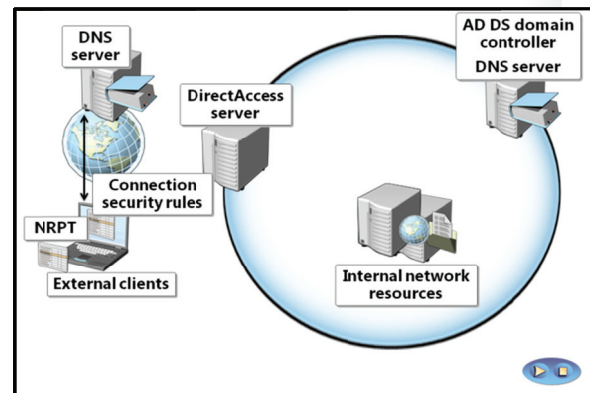
6. Based on the computer's successful logon to the domain, the DirectAccess client assigns the domain (firewall network) profile to the attached network.

By design the DirectAccess Connection Security tunnel rules are scoped for the public and private firewall profiles, they are disabled from the list of active connection security rules.

The DirectAccess client has successfully determined that it is connected to its intranet and does not use DirectAccess settings (NRPT rules or Connection Security tunnel rules). The DirectAccess client can access intranet resources normally. It can also access Internet resources through normal means, such as a proxy server.

## How DirectAccess Works for External Client Computers

When a DirectAccess client starts, the DirectAccess client assumes that it is not connected to the intranet by trying to reach the URL address specified for NLS. Because the client computer cannot communicate with NLS, it starts to use NRPT and connection security rules. The NRPT has DirectAccess-based rules for name resolution, and connection security rules define DirectAccess IPsec tunnels for communication with intranet resources. Internet-connected DirectAccess clients use the following process to connect to intranet resources.



The DirectAccess client first attempts to access the NLS. Then, the client attempts to locate a domain controller. Afterwards, the client attempts to access intranet resources and internet resources.

### DirectAccess Client Attempts To Access the Network Location Server

The DirectAccess client attempts to access the NLS as follows:

1. The client tries to resolve the FQDN of the NLS URL. Because the FQDN of the NLS URL corresponds to an exemption rule in the NRPT, the DirectAccess client does not send the DNS query to a locally-configured DNS server (an Internet-based DNS server). An external Internet-based DNS server would not be able to resolve the name.
2. The DirectAccess client processes the name resolution request as defined in the DirectAccess exemption rules in the NRPT.
3. Because the NLS is not found on the same network as the DirectAccess client is currently located on, the DirectAccess client applies a public or private firewall network profile to the attached network.
4. The Connection Security tunnel rules for DirectAccess, scoped for the public and private profiles, provide the public or private firewall network profile.

The DirectAccess client uses a combination of NRPT rules and connection security rules to locate and access intranet resources across the Internet through the DirectAccess server.

## DirectAccess Client Attempts To Locate a Domain Controller

After starting up and determining its network location, the DirectAccess client attempts to locate and log on to a domain controller. This process creates an IPsec tunnel or infrastructure tunnel by using the IPsec tunnel mode and Encapsulating Security Payload (ESP) to the DirectAccess server. The process is as follows:

1. The DNS name for the domain controller matches the intranet namespace rule in the NRPT, which specifies the IPv6 address of the intranet DNS server. The DNS client service constructs the DNS name query that is addressed to the IPv6 address of the intranet DNS server and forwards it to the DirectAccess client's TCP/IP stack for sending.
2. Before sending the packet, the TCP/IP stack checks to determine if there are Windows Firewall outgoing rules or connection security rules for the packet.
3. Because the destination IPv6 address in the DNS name query matches a connection security rule that corresponds with the infrastructure tunnel, the DirectAccess client uses AuthIP and IPsec to negotiate and authenticate an encrypted IPsec tunnel to the DirectAccess server. The DirectAccess client (both the computer and the user) authenticates itself with its installed computer certificate and its NT LAN Manager (NTLM) credentials, respectively.



**Note:** AuthIP enhances authentication in IPsec by adding support for user-based authentication with Kerberos v5 or SSL certificates. AuthIP also supports efficient protocol negotiation and usage of multiple sets of credentials for authentication.

4. The DirectAccess client sends the DNS name query through the IPsec infrastructure tunnel to the DirectAccess server.
5. The DirectAccess server forwards the DNS name query to the intranet DNS server. The DNS name query response is sent back to the DirectAccess server and back through the IPsec infrastructure tunnel to the DirectAccess client.

Subsequent domain logon traffic goes through the IPsec infrastructure tunnel. When the user on the DirectAccess client logs on, the domain logon traffic goes through the IPsec infrastructure tunnel.

## DirectAccess Client Attempts To Access Intranet Resources

The first time that the DirectAccess client sends traffic to an intranet location that is not on the list of destinations for the infrastructure tunnel (such as an email server), the following process occurs:

1. The application or process that attempts to communicate constructs a message or payload and hands it off to the TCP/IP stack for sending.
2. Before sending the packet, the TCP/IP stack checks to determine if there are Windows Firewall outgoing rules or connection security rules for the packet.
3. Because the destination IPv6 address matches the connection security rule that corresponds with the intranet tunnel (which specifies the IPv6 address space of the entire intranet), the DirectAccess client uses AuthIP and IPsec to negotiate and authenticate an additional IPsec tunnel to the DirectAccess server. The DirectAccess client authenticates itself with its installed computer certificate and the user account's Kerberos credentials.
4. The DirectAccess client sends the packet through the intranet tunnel to the DirectAccess server.
5. The DirectAccess server forwards the packet to the intranet resources. The response is sent back to the DirectAccess server and back through the intranet tunnel to the DirectAccess client.

Any subsequent intranet access traffic that does not match an intranet destination in the infrastructure tunnel connection security rule goes through the intranet tunnel.

## DirectAccess Client Attempts To Access Internet Resources

When the user or a process on the DirectAccess client attempts to access an Internet resource (such as an Internet web server), the following process occurs:

1. The DNS client service passes the DNS name for the Internet resource through the NRPT. There are no matches. The DNS client service constructs the DNS name query that is addressed to the IP address of an interface-configured Internet DNS server and hands it off to the TCP/IP stack for sending.
2. Before sending the packet, the TCP/IP stack checks to determine if there are Windows Firewall outgoing rules or connection security rules for the packet.
3. Because the destination IP address in the DNS name query does not match the connection security rules for the tunnels to the DirectAccess server, the DirectAccess client sends the DNS name query normally.
4. The Internet DNS server responds with the IP address of the Internet resource.
5. The user application or process constructs the first packet to send to the Internet resource. Before sending the packet, the TCP/IP stack checks to determine if there are Windows Firewall outgoing rules or connection security rules for the packet.
6. Because the destination IP address in the DNS name query does not match the connection security rules for the tunnels to the DirectAccess server, the DirectAccess client sends the packet normally.

Any subsequent Internet resource traffic that does not match a destination in either the infrastructure intranet tunnel or connection security rules is sent and received normally.

Like the connection process, accessing the domain controller and intranet resources is also a very similar process, because both of these processes are using NRPT tables to locate appropriate DNS server to resolve the name queries, with the differences of the IPsec tunnel that is established between the client and DirectAccess server. When accessing the domain controller, all the DNS queries are sent through the IPsec infrastructure tunnel, and when accessing intranet resources, a second IPsec tunnel is established (intranet tunnel).



## Lesson 2

# Installing and Configuring DirectAccess Components

In order to install and configure DirectAccess in your organization, you need to meet a number of requirements pertaining to Active Directory configuration, DNS configuration, and certificate services. After these requirements are met, you then install and configure the DirectAccess role. Finally, you configure client computers, and verify that DirectAccess is functional when connecting from both the internal network and the Internet.

In this lesson, you will learn about DirectAccess requirements, how to plan the DirectAccess solution, and the process of installation and deployment of DirectAccess. You will also learn about the new features for implementing DirectAccess in Windows 8.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe the prerequisites for implementing DirectAccess.
- Describe the process of configuring DirectAccess.
- Configure AD DS services for DirectAccess.
- Install and configure DirectAccess Server.
- Configure the DirectAccess clients.
- Describe the differences in DirectAccess between Windows 7 and Windows 8.

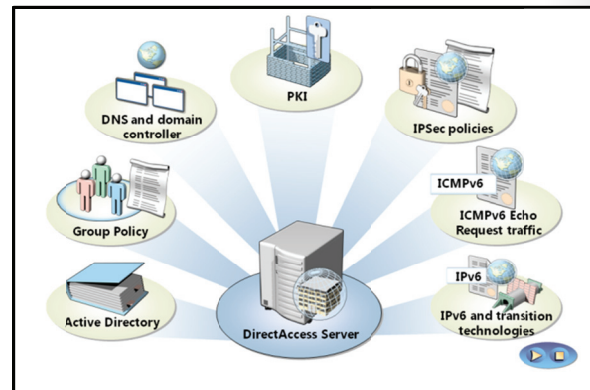
### Prerequisites for Implementing DirectAccess

To deploy DirectAccess, the DirectAccess server, the client computer, and infrastructure should meet certain requirements.

#### Requirements for DirectAccess Server

In order to deploy DirectAccess, you need to ensure that the server meets the hardware and network requirements:

- The server must be joined to an Active Directory domain.
- The server must have Windows Server 2012 or Windows Server 2008 R2 operating system installed.
- The Windows Server 2012 that will be installed as the DirectAccess Server can have a single network adapter installed which is connected to the intranet and published over Microsoft Forefront Threat Management Gateway 2010 (TMG) or Microsoft Forefront Unified Access Gateway 2010 (UAG) for Internet connection. In the deployment scenario where DirectAccess is installed on an Edge server, it needs to have two network adapters, one connected to the internal network and the other connected to the external network.



**Note:** An Edge server is any server that resides on the edge between two or more networks, typically a private network and Internet.



- Implementation of DirectAccess in Windows Server 2012 does not require two consecutive static, public IPv4 addresses be assigned to the network adapter. However, to achieve two-factor authentication with smart card or OTP deployment, DirectAccess server will still need two public IP addresses.
- You can even deploy Windows Server 2012 DirectAccess behind a NAT device, with support for a single or multiple interfaces, thereby circumnavigating the need for an additional public address. In this configuration, only IP over HTTPS (IP-HTTPS) is deployed which allows a secure IP tunnel to be established using a secure HTTP connection.
- On the DirectAccess server, you can install the Remote Access role to configure DirectAccess settings for the DirectAccess server and clients, and monitor the status of the DirectAccess server. The Remote Access wizard provides you with the option to configure only DirectAccess, only VPN, or both scenarios on the same server running Windows Server 2012. This was not possible in Windows Server 2008 R2 deployment of DirectAccess.
- For Load Balancing Support, Windows Server 2012 has the ability to use NLB (up to 8 nodes) to achieve high availability and scalability for both DirectAccess and RRAS.

### Requirements for DirectAccess Client

To deploy DirectAccess, you also need to ensure that the client computer meets certain requirements:

- The client computer should be joined to an Active Directory domain.
- With the new 2012 DirectAccess scenario it is possible to offline provision computers for domain membership without the need for the computer to be on premises.
- The client computer can be loaded with Windows 8, Windows 7 Enterprise Edition, Windows 7 Ultimate Edition, Windows Server 2012, or Windows Server 2008 R2 operating system.

You cannot deploy DirectAccess on clients running Windows Vista®, Windows Server 2008, or other earlier versions of the Windows operating systems.

### Infrastructure Requirements

The following are the infrastructure requirements to deploy DirectAccess:

- *Active Directory.* You must deploy at least one Active Directory domain. Workgroups are not supported.
- *Group Policy.* You need Group Policy for centralized administration and deployment of DirectAccess client settings. The DirectAccess Setup Wizard creates a set of GPOs and settings for DirectAccess clients, DirectAccess servers, and management servers.
- *DNS and domain controller.* You must have at least one domain controller and DNS server running Windows Server 2012, or Windows Server 2008 SP2 or Windows Server 2008 R2.
- *PKI.* You need to use PKI to issue computer certificates for authentication and health certificates only when NAP is deployed. You do not need external certificates. The SSL certificate installed on the DirectAccess server must have a CRL distribution point that is reachable from the Internet. The certificate Subject field must contain the FQDN that can be resolved to a public IPv4 address assigned to the DirectAccess server by using the Internet DNS.
- *IPsec policies.* DirectAccess utilizes IPsec policies that are configured and administered as part of Windows Firewall with Advanced Security.

- *Internet Control Message Protocol Version 6 (ICMPv6) Echo Request traffic.* You must create separate inbound and outbound rules that allow ICMPv6 Echo Request messages. The inbound rule is required to allow ICMPv6 Echo Request messages and is scoped to all profiles. The outbound rule to allow ICMPv6 Echo Request messages is scoped to all profiles and is only required if the Outbound block is turned on. DirectAccess clients that use Teredo for IPv6 connectivity to the intranet use the ICMPv6 message when establishing communication.
- *IPv6 and transition technologies.* IPv6 and the transition technologies such as ISATAP, Teredo, and 6to4 must be available for use on the DirectAccess server. For each DNS server running Windows Server 2008 or Windows Server 2008 R2, you need to remove the ISATAP name from the global query block list.

**Question:** You have Windows Server 2003 Certificate Authority server in your domain. Can you use the existing PKI infrastructure for DirectAccess or should you set up the new Certificate Authority server on Windows Server 2008 R2?

## Process of Configuring DirectAccess

To configure DirectAccess, perform the following steps:

### 1. Configure AD DS and DNS requirements

- Create a security group in Active Directory and add all client computer accounts that will be accessing intranet through DirectAccess.
- Configure both internal and external DNS servers with appropriate host names and IP addresses.

#### To configure DirectAccess:

1. Configure the AD DS domain controller and DNS
2. Configure the PKI environment
3. Configure the DirectAccess server
4. Configure the DirectAccess clients and test intranet and Internet access

### 2. Configure the PKI environment

- Add and configure the Certificate Authority server role, create the certificate template and CRL distribution point, publish the CRL list, and distribute the computer certificates.

### 3. Configure DirectAccess Server

- Install Windows Server 2012 on a server computer with one or two physical network adapters (depends on DirectAccess design scenario).
- Join the DirectAccess server to an Active Directory domain.
- Install the Remote Access role and configure the DirectAccess server so that it is either one of the following:
  - The DirectAccess server is on the perimeter network with one network adapter connected to the perimeter network and at least one other network adapter connected to the intranet. In this deployment scenario, DirectAccess server is placed between a front-end firewall and back-end firewall.
  - The DirectAccess server is published by using IPsec Gateway (TMG or UAG). In this deployment scenario, DirectAccess is placed behind a front-end firewall and it has one network adapter connected to internal network.
  - The DirectAccess server is installed on an Edge server (typically front end firewall) with one network adapter connected to the Internet and at least one other network adapter connected to the intranet.

An alternative design is that the DirectAccess server has only one, and not two, network interface. For this design, perform the following steps:

- Verify that the ports and protocols needed for DirectAccess and Internet Control Message Protocol (ICMP) Echo Request are enabled in the firewall exceptions and opened on the perimeter and Internet-facing firewalls.
  - The DirectAccess server in simplified implementation can use a single public IP address in combination with Kerberos Proxy services for client authentication against domain controllers. For two-factor authentication and integration with NAP, you need to configure at least two consecutive public static IPv4 addresses that are externally resolvable through DNS. Ensure that you have an IPv4 address available and that you have the ability to publish that address in your externally-facing DNS server.
  - If you have disabled IPv6 on clients and servers, enable IPv6 because it is required for DirectAccess.
  - Install a web server on the DirectAccess server to enable DirectAccess clients and determine if they are inside or outside the intranet. You can install this web server on a separate internal server for determining the network location.
  - Based on the deployment scenario, you need to designate one of the server network adapters as the Internet-facing interface (in deployment with two network adapters) or publish the DirectAccess server which is deployed behind NAT for Internet access.
  - On the DirectAccess server, ensure that the Internet-facing interface is configured to be either a Public or a Private interface, depending on your network design. Configure the intranet interfaces as domain interfaces. If you have more than two interfaces, ensure that no more than two classification types are selected.
4. **Configure the DirectAccess clients and test intranet and Internet access**
- Verify that DirectAccess group policy has been applied and certificates have been distributed to client computers:
  - Test whether you can connect to DirectAccess server from an intranet.
  - Test whether you can connect to DirectAccess server from the Internet.

## Demonstration: Configuring AD DS and Network Services for DirectAccess

In this demonstration, you will see how to:

- Create a security group for DirectAccess computers.
- Configure firewall rules for ICMPv6 traffic.
- Create required DNS records.
- Configure the PKI environment.

### Demonstration Steps

#### Create a security group for DirectAccess client computers

1. On LON-DC1, open the Active Directory Users and Computers console, and create an organizational unit with the name **DA\_Clients OU** and inside that organizational unit, create a Global Security group with the name **DA\_Clients**.
2. Add **LON-SVR3** to the **DA\_Clients** security group.

3. Close the Active Directory Users and Computers console.

**Question:** Why did you create the DA\_Clients group?

### Configure firewall rules GPO for ICMPv6 traffic

1. Open the Group Policy Management console, and then right-click **Default Domain Policy**.
2. In the console tree of the Group Policy Management Editor, navigate to Computer Configuration \Policies\Windows Settings\Security Settings\Windows Firewall with Advanced Security \Windows Firewall with Advanced Security.
3. Create a new inbound rule with the following settings:
  - o Rule Type: **Custom**
  - o Protocol type: **ICMPv6**
  - o Specific ICMP types: **Echo Request**
  - o Name: **Inbound ICMPv6 Echo Requests**
4. Create a new outbound rule with the following settings:
  - o Rule Type: **Custom**
  - o Protocol type: **ICMPv6**
  - o Specific ICMP types: **Echo Request**
  - o Action: **Allow the connection**
  - o Name: **Outbound ICMPv6 Echo Requests**
5. Close the Group Policy Management Editor and Group Policy Management consoles.

### Create required DNS records

1. Open the DNS Manager console and then create two new host records with the following settings:
  - o Name: **nls**; IP Address: **172.16.0.22**
  - o Name: **crl**; IP Address: **172.16.0.22**
2. Close the DNS Manager console.

**Question:** What is the purpose of the nls.adatum.com DNS host record that you associated with an internal IP address?

### Configure the PKI environment

1. Switch to LON-DC1.
2. Open the Certification Authority console.
3. Configure the **AdatumCA** certification authority with the following extension settings:
  - o Add Location: **http://crl.adatum.com/crld/**
  - o Variable: **CAName**, **CRLNameSuffix**, and **DeltaCRLAllowed**
  - o Location: **.crl**
  - o Select **Include in CRLs. Clients use this to find Delta CRL locations** and **Include in the CDP extension of issued certificates**
  - o Do not restart Certificate Services.
  - o Add Location: **\\lon-svr2\crl-dist\$\**

- Variable: **CAName**, **CRLNameSuffix**, and **DeltaCRLAllowed**
  - Location: **.crl**
  - Select **Publish CRLs to this location** and **Publish Delta CRLs to this location**
4. Restart Certificate Services.
  5. Close the Certificate Authority console.

### Configure permissions on the web server certificate template



**Note:** Users require the Enroll permission on the certificate.

1. Right-click **Certificate Template** in the Certification Authority console and then click **manage**.
2. In the Certificate Template console, in **Web Server** template **Properties**, configure security settings for **Authenticated Users** to be allowed to **Enroll** for a certificate.
3. Close the Certificate Templates console.

### Configure computer certificate auto-enrollment

1. On LON-DC1, open Group Policy Management console.
2. In the console tree, expand **Forest: Adatum.co\Domains\Adatum.com**.
3. Edit the Default Domain Policy and in the console tree of the Group Policy Management Editor, open Computer Configuration\Policies\Windows Settings\Security Settings\Public Key Policies.
4. At **Automatic Certificate Request Settings**, configure **Automatic Certificate Request** with a **Computer**.
5. On the **Certificate Template** page, click **Computer**, click **Next**, and then click **Finish**.
6. Close the Group Policy Management Editor and close the Group Policy.

## Demonstration: Configuring the DirectAccess Server

In this demonstration, you will see how to:

- Obtain certificates for IPsec.
- Configure DirectAccess.

### Demonstration Steps

#### Obtain the required certificates for LON-SVR2

1. Switch to LON-SVR2.
2. Open Microsoft Management Console by typing the **mmc** command, and then add the **Certificates** snap-in for **Local computer**.
3. In the Certificates snap-in, in the Microsoft Management Console, request a new certificate with the following settings:
  - Certificate template: **Web Server**
  - Common name: **131.107.0.2**
4. Verify that a new certificate with the name **131.107.0.2** has been issued with **Intended Purposes** of **Server Authentication**.

5. For the **131.107.0.2** certificate, in **Properties**, specify the **Friendly Name** as **IP-HTTPS Certificate**, and then click **OK**.
6. In the Certificates console, right-click the certificate with the name **lon-svr2.adatum.com**, and then click **delete**.
7. Close the Certificates snap-in console without saving it.
8. Close the console.

### Complete the DirectAccess setup wizard on LON-SVR2

1. Open the Server Manager console.
2. In the Server Manager console, open the Remote Access Management console.
3. Click **Configuration**; the **Enable Direct Access Wizard** will start automatically.
4. Click **Next**. Wait until the **DirectAccess** prerequisites page completes loading.
5. Complete the **Enable Direct Access Wizard** by using the following settings:
  - o DirectAccess Client Setup page; Enter the object names to select: **DA\_clients**
  - o Remote Access Server setup page,
    - Network Topology: **Edge**
    - Type the public name or IPv4 address used by clients to connect to the Remote Access server: **131.107.0.2**



**Note:** On this page, you might notice that you are using IP address of the Edge server instead of FQDN. This is because in this lab environment there is no public DNS server, as it would exist in real-life scenario.

- Infrastructure Server Setup page: Accept default values
  - Configure Remote Access page: Accept default values
6. Wait until **Enable DirectAccess Wizard Apply** completes, and then click **Close**.
  7. At the command prompt, type the following command:

```
GPUpdate /force
```

8. Close the Server Manager console.

### Demonstration: Configuring the DirectAccess Client

To prepare the DirectAccess clients and test the DirectAccess environment, complete the following tasks:

- Configure the DirectAccess client.
- Verify that DirectAccess clients have the computer certificate that is required for DirectAccess authentication. This should have been distributed with Group Policy.
- Verify that the client can connect to intranet resources.

## Demonstration Steps

### Configure the DirectAccess client

1. Switch to LON-SVR3.
2. Open the Command Prompt window and type **gpupdate/force** to force apply Group Policy on LON-SVR3.
3. At command prompt, type **gpresult /R** to verify that the **DirectAccess Client Settings** GPO is applied to the Computer Settings.



**Note:** If **DirectAccess Client Settings** GPO is not applied, restart LON-SVR3, and then repeat step 2 on LON-SVR3.

4. Verify that **DNS Effective Name Resolution Policy Table Settings** is applied by typing the following command at the command prompt:

```
netsh name show effectivepolicy
```

5. Verify that **DNS Effective Name Resolution Policy Table Settings** is displayed in the Command Prompt window.
6. Simulate moving the client computer **LON-SVR3** out of the corporate network, that is to the Internet, by changing the network adapter settings with external IP address to the following values:
  - o IP address: **131.107.0.10**
  - o Subnet mask: **255.255.0.0**
  - o Default gateway: **131.107.0.2**
7. Disable and then again enable the **Local Area Connection** network adapter.
8. In Hyper-V Manager, right-click **20417A-LON-SVR3** and then click **Settings**. Change the **Legacy Network Adapter** to be on the **Private Network 2** network.

### Verify connectivity to the internal network resources

1. Move the mouse to the lower-left part of screen, click **Start**, and then click the **Internet Explorer** icon.
2. In the Address bar, type **http://lon-svr1.adatum.com** and then press Enter. The default IIS 8 web page for LON-SVR1 appears.
3. Leave the Internet Explorer window open.
4. Click **Start**, type **\\Lon-SVR1\Files**, and then press Enter. A folder window with the contents of the **Files** shared folder appears.
5. In the Files shared folder window, double-click the **example.txt** file. The content of the example.txt file is displayed.
6. Close all open windows.
7. Move the mouse pointer to the lower-right corner of the screen, and in the notification area, click **search**, and in the **search** box, type **cmd**.
8. At the command prompt, type **ipconfig**.
9. Notice the IP address for Tunnel adapter iphttpsinterface starts with **2002**. This is an IP-HTTPS address.

### Verify connectivity to the DirectAccess server

1. At the command prompt, type the following command:

```
Netsh name show effectivepolicy
```

Verify that **DNS Effective Name Resolution Policy Table Settings** present two entries for **adatum.com** and **Directaccess-NLS.Adatum.com**.

2. At the PowerShell prompt, type the following command, and then press Enter.

```
Get-DAClientExperienceConfiguration
```

Notice the DirectAccess client settings.

### Verify client connectivity on DirectAccess Server

1. Switch to LON-SVR2.
2. In the Remote Access Management console pane, click **Remote Client Status**.

Notice that **Client** is connected via **IPHttps**. In the Connection Details pane, in the bottom right of the screen, note the use of Kerberos for the Machine and the User.

3. Close all open programs.

**Question:** How will you configure IPv6 address for Windows 8 to use DirectAccess?

## Windows 7 Client vs. Windows 8 Client Implementation

Users working with DirectAccess in the Windows 8 operating system will have a better user experience than those working in Windows 7.

In Windows 8, the DirectAccess solution is completely transparent for the user. However, in Windows 7, it is hard to troubleshoot the network connectivity problems. Usually, when problems start, there are no native tools that can easily track the network behavior and so administrators often use network monitoring tools to get information regarding connectivity issues.

WINDOWS 7	WINDOWS 8
<ul style="list-style-type: none"> <li>• No tool from the client site for monitoring user interface for DirectAccess</li> <li>• Needs to be setup manually for selected site in multisite deployment</li> <li>• Needs certificate</li> </ul>	<ul style="list-style-type: none"> <li>• Includes an in-box user interface for DirectAccess troubleshooting</li> <li>• Automatically choose a site in multisite deployment</li> <li>• Can be used in deployments that does not require full PKI implementations</li> </ul>

### Windows 8 Client Implementation

- Windows 8 includes an in-box user interface for DirectAccess clients that help users understand network connectivity experience. Simplified user interface that run above the Windows PowerShell commands provide basic information regarding connectivity.
- Users can easily check their connectivity status. Users can even customize the look of the interface providing additional information such as support email addresses.
- Users might choose the site that they want to connect to in the multisite environment and even choose not to be connected to any site.



- Remediation options for actionable problems are presented clearly to the user. Instead of using other tools, remediation and problem solving can be done in the same user interface for DirectAccess. Typical problems that can be flagged for remediation are:
  - Credentials (Smartcard, TPM, and OTP)
  - NAP
  - Proxy authentication issue
  - Proxy configuration issue
  - Lack of Internet connectivity
- Users can easily send customized logs to their helpdesk by using the properties of Network Connectivity Assistance. Users can manually select the DirectAccess entry point that should be used. They can collect logs (HTML plus custom logs) and send these logs to already configured email addresses.
- When using Windows 7 in a multi-site deployment, you need to create multiple GPOs with different settings. However, in Windows 8, clients can easily select the closest DirectAccess server in a multisite deployment.
- Easy setup of DirectAccess automatically configures Windows 8 computers to participate in a DirectAccess scenario without the need for additional configuration.
- The receive side scaling concept for UDP traffic helps in improving performance in enterprise deployment.

## Lab: Implementing DirectAccess

### Scenario

Because A. Datum has expanded, many of the employees are now frequently out of the office, either working from home or traveling. A. Datum wants to implement a remote access solution for its employees so they can connect to the corporate network while they are away from the office. Although the VPN solution implemented with NAP provides a high level of security, business management is concerned about the complexity of the environment for end users. Also IT management is concerned that they are not able to manage the remote clients effectively.

To address these issues, A. Datum has decided to implement DirectAccess on client computers running Windows 8.

As a senior network administrator, you are required to deploy and validate the DirectAccess deployment. You will configure the DirectAccess environment and validate that the client computers can connect to the internal network when operating remotely.

### Objectives

After completing this lab, you will be able to:

- Configure the server infrastructure to deploy DirectAccess.
- Configure the DirectAccess clients.
- Validate the DirectAccess implementation.

### Lab Setup

Estimated time: **90 minutes**

Virtual Machine(s)	20417A-LON-DC1 20417A-LON-SVR1 20417A-LON-SVR2 20417A-LON-SVR3
User Name	Adatum\Administrator
Password	Pa\$\$w0rd

For this lab, you will use the available virtual machine environment. Before you begin the lab, you must complete the following steps:

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, click **20417A-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on using the following credentials:
  - a. User name: **Adatum\Administrator**
  - b. Password: **Pa\$\$w0rd**
5. Repeat steps 2-4 for **20417A-LON-SVR1**, **20417A-LON-SVR2**, and **20417A-LON-SVR3**.

## Exercise 1: Configuring the DirectAccess Infrastructure

### Scenario

You decided to implement DirectAccess as a solution for remote client computers that are not able to connect through VPN. Also, you want to address management problems, such as GPO application for remote client computers. For this purpose, you will configure the prerequisite components of DirectAccess, and configure the DirectAccess server.

The main tasks for this exercise are as follows:

1. Configure the AD DS and DNS requirements.
2. Configure certificate requirements.
3. Configure the internal resources for DirectAccess.
4. Configure DirectAccess server.

### ► Task 1: Configure the AD DS and DNS requirements

1. Create a security group for DirectAccess client computers by performing the following steps:
  - a. Switch to LON-DC1.
  - b. Open the Active Directory Users and Computers console, and create an Organizational Unit named **DA\_Clients OU**, and within that organizational unit, create a Global Security group named **DA\_Clients**.
  - c. Modify the membership of the **DA\_Clients** group to include **LON-SVR1**.
  - d. Close the Active Directory Users and Computers console.
2. Configure firewall rules for ICMPv6 traffic by performing the following steps:
  - a. Open the Group Policy Management console, and then open **Default Domain Policy**.
  - b. In the console tree of the Group Policy Management Editor, navigate to **Computer Configuration\Policies\Windows Settings\Security Settings\Windows Firewall with Advanced Security\Windows Firewall with Advanced Security**.
  - c. Create a new inbound rule with the following settings:
    - Rule Type: **Custom**
    - Protocol type: **ICMPv6**
    - Specific ICMP types: **Echo Request**
    - Name: **Inbound ICMPv6 Echo Requests**
  - d. Create a new outbound rule with the following settings:
    - Rule Type: **Custom**
    - Protocol type: **ICMPv6**
    - Specific ICMP types: **Echo Request**
    - Action: **Allow the connection**
    - Name: **Outbound ICMPv6 Echo Requests**
  - e. Close the Group Policy Management Editor and Group Policy Management consoles.

3. Create required DNS records by performing the following steps:
  - a. Open the DNS Manager console, and then create new host records with the following settings:
    - Name: **nls**; IP Address: **172.16.0.21**
    - Name: **crl**; IP Address: **172.16.0.22**
  - b. Close the DNS Manager console.
4. Remove ISATAP from the DNS global query block list by performing the following steps:
  - a. Open the Command Prompt window, type the following command, and then press Enter:
 

```
dnscmd /config /globalqueryblocklist wpad
```

Ensure that the **Command completed successfully** message appears.
  - b. Close the Command Prompt window.
5. Configure the DNS suffix on LON-SVR2 by performing the following steps:
  - a. Switch to LON-SVR2, and in the **Local Area Connection Properties** dialog box, in the **Internet Protocol Version 4 (TCP/IPv4)** dialog box, add the **Adatum.com** DNS suffix.
  - b. Close the Local Area Connection Properties dialog box.

## ► Task 2: Configure certificate requirements

1. Configure the CRL distribution settings by performing the following steps:
  - a. Switch to LON-DC1 and open the Certification Authority console.
  - b. Configure **Adatum-LON-DC1-CA** certification authority with the following extension settings:
    - Add Location: **http://crl.adatum.com/crld/**
    - Variable: **CAName, CRLNameSuffix, DeltaCRLAllowed**
    - Location: **.crl**
    - Select **Include in CRLs. Clients use this to find Delta CRL locations** and **Include in the CDP extension of issued certificates**
    - Do not restart Certificate Services.
    - Add Location: **\\lon-svr2\crldist\$\**
    - Variable: **CAName, CRLNameSuffix, DeltaCRLAllowed**
    - Location: **.crl**
    - Select **Include in CRLs. Clients use this to find Delta CRL locations** and **Include in the CDP extension of issued certificates**
    - Restart Certificate Services.
    - Close the Certificate Authority console.
2. To duplicate the web certificate template and configure appropriate permission by performing the following steps:
  - a. In the Certificate Templates console, in the contents pane, duplicate the **Web Server** template by using the following options:
    - Template display name: **Adatum Web Server Certificate**
    - Request Handling: **Allow private key to be exported**
    - Authenticated Users permissions: under **Allow**, click **Enroll**

- b. Close the Certificate Templates console.
  - c. In the Certification Authority console, choose to issue a New Certificate Template and select the **Adatum Web Server Certificate** template.
  - d. Close the Certification Authority console.
3. Configure computer certificate auto-enrollment by performing the following steps:
  - a. On LON-DC1, open the Group Policy Management console.
  - b. In the console tree, navigate to **Forest: Adatum.com, Domains, and Adatum.com**.
  - c. Edit the Default Domain Policy and in the console tree of the Group Policy Management Editor, navigate to **Computer Configuration\Policies\Windows Settings\Security Settings\Public Key Policies**.
  - d. Under **Automatic Certificate Request Settings**, configure **Automatic Certificate Request** to issue the **Computer** certificate.
  - e. Close the Group Policy Management Editor and close the Group Policy Management console.

► **Task 3: Configure the internal resources for DirectAccess**

1. To request a certificate for LON-SVR1 by performing the following steps:
  - a. On LON-SVR1, open a command prompt, type the following command, and then press Enter.

```
gpupdate /force
```
  - b. At the command prompt, type the following command, and then press Enter.

```
mmc
```
  - c. Add the **Certificates** snap-in for **Local computer**.
  - d. In the console tree of the Certificates snap-in, navigate to Certificates (Local Computer) \Personal\Certificates, request a new certificate, and then under Request Certificates, select Adatum Web Server Certificate with the following setting:
    - Subject name: Under **Common name**, type **nls.adatum.com**
  - e. In the details pane of the Certificates snap-in, verify that a new certificate with the name **nls.adatum.com** was enrolled with **Intended Purposes** of **Server Authentication**.
  - f. Close the console window. When you are prompted to save settings, click **No**.
2. To change the HTTPS bindings, perform the following steps:
  - a. Open Internet Information Services (IIS) Manager.
  - b. In the console tree of Internet Information Services (IIS), navigate to and click **Default Web site**.
  - c. Configure Site Bindings by selecting **nls.adatum.com** for **SSL Certificate**.
  - d. Close the Internet Information Services (IIS) Manager console.

► **Task 4: Configure DirectAccess server.**


1. Obtain required certificates for LON-SVR2 by performing the following steps:
  - a. Switch to LON-SVR2.
  - b. Open a command prompt and refresh group policy by typing **gpupdate /force**.

- c. Open Microsoft Management Console by typing **mmc** command, and then add the **Certificates** snap-in for **Local computer**.
  - d. In the **Certificates** snap-in, in the mmc console, request a new certificate with the following settings:
    - Certificate template: Adatum Web Server Certificate
    - Common name: 131.107.0.2
    - Friendly name: IP-HTTPS Certificate
  - e. Close the console.
2. Create CRL distribution point on LON-SVR2 by performing the following steps:
    - a. Switch to Server Manager
    - b. In Internet Information Services (IIS) Manager, create new virtual directory CRLD and assign c:\crl-dist as a home directory.
  3. Share and secure the CRL distribution point by performing the following step:



**Note:** You perform this step to assign permissions to the CRL distribution point. In the details pane of Windows Explorer, right-click the **CRLDist** folder, and then click **Properties**, and grant Full Share and NTFS permission.

4. Publish the CRL to LON-SVR2 by performing the following steps:
 



**Note:** This step makes the CRL available on the edge server for Internet-based DirectAccess clients.

  - a. Switch to LON-DC1.
  - b. Start the **Certification Authority** console.
  - c. In the console tree, open **ADATUMCA**, right-click **Revoked Certificates**, point to **All Tasks**, and then click **Publish**.
5. Complete DirectAccess setup wizard on LON-SVR2 by performing the following steps:
  - a. On LON-SVR2, open the **Server Manager** console.
  - b. In the Server Manager console, start the **Remote Access Management** console, click **Configuration**, and start the **Enable Direct Access Wizard** with following settings:
    - Select Groups: **DA\_Clients**
    - Network Topology: **Edge** is selected, and verify that **131.107.0.2** is used by clients to connect to the Remote Access server.
    - **Infrastructure Server Setup** page, click **Next**
    - **Configure Remote Access** page, click **Next**
    - In **Summary**, click **Finish**, to apply DirectAccess Settings



**Note:** Since the server you already configured is a VPN server, you can only use the getting started wizard which generates self-signed certificate for DirectAccess communication. Next steps will modify default DirectAccess settings to include already deployed certificates from the internal Certification Authority.

- c. In the details pane of the Remote Access Management console, under **Step 2**, click **Edit**.

- d. On the **Network Topology** page, verify that **Edge** is selected, and type **131.107.0.2**.
  - e. On the **Network Adapters** page, verify that **CN=131.107.0.2** is used as a certificate to authenticate IP-HTTPS connection.
  - f. On the **Authentication** page, select **Use computer certificates**, click **Browse**, and then select **Adatum Lon-Dc1 CA**.
  - g. On the VPN Configuration page, click **Finish**.
  - h. In details pane of the Remote Access Management console, under **Step 3**, click **Edit**.
  - i. On the **Network Location Server** page, select the **The network location server is deployed on a remote web server (recommended)** and in the URL of the NLS, type **https://nls.adatum.com**, and then click **Validate**.
  - j. Ensure that URL is validated.
  - k. On the **DNS** page, examine the values, and then click **Next**.
  - l. In the **DNS Suffix Search List**, select **Next**.
  - m. On the **Management** page, click **Finish**.
  - n. In details pane of the Remote Access Management console, review the setting for **Step 4**.
  - o. In **Remote Access Review**, click **Apply**.
  - p. Under **Applying Remote Access Setup Wizard Settings**, click **Close**.
6. Update Group Policy settings on LON-SVR2 by performing the following step:
- Open the command prompt, and type the following commands:

```
gpupdate /force
Ipconfig
```



**Note:** Verify that **LON-SVR2** has an IPv6 address for **Tunnel adapter IPHTTPSInterface** starting with **2002**.

**Results:** After completing this exercise, you will have configured the DirectAccess infrastructure.

## Exercise 2: Configuring the DirectAccess Clients

### Scenario

After you configured the DirectAccess server and the required infrastructure, you must configure DirectAccess clients. You decide to use Group Policy mechanism to apply DirectAccess settings to the clients and for certificate distribution.

The main tasks for this exercise are as follows:

1. Configure Group Policy to configure client settings for DirectAccess.
2. Verify client computer certificate distribution.
3. Verify IP address configuration.

► **Task 1: Configure Group Policy to configure client settings for DirectAccess.**

1. Switch to LON-SVR3.
2. Restart LON-SVR3 and then log back on as **Adatum\Administrator** with the password of **Pa\$\$w0rd**. Open the Command Prompt window and then type the following commands:

```
gpupdate /force  
gpresult /R
```

3. Verify that **DirectAccess Client Settings GPO** is displayed in the list of the Applied Policy objects for the Computer Settings.

► **Task 2: Verify client computer certificate distribution.**

1. On LON-SVR3, open the **Certificates** MMC.
2. Verify that a certificate with the name **LON-SVR3.adatum.com** is present with **Intended Purposes of Client Authentication and Server Authentication**.
3. Close the console window without saving it.

**Question:** Why did you install a certificate on the client computer?

► **Task 3: Verify IP address configuration.**

1. On LON-SVR3, open Internet Explorer and go to **http://lon-svr1.adatum.com/**. The default IIS 8 web page for LON-SVR1 appears.
2. In Internet Explorer, go to **https://nls.adatum.com/**. The default IIS 8 web page for LON-SVR1 appears.
3. Open Windows Explorer, and type **\\Lon-SVR1\Files**, and then press Enter. You should see a folder window with the contents of the **Files** shared folder.
4. Close all open windows.

**Results:** After completing this exercise, you will have configured the DirectAccess clients.

## Exercise 3: Verifying the DirectAccess Configuration

### Scenario

When client configuration is completed, it is important to verify that DirectAccess works. You do this by moving the DirectAccess client to the Internet and trying to access internal resources.

The main tasks for this exercise are as follows:

1. Move the client computer to the Internet virtual network.
2. Verify connectivity to the DirectAccess server.
3. Verify connectivity to the internal network resources.



### ► Task 1: Move the client computer to the Internet virtual network



**Note:** To verify the DirectAccess functionality, you must move the client computer to the Internet.

1. Switch to LON-SVR3.
2. Change the network adapter configuration with the following settings:
  - IP address: **131.107.0.10**
  - Subnet mask: **255.255.0.0**
  - Default gateway: **131.107.0.2**
3. Disable and then again enable the **Local Area Network** network adapter.
4. Close the Network Connections window.
5. In Hyper-V Manager, right-click **20417A-LON-SVR3** and then click **Settings**. Change the Legacy Network Adapter to be on the **Private Network 2** network. Click **OK**.

### ► Task 2: Verify connectivity to the DirectAccess server

1. On LON-SVR3, open a command prompt, and type the following command:

```
ipconfig
```

2. Notice the IP address that starts with **2002**. This is IP-HTTPS address.
3. At the command prompt, type the following command, and then press Enter.

```
Netsh name show effectivepolicy
```

4. At the command prompt, type the following command, and then press Enter.

```
powershell
```

5. At the Windows PowerShell command prompt, type the following command, and then press Enter.

```
Get-DAClientExperienceConfiguration
```

### ► Task 3: Verify connectivity to the internal network resources

1. Open Internet Explorer and go to **http://lon-svr1.adatum.com/**. You should see the default IIS 8 web page for LON-SVR1.
2. Open Windows Explorer, type **\\LON-SVR1\Files**, and then press Enter.
3. You should see a folder window with the contents of the Files shared folder.
4. At the command prompt, type the following command:

```
ping lon-dc1.adatum.com
```

Verify that you are receiving replies from lon-dc1.adatum.com.

5. At the command prompt, type the following command, and then press Enter.

```
gpupdate /force
```

6. Close all open windows.

7. Switch to LON-SVR2.
8. Start the **Remote Access Management** console and review the information on **Remote Client Status**.



**Note:** Notice that LON-SVR3 is connected via IPHttps. In the Connection Details pane, in the bottom-right of the screen, note the use of Kerberos for the Machine and the User.

9. Close all open windows.

**Results:** After completing this exercise, you will have verified the DirectAccess configuration.

#### ► To prepare for the next module

When you finish the lab, revert the virtual machines to their initial state. To do this, perform the following steps:

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20410A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20410A-LON-SVR1**, **20410A-LON-SVR2**, and **20410A-LON-SVR3**.

## Module Review and Takeaways

### Review Questions

**Question:** What are the main benefits of using DirectAccess for providing remote connectivity?

**Question:** How do you configure a DirectAccess server?

**Question:** How do you configure DirectAccess clients?

**Question:** How does the DirectAccess client determine if it is connected to the intranet or the Internet?

**Question:** What is the use of an NRPT?

### Best Practices

Although DirectAccess was present in previous Windows 7 and Windows 2008 R2 edition, Windows 8 introduces new features for improved manageability, ease of deployment, and improved scale and performance.

Monitoring of the environment is now much easier with support of PowerShell, Windows Management Instrumentation (WMI), GUI monitoring, along with Network Connectivity Assistant on the client side.

One of the best enhancements is that DirectAccess can now access IP4 servers on your network and your servers do not need to have IP6 addresses to be exposed through DirectAccess, because your DirectAccess server acts as a proxy.

For ease of deployment you do not need to have IP addresses on the Internet-facing network. Therefore, this is a good scenario for proof of concept. However, if you are concerned about security and if you want to integrate with NAP, you still need two public addresses.

Consider integrating DirectAccess with your existing Remote Access solution because Windows Server 2012 can implement DirectAccess server behind the NAT device which is the most common Remote Access Server (RAS) solution for companies.

### Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
You have configured DirectAccess, but users are complaining about connectivity issues. You want to troubleshoot those issues more efficiently.	
The DirectAccess client tries to connect to the DirectAccess server by using IPv6 and IPsec with no success.	

### Real-world Issues and Scenarios

You are considering implementing DirectAccess in your organization. You are planning to implement Windows Server 2012 servers. What are the other considerations that you should be aware of?

**Tools**

Tool	Use for	Where to find it
Express Setup, Remote Access Configuration	A graphical tool that simplifies the configuration of DirectAccess	Server Manager/Tools
Dnscmd.exe	A command-line tool used for DNS management	Run from command-line
Services.msc	Helps in managing Windows services	Server Manager/Tools
Gpedit.msc	Helps in editing the Local Group Policy	Run from command-line
IPconfig.exe	A command-line tool that displays current TCP/IP network configuration	Run from command-line
DNS Manager console	Helps in configuring name resolution	Server Manager/Tools
Mmc.exe	Helps in the creation and management of the Management Console	Run from command-line
Gpupdate.exe	Helps in managing Group Policy application	Run from command-line
Active Directory Users and Computers	Is useful in configuring group membership for client computers that will be configured with DirectAccess	Server Manager/Tools

# Module 7

## Implementing Failover Clustering

### Contents:

Module Overview	7-1
<b>Lesson 1:</b> Overview of Failover Clustering	7-2
<b>Lesson 2:</b> Implementing a Failover Cluster	7-13
<b>Lesson 3:</b> Configuring Highly-Available Applications and Services on a Failover Cluster	7-18
<b>Lesson 4:</b> Maintaining a Failover Cluster	7-22
<b>Lesson 5:</b> Implementing a Multi-Site Failover Cluster	7-27
<b>Lab:</b> Implementing Failover Clustering	7-32
Module Review and Takeaways	7-37

## Module Overview

Providing high availability is very important for any organization that wants to provide continuous services to its users. Failover Clustering is one of the main technologies in Windows Server® 2012 that can provide high availability for various applications and services. In this module, you will learn about Failover Clustering, Failover Clustering components, and implementation techniques.

### Objectives

After completing this module, you will be able to:

- Describe Failover Clustering.
- Implement a failover cluster.
- Configure highly-available applications and services.
- Maintain a failover cluster.
- Implement multi-site Failover Clustering.

## Lesson 1

# Overview of Failover Clustering

Failover clusters in Windows Server 2012 provide a high-availability solution for many server roles and applications. By implementing failover clusters, you can maintain application or service availability if one or more computers in the failover cluster fail. Before you implement Failover Clustering, you should be familiar with general high-availability concepts. You must understand clustering terminology and also how failover clusters work.

Also, it is important to be familiar with new clustering features in Windows Server 2012.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe availability.
- Describe Failover Clustering improvements in Windows Server 2012.
- Describe failover cluster components.
- Define failover and failback.
- Describe failover cluster networks.
- Describe failover cluster storage.
- Describe a quorum.
- Describe quorum modes.
- Describe Cluster Shared Volumes (CSVs).

### What Is Availability?

Availability refers to a level of service that applications, services, or systems provide, and is expressed as the percentage of time that a service or system is available. Highly-available systems have minimal downtime—whether planned or unplanned—and are available more than 99 percent of the time, depending on the needs and the budget of the organization. For example, a system that is unavailable for 8.75 hours per year would have a 99.9 percent availability rating.

To improve availability, you must implement fault-tolerance mechanisms that mask or minimize how failures of the service's components and dependencies affect the system. You can achieve fault tolerance by implementing redundancy to single points of failure.

Availability requirements must be expressed so that there are no misunderstandings about the implications. Miscommunication about service level expectations between the customer and the IT organization can result in poor business decisions, such as unsuitable investment levels and customer dissatisfaction.

- Availability is a level of service expressed as a percentage of time
- Highly-available services or systems are available more than 99 percent of the time
- High availability requirements differ based on how availability is measured
- Planned outages typically are not included when calculating availability



The availability measurement period can also have a significant effect on the definition of availability. For example, a requirement for 99.9 percent availability over a one-year period allows for 8.75 hours of downtime, whereas a requirement for 99.9 percent availability over a rolling four-week window allows for only 40 minutes of downtime per period.

You also have to identify and negotiate planned outages maintenance activities, service pack updates, and software updates. These are scheduled outages, and typically are not included as downtime when calculating the system's availability. You typically calculate availability based on unplanned outages only. However, you have to negotiate exactly which planned outages you consider as downtime.

## Failover Clustering Improvements in Windows Server 2012

Failover Clustering has not significantly changed since Windows Server 2008 R2. However, there are some new features and technologies in Windows Server 2012 that help increase scalability and cluster storage availability, and provide better and easier management and faster failover.

The important new features in Windows Server 2012 Failover Clustering include:

- Increased scalability.** In Windows Server 2012, failover cluster can have 64 physical nodes and can run 4,000 virtual machines on each cluster. This is a significant improvement over Windows Server 2008 R2 which supports only 16 physical nodes and 1,000 virtual machines per cluster. Each cluster you create is now available from Server Manager console. Server Manager in Windows Server 2012 can discover and manage all clusters created in an Active Directory® Domain Services (AD DS) domain. If the cluster is deployed in multi-site scenario, the administrator can now control which nodes in a cluster have votes for establishing quorum. Failover Clustering scalability is also improved for virtual machines that are running on clusters. This will be discussed in more detail in *Module 8: Implementing Hyper-V*.
- Improved Cluster Shared Volumes (CSVs) volumes.** This technology was introduced in Windows Server 2008 R2, and it became very popular for providing virtual machine storage. In Windows Server 2012, CSV volumes appear as CSV File System and it supports server message block (SMB) version 2.2 storage for Hyper-V and other applications. Also, CSV can use SMB multichannel and SMB Direct to enable traffic to stream across multiple networks in a cluster. For additional security, you can use BitLocker Drive Encryption for CSV disks, and you can also make CSV storage visible only to a subset of nodes in a cluster. For reliability, CSV volumes can be scanned and repaired with zero offline time.
- Cluster-aware updating.** Updating cluster nodes required a lot of preparation and planning in earlier versions of Windows Server, to minimize or avoid downtime. Also, procedure of updating cluster nodes was mostly manual, which caused additional administrative effort. In Windows Server 2012, a new technology is introduced for this purpose. This technology is called Cluster-Aware Updating. This technology automatically updates cluster nodes with Windows Update hotfix, by keeping the cluster online, and minimizing downtime. This technology will be explained in more detail in *Lesson 4: Maintaining a Failover Cluster*.
- Active Directory® integration improvements.** Because Windows Server 2008, Failover Clustering is integrated in Active Directory Domain Services (AD DS). In Windows Server 2012, this integration is improved. Administrators can create cluster computer objects in targeted organizational units (OUs), or by default in the same OUs as the cluster nodes. This aligns failover cluster dependencies on AD DS

The improvements in Failover Clustering in Windows Server 2012 include:

- Increased scalability
- Improved CSVs
- Cluster-Aware Updating
- Active Directory integration improvements
- Management improvements

Cluster.exe command-line tool, Cluster Automation Server (MSCluster) COM interface, and Add-ClusterPrintServerRole cmdlet are some of the removed features in Windows Server 2012

with the delegated domain administration model that is used in many IT organizations. Also, now failover clusters can be deployed with access only to read-only domain controllers.

- **Management improvements.** Although Failover Clustering in Windows Server 2012 still uses almost the same management console and the same administrative techniques, it brings some important management improvements. Validation wizard is improved in which the validation speed for large failover clusters is improved and new tests for CSVs, the Hyper-V role, and virtual machines are added. Also, new Windows PowerShell cmdlets are available for managing clusters, monitoring clustered virtual machine applications, and creating highly available iSCSI target.

## Removed and Deprecated Features

In Windows Server 2012 clustering, some features are removed or deprecated. If you are moving from an older version of Failover Clustering, you should be aware of these features:


- The Cluster.exe command-line tool is deprecated. However, it can be optionally installed with the Failover Clustering Tools. Failover Clustering Windows PowerShell cmdlets provide a functionality that is generally the same as Cluster.exe commands.
- The Cluster Automation Server (MSClus) COM interface is deprecated, but it can be optionally installed with the Failover Clustering Tools.
- The Support for 32-bit cluster resource DLLs is deprecated, but 32-bit DLLs can be optionally installed. Cluster resource DLLs should be updated to 64 bit.
- The Print Server role is removed from the High Availability Wizard, and it cannot be configured in Failover Cluster Manager.
- The Add-ClusterPrintServerRole cmdlet is deprecated, and it is not supported in Windows Server 2012.

## Failover Cluster Components

A *failover cluster* is a group of independent computers that work together to increase the availability of applications and services. Physical cables and software connect the clustered servers, known as *nodes*. If one of the cluster nodes fails, another node begins to provide service. This process is known as *failover*. With failover, users experience a minimum of service disruptions.

A Failover Clustering solution consists of several components, which include:

- **Nodes.** These are computers that are members of a failover cluster. These computers run cluster service and resources and applications associated to cluster.
- **Network.** This is a network across which cluster nodes can communicate with one another and with *clients*. There are three types of networks that can be used in a cluster. These networks are discussed in more detail in the "Failover Cluster Networks" section.
- **Resource.** This is an entity that is hosted by a node. It is managed by the Cluster service the Cluster service and can be started, stopped, and moved to another node.





- *Cluster storage.* This is a storage system that is usually shared between cluster nodes. In some scenarios, such as clusters of servers running Microsoft® Exchange Server, shared storage is not required.
- *Clients.* These are computers (or users) that are using the Cluster service.
- *Service or application.* This is a software entity that is presented to clients and used by clients.
- *Witness.* This can be a file share or disk which is used to maintain quorum. Ideally the witness should be located a network that is both logically and physically separate from those used by the failover cluster. However, the witness must remain accessible by all cluster node members. The concepts of quorum and how the witness comes into play will be examined more closely in the coming lessons of this module.

In a failover cluster, each node in the cluster:

- Has full connectivity and communication with the other nodes in the cluster.
- Is aware when another node joins or leaves the cluster.
- Is connected to a network through which client computers can access the cluster.
- Is connected through a shared bus or iSCSI connection to shared storage.
- Is aware of the services or applications that are running locally, and the resources that are running on all other cluster nodes.

Cluster storage usually refers to logical devices—typically hard disk drives or logical unit numbers (LUN)—that all the cluster nodes attach to, through a shared bus. This bus is separate from the bus that contains the system and boot disks. The shared disks store resources such as applications and file shares that the cluster will manage.

A failover cluster typically defines at least two data communications networks: one network enables the cluster to communicate with clients, and the second, isolated network enables the cluster node members to communicate directly with one another. If a directly-connected shared storage is not being used, then a third network segment (for iSCSI or Fibre Channel) can exist between the cluster nodes and a data storage network.

Most clustered applications and their associated resources are assigned to one cluster node at a time. The node that provides access to those cluster resources is the active node. If the nodes detect the failure of the active node for a clustered application, or if the active node is taken offline for maintenance, the clustered application is started on another cluster node. To minimize the impact of the failure, client requests are immediately and transparently redirected to the new cluster node.

## What Are Failover and Failback?

Failover transfers the responsibility of providing access to resources in a cluster from one node to another. Failover can occur when an administrator intentionally moves resources to another node for maintenance, or when unplanned downtime of one node happens because of hardware failure or other reasons. Also, service failure on an active node can initiate failover to another node.

- During failover, the clustered instance and all associated resources are moved from one node to another
- Failover occurs when:
  - The node that currently hosts the instance becomes inactive for any reason
  - One of the resources within the instance fails
  - An administrator forces a failover
- Cluster service can failback after the offline node becomes active again

A failover attempt consists of the following steps:

1. The Cluster service takes all the resources in the instance offline in an order that is determined by the instance's dependency hierarchy. That is, dependent resources first, followed by the resources on which they depend. For example, if an application depends on a physical disk resource, the Cluster service takes the application offline first, which enables the application to write changes to the disk before the disk is taken offline.
2. After all the resources are offline, the Cluster service attempts to transfer the instance to the node that is listed next on the instance's list of preferred owners.
3. If the Cluster service successfully moves the instance to another node, it attempts to bring all the resources online. This time, it starts at the lowermost part of the dependency hierarchy. Failover is complete when all the resources are online on the new node.

The Cluster service can *failback* instances that were originally hosted on the offline node, after the offline node becomes active again. When the Cluster service fails back an instance, it uses the same procedures that it performs during failover. That is, the Cluster service takes all the resources in the instance offline, moves the instance, and then brings all the resources in the instance back online.

## Failover Cluster Networks

Network and network adapters are important parts of each cluster implementation. You cannot configure a cluster without configuring the networks that the cluster will use. A network can perform one of the following roles in a cluster:

- **Private network.** A private network carries internal cluster communication. By using this network, cluster nodes exchange heartbeats and check for another node or nodes. The failover cluster authenticates all internal communication. However, administrators who are especially concerned about security may want to restrict internal communication to physically secure networks.
- **Public network.** A public network provides client systems with access to cluster application services. IP address resources are created on networks that provide clients with access to the Cluster service.
- **Public-and-private network.** A public-and-private network (also known as a mixed network) carries internal cluster communication and connects clients to cluster application services.

Network	Description
Public network	Clients use this network to connect to the clustered service
Private network	Nodes use this network to communicate with each other
Public-and-private network	Required to communicate with external storage systems

- One network can support both client and node communications
- Multiple network cards are recommended to provide enhanced performance and redundancy

When you configure networks in failover clusters, you must also dedicate a network to connect to the shared storage. If you use iSCSI for the shared storage connection, the network will use an IP-based Ethernet communications network. However, you should not use this network for node or client communication. Sharing the iSCSI network in this manner may result in contention and latency issues for both users and for the resource that is being provided by the cluster.

Though not a best practice, you can use the private and public networks for both client and node communications. Preferably, you should dedicate an isolated network for the private node communication. The reasoning for this is similar using a separate Ethernet network for iSCSI – namely to avoid issues resource bottleneck and contention issues. The public network is configured to allow client connections to the failover cluster. Although the public network can provide backup for the private network, a better design practice is to define alternative networks for the primary private and public networks or at least team the network interfaces used for these networks.

The networking features in Windows Server 2012–based clusters include the following:

- The nodes transmit and receive heartbeats by using User Datagram Protocol (UDP) unicast, instead of UDP broadcast (which was used in legacy clusters). The messages are sent on port 3343.
- You can include clustered servers on different IP subnets, which reduces the complexity of setting up multi-site clusters.
- The Failover Cluster Virtual Adapter is a hidden device that is added to each node when you install the Failover Clustering feature. The adapter is assigned a media access control (MAC) address based on the MAC address that is associated with the first enumerated physical network adapter in the node.
- Failover clusters fully support IPv6 for both node-to-node and node-to-client communication.
- You can use Dynamic Host Configuration Protocol (DHCP) to assign IP addresses, or assign static IP addresses to all nodes in the cluster. However, if some nodes have static IP addresses and you configure others to use DHCP, the Validate a Configuration Wizard will raise an error. The cluster IP address resources are obtained based on the configuration of the network interface supporting that cluster network.

## Failover Cluster Storage

Most Failover Clustering scenarios require shared storage to provide consistent data to a highly-available service or application after failover. There are three shared storage options for a failover cluster:

- *Shared serial attached SCSI (SAS)*. Shared serial attached SAS is the lowest cost option. However, it is not very flexible for deployment because the two cluster nodes must be physically close together. In addition, the shared storage devices that are supporting SAS have a limited number of connections for cluster nodes.
- *Internet SCSI (iSCSI)*. iSCSI is a type of storage area network (SAN) that transmits SCSI commands over IP networks. Performance is acceptable for most scenarios when 1 gigabit per second (Gbps) or 10 Gbps Ethernet is used as the physical medium for data transmission. This type of SAN is fairly inexpensive to implement because no specialized networking hardware is required. In Windows Server 2012, you can implement iSCSI target software on any server, and present local storage over iSCSI interface to clients.
- *Fibre channel*. Fibre channel SANs typically have better performance than iSCSI SANs, but are much more expensive. Specialized knowledge and hardware are required to implement a fibre channel SAN.

- Failover clusters require shared storage to provide consistent data to a virtual server after failover
- Shared storage options include:
  - SAS
  - iSCSI
  - Fibre channel



**Note:** The Microsoft iSCSI Software Target is now an integrated feature in Windows Server 2012. It can provide storage from a server over a TCP/IP network, including shared storage for applications that are hosted in a failover cluster. Also, in Windows Server 2012, a highly-available iSCSI Target Server can be configured as a clustered role by using Failover Cluster Manager or Windows PowerShell®.

## Storage Requirements

After you choose the type of storage, you should also be aware of the following storage requirements:

- To use the native disk support included in Failover Clustering, use basic disks and not dynamic disks.
- We recommend that you format the partitions with NTFS. For the disk witness, the partition must be NTFS, because FAT is not supported.
- For the partition style of the disk, you can use either master boot record (MBR) or GUID partition table (GPT).
- Because improvements in failover clusters require that the storage respond correctly to specific SCSI commands, the storage must follow the SCSI Primary Commands-3 (SPC-3) standard. In particular, the storage must support Persistent Reservations, as specified in the SPC-3 standard.
- The miniport driver used for the storage must work with the Microsoft Storport storage driver. Storport offers a higher performance architecture and better Fiber Channel compatibility in Windows systems.
- You must isolate storage devices. That is, one cluster per device. Servers from different clusters must be unable to access the same storage devices. In most cases, a logical unit number (LUN) that is used for one set of cluster servers should be isolated from all other servers through LUN masking or zoning.
- Consider using multipath I/O software. In a highly-available storage fabric, you can deploy failover clusters with multiple host bus adapters by using multipath I/O software. This provides the highest level of redundancy and availability. For Windows Server 2012, your multipath solution must be based on Microsoft Multipath I/O (MPIO). Your hardware vendor usually supplies an MPIO device-specific module (DSM) for your hardware, although Windows Server 2012 includes one or more DSMs as part of the operating system.

## What Is Quorum?

*Quorum* is the number of elements that must be online for a cluster to continue running. In effect, each element can cast one *vote* to determine whether the cluster continues to run. Each cluster node is an element that has one vote. In case, there is an even number of nodes, then an additional element, which is known as a *witness* is assigned to the cluster. The witness element can be either a disk or a file share. Each voting element contains a copy of the cluster configuration; and the Cluster service works to keep all copies synchronized at all times.

- In failover clusters, quorum defines the consensus that enough cluster members are available to provide services
- Quorum:
  - Is based on votes in Windows Server 2008
  - Allows nodes, file shares, or a shared disk to have a vote, depending on the quorum mode
  - Allows the failover cluster to remain online when sufficient votes are available

The cluster will stop providing failover protection if most of the nodes fail or if there is a problem with communication between the cluster nodes. Without a quorum mechanism, each set of nodes could continue to operate as a failover cluster. This results in a partition within the cluster. Quorum prevents two or more nodes from concurrently operating a failover cluster resource. If a clear majority is not achieved between the node members, then the vote of the witness becomes crucial to maintain the validity of the cluster. Concurrent operation could occur when network problems prevent one set of nodes from communicating with another set of nodes. That is, a situation might occur where more than one node tries to control access to a resource. If that resource is, for example, a database application, damage could result. Imagine the consequence if two or more instances of the same database are made available on the

network, or if data was accessed and written to a target from more than one source at a time. If the application itself is not damaged, the data could easily become corrupted.

Because a given cluster has a specific set of nodes and a specific quorum configuration, the cluster can calculate the number of votes that are required for the cluster to continue providing failover protection. If the number of votes drops below the majority, the cluster stops running. That is, it will not provide failover protection if there is a node failure. Nodes will still listen for the presence of other nodes, in case another node appears again on the network, but the nodes will not function as a cluster until a majority consensus or quorum is achieved.



**Note:** The full functioning of a cluster depends not just on quorum, but on the capacity of each node to support the services and applications that fail over to that node. For example, a cluster that has five nodes could still have quorum after two nodes fail, but each remaining cluster node would continue serving clients only if it has enough capacity (such as disk space, processing power, network bandwidth, RAM) to support the services and applications that failed over to it. An important part of the design process is planning each node's failover capacity. A failover node must be able to run its own load and also the load of additional resources that might failover to it.

### The Process of Achieving Quorum

Because a given cluster has a specific set of nodes and a specific quorum configuration, the cluster software on each node stores information about how many votes constitute a quorum for that cluster. If the number drops below the majority, the cluster stops providing services. Nodes will continue listening for incoming connections from other nodes on port 3343, in case they appear again on the network, but the nodes will not begin to function as a cluster until quorum is achieved.

There are several phases a cluster must complete to achieve quorum. As a given node comes up, it determines whether there are other cluster members that can be communicated with. This process may be in progress on multiple nodes at the same time. After communication is established with other members, the members compare their membership "views" of the cluster until they agree on one view (based on timestamps and other information). A determination is made whether this collection of members "has quorum;" or has enough members the total of which creates sufficient votes so that a "split" scenario cannot exist. A "split" scenario means that another set of nodes that are in this cluster are running on a part of the network inaccessible to these nodes. Therefore, more than one node could be actively trying to provide access to the same clustered resource. If there are not enough votes to achieve quorum, the voters (the currently recognized members of the cluster) wait for more members to appear. After at least the minimum vote total is attained, the Cluster service the Cluster service begins to bring cluster resources and applications into service. With quorum attained, the cluster becomes fully functional.

## Quorum Modes in Windows Server 2012 Failover Clustering

Same quorum modes from Windows Server 2008 are also present in Windows Server 2012. As before, a majority of votes determines whether a cluster achieves quorum. Nodes can vote, and where appropriate, either a disk in cluster storage (known as a *disk witness*) or a file share (known as a *file share witness*) can vote. There is also a quorum mode called No Majority: Disk Only, which functions like the disk-based quorum in Windows Server 2003. Other than that mode, there is no single point of failure with the quorum modes, because only the number of votes is important and not whether a particular element is available to vote.

Quorum Mode	What Has the Vote?	When Is Quorum Maintained?
Node Majority	Only nodes in the cluster have a vote	Quorum is maintained when more than half of the nodes are online
Node and Disk Majority	The nodes in the cluster and a disk witness have a vote	Quorum is maintained when more than half of the votes are online
Node and File Share Majority	The nodes in the cluster and a file share witness have a vote	Quorum is maintained when more than half of the votes are online
No Majority: Disk Only	Only the quorum-shared disk has a vote	Quorum is maintained when the shared disk is online

This quorum mode is flexible. You can choose the mode best suited to your cluster.

Be aware that, most of the time, it is best to use the quorum mode selected by the cluster software. If you run the Quorum Configuration Wizard, the quorum mode that the wizard lists as “recommended” is the quorum mode chosen by the cluster software. We recommend changing the quorum configuration only if you have determined that the change is appropriate for your cluster.

There are four quorum modes:

- **Node Majority.** Each node that is available and in communication can vote. The cluster functions only with a majority of the votes. That is, more than half. This model is preferred when the cluster consists of an odd number of server nodes (no witness is needed to maintain or achieve quorum).
- **Node and Disk Majority.** Each node plus a designated disk in the cluster storage, the disk witness, can vote, when they are available and in communication. The cluster functions only with a majority of the votes. That is, more than half. This model is based on an even number of server nodes being able to communicate with one another in the cluster in addition to the disk witness.
- **Node and File Share Majority.** Each node plus a designated file share created by the administrator, which is the file share witness, can vote when they are available and in communication. The cluster functions only with a majority of the votes. That is, more than half. This model is based on an even number of server nodes being able to communicate with one another in the cluster, in addition to the file share witness.
- **No Majority: Disk Only.** The cluster has quorum if one node is available and in communication with a specific disk in the cluster storage. Only the nodes that are also in communication with that disk can join the cluster.

Except for the No Majority: Disk Only mode, all quorum modes in Windows Server 2012 failover clusters are based on a simple majority vote model. As long as a majority of the votes are available, the cluster continues to function. For example, if there are five votes in the cluster, the cluster continues to function as long as there are at least three available votes. The source of the votes is not relevant—the vote could be a node, a disk witness, or a file share witness. The cluster will stop functioning if a majority of votes is not available.

In the No Majority: Disk Only mode, the quorum-shared disk can veto all other possible votes. In this mode, the cluster will continue to function as long as the quorum-shared disk and at least one node are available. This type of quorum also prevents more than one node from assuming the primary role.





**Note:** If the quorum-shared disk is not available, the cluster will stop functioning, even if all nodes are still available. In this mode, the quorum-shared disk is a single point of failure, so this mode is not recommended.

When you configure a failover cluster in Windows Server 2012, the Installation Wizard automatically selects one of two default configurations. By default, Failover Clustering selects:

- Node Majority if there is an odd number of nodes in the cluster.
- Node and Disk Majority if there is an even number of nodes in the cluster.

Modify this setting only if you determine that a change is appropriate for your cluster, and ensure that you understand the implications of making the change.

In addition to planning your quorum mode, you should also consider the capacity of the nodes in your cluster, and their ability to support the services and applications that may fail over to that node. For example, a cluster that has four nodes and a disk witness will still have quorum after two nodes fail. However, if you have several applications or services deployed on the cluster, each remaining cluster node may not have the capacity to provide services.

## What Are Cluster Shared Volumes?

In a classic failover cluster deployment, only a single node at a time controls an LUN on the shared storage. This means that the other nodes cannot “see” shared storage, until each node becomes an active node. CSV is a technology introduced in Windows Server 2008 R2 which enables multiple nodes to concurrently share a single LUN. Each node obtains exclusive access to individual files on the LUN instead of the whole LUN. In other words, CSVs provide a distributed file access solution so that multiple nodes in the cluster can simultaneously access the same NTFS file system.

The benefits of CSVs include:

- Fewer LUNs required
- Better use of disk space
- Resources in a single logical location
- No special hardware required
- Increased resiliency

To implement CSV:

1. Create and format volumes on shared storage
2. Add the disks to failover cluster storage
3. Add the storage to the CSV

In Windows Server 2008 R2, CSVs were designed only for hosting virtual machines running on a Hyper-V server in a failover cluster. This enabled administrators to have a single LUN that hosts multiple virtual machines in a failover cluster. Multiple cluster nodes have access to the LUN, but each virtual machine runs only on one node at a time. If the node on which the virtual machine was running fails, CSV lets the virtual machine to be restarted on a different node in the failover cluster. Additionally, this provides simplified disk management for hosting virtual machines compared to each virtual machine requiring a separate LUN.

In Windows Server 2012, CSVs have been additionally enhanced. It is now possible to use CSVs for other roles, and not just Hyper-V. For example, you can now configure file server role in a failover cluster in a Scale-Out File Server scenario. The Scale-Out File Server is designed to provide scale-out file shares that are continuously available for file-based server application storage. Scale-out file shares provides the ability to share the same folder from multiple nodes of the same cluster. In this context, CSVs in Windows Server 2012 introduces support for a read cache, which can significantly improve performance in certain scenarios. Also, a CSV File System (CSVFS) can perform CHKDSK without affecting applications with open handles on the file system.

Other important improvements in Cluster Shared Volumes in Windows Server 2012 are:

- *CSVFS benefits.* In Disk Management, CSV volumes now appear as CSVFS. However, this is not a new file system. The underlying technology is still the NTFS file system, and CSVFS volumes are still formatted with NTFS. However, because volumes appear as CSVFS, applications can discover that they are running on CSVs, which helps improve compatibility. And because of a single file namespace, all files have the same name and path on any node in a cluster.
- *Multisubnet support for CSVs.* CSVs have been enhanced to integrate with SMB Multichannel to help achieve faster throughput for CSV volumes.
- *Support for BitLocker drive encryption.* Windows Server 2012 support BitLocker volume encryption for both traditional clustered disks and CSVs. Each node performs decryption by using the computer account for the cluster itself.
- *Support for SMB 3.0 storage.* CSVs in Windows Server 2012 provide support for SMB 3.0 storage for Hyper-V and applications such as Microsoft SQL Server.
- *Integration with SMB Multichannel and SMB Direct.* This allows CSV traffic to stream across multiple networks in the cluster and to take advantage of network adapters that support Remote Direct Memory Access (RDMA).
- *Integration with the Storage Spaces feature in Windows Server 2012.* This can provide virtualized storage on clusters of inexpensive disks.
- *Ability to scan and repair volumes.* CSVs in Windows Server 2012 support the ability to scan and repair volumes with zero offline time.

## Implementing Cluster Shared Volumes

You can configure a CSV only when you create a failover cluster. After you create the failover cluster, you can enable the CSV for the cluster, and then add storage to the CSV.

Before you can add storage to the CSV, the LUN must be available as shared storage to the cluster. When you create a failover cluster, all the shared disks configured in Server Manager are added to the cluster, and you can add them to a CSV. If you add more LUNs to the shared storage, you must first create volumes on the LUN, add the storage to the cluster, and then add the storage to the CSV.

As a best practice, you should configure CSV before you make any virtual machines highly available. However, you can convert from regular disk access to CSV after deployment. The following considerations apply:

- When you convert from regular disk access to CSV, the LUN's drive letter or mount point is removed. This means that you must re-create all virtual machines that are stored on the shared storage. If you must retain the same virtual machine settings, consider exporting the virtual machines, switching to CSV, and then importing the virtual machines in Hyper-V.
- You cannot add shared storage to CSV if it is in use. If you have a running virtual machine that is using a cluster disk, you must shut down the virtual machine, and then add the disk to CSV.



### Additional Reading:

Server Message Block overview

<http://technet.microsoft.com/en-us/library/hh831795.aspx>

Storage Spaces Overview

<http://technet.microsoft.com/en-us/library/hh831739.aspx>



## Lesson 2

# Implementing a Failover Cluster

Failover clusters Windows Server 2012 have specific recommended hardware and software configurations that enable Microsoft to support the cluster. Failover clusters are intended to provide a higher level of service than stand-alone servers. Therefore, cluster hardware requirements are frequently stricter than requirements for stand-alone servers.

This lesson describes how to prepare for cluster implementation and also discusses the hardware, network, storage, infrastructure, and software requirements for Windows Server 2012 failover clusters. This lesson also outlines the steps for using the Validate a Configuration Wizard to ensure correct cluster configuration.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe how to prepare for implementing Failover Clustering.
- Describe hardware requirements for Failover Clustering.
- Describe network requirements for Failover Clustering.
- Describe infrastructure requirements for Failover Clustering.
- Describe software requirements for Failover Clustering.
- Validate and configure a cluster.

### Preparing for Implementing Failover Clustering

Before you implement Failover Clustering technology, you must identify services and applications that you want to make highly available. Failover clustering cannot be applied to all applications. Also, you should be aware that Failover Clustering does not provide improved scalability by adding nodes. You can only obtain scalability by scaling up and using more powerful hardware for the individual nodes. Therefore, you should only use Failover Clustering when your goal is high availability, instead of scalability.

Failover clustering is best suited for stateful applications that are restricted to a single set of data. One example of such an application is a database. Data is stored in a single location and can only be used by one database instance. You can also use Failover Clustering for Hyper-V virtual machines.

Failover clustering uses only IP-based protocols and is, therefore, suited only to IP-based applications. Both IP version 4 (IPv4) and IP version 6 (IPv6) are supported.

The best results for Failover Clustering occur when the client can do reconnecting to the application automatically after failover. If the client does not reconnect automatically, then the user must restart the client application.

Use failover clustering when:

- High availability is required
- Scalability is not required
- Application is stateful
- Client automatically reconnects to the application
- Application uses IP-based protocols



Consider the following guidelines when planning node capacity in a failover cluster:

- Spread out the highly-available applications from a failed node. When all nodes in a failover cluster are active, the highly-available services or applications from a failed node should be spread out among the remaining nodes to prevent a single node from being overloaded.
- Ensure that each node has sufficient idle capacity to service the highly-available services or applications that are allocated to it when another node fails. This idle capacity should be a sufficient buffer to avoid nodes running at near capacity after a failure event. Failure to adequately plan resource utilization can result in decrease in performance following node failure.
- Use hardware with similar capacity for all nodes in a cluster. This simplifies the planning process for failover because the failover load will be evenly distributed among the surviving nodes.
- Use standby servers to simplify capacity planning. When a passive node is included in the cluster, then all highly-available services or applications from a failed node can be failed over to the passive node. This avoids the need for complex capacity planning. If this configuration is selected, it is important that the standby server has sufficient capacity to run the load from more than one node failure.

You should also examine all cluster configuration components to identify single points of failure. You can remedy many single points of failure with simple solutions, such as adding storage controllers to separate and stripe disks, or teaming network adapters, and using multipathing software. These solutions reduce the probability that a failure of a single device causing a failure in the cluster. Typically, server class computer hardware has options for multiple power supplies for power redundancy, and for creating redundant array of independent disks (RAID) sets for disk data redundancy.

## Hardware Requirements for Failover Cluster Implementation

It is very important to make good decisions when you select hardware for cluster nodes. Failover clusters have to satisfy the following criteria to meet availability and support requirements:

- All hardware that you select for a failover cluster should meet the “Certified for Windows Server 2012” logo requirements. Hardware that has this logo was independently tested to meet the highest technical bar for reliability, availability, stability, security, and platform compatibility. Also, this means that official support options exist in case malfunctions arise.
- You should install the same or similar hardware on each failover cluster node. For example, if you choose a specific model of network adapter, you should install this adapter on each of the cluster nodes.
- If you are using Serial Attached SCSI or Fiber Channel storage connections, the mass-storage device controllers that are dedicated to the cluster storage should be identical in all clustered servers. They should also use the same firmware version.
- If you are using iSCSI storage connections, each clustered server must have one or more network adapters or host bus adapters dedicated to the cluster storage. The network that you use for iSCSI storage connections should not be used for network communication. In all clustered servers, the network adapters that you use to connect to the iSCSI storage target should be identical, and we recommend that you use Gigabit Ethernet or more.

The hardware requirements for a failover implementation include:

- Server hardware components must be marked with the Certified for Windows Server 2012 logo
- Server nodes should all have the same configuration and contain the same or similar components
- All tests in the Validate a Configuration Wizard must be passed

- After you configure the servers with the hardware, all tests provided in the Validate a Configuration Wizard must be passed before the cluster is considered a configuration that is supported by Microsoft.

## Network Requirements for Failover Cluster Implementation

Failover cluster network components must have the Certified for Windows Server 2012 logo and also pass the tests in the Validate a Configuration Wizard. Additionally:

- The network adapters in each node should be identical and have the same IP protocol version, speed, duplex, and flow control capabilities that are available.
- The networks and network equipment to which you connect the nodes should be redundant so that even a single failure allows for the nodes to continue communicating with one another. You can use network adapter teaming to provide single network redundancy. We recommend multiple networks to provide multiple paths between nodes for inter-node communication; otherwise, a warning will be generated during the validation process.
- The network adapters in a cluster network must have the same IP address assignment method, which means either that they all use static IP addresses or that they all use DHCP.

The network requirements for a failover implementation include:

- The network hardware components must be marked with the Certified for Windows Server 2012 logo
- The server should be connected to multiple networks for communication redundancy, or to a single network with redundant hardware, to remove single points of failure
- The network adapters should be identical and have the same IP protocol versions, speed, duplex, and flow control capabilities



**Note:** If you connect cluster nodes with a single network, the network passes the redundancy requirement in the Validate a Configuration Wizard. However, the report from the wizard will include a warning that the network should not have single points of failure.

## Infrastructure Requirements for Failover Cluster

Failover clusters depend on infrastructure services. Each server node must be in the same Active Directory domain, and if you use Domain Name System (DNS), the nodes should use the same DNS servers for name resolution.

We recommend that you install the same Windows Server 2012 features and roles on each node. Inconsistent configuration on cluster nodes can cause instability and performance issues. In addition, you should not install the AD DS role on any of the cluster nodes because AD DS has its own fault-tolerance mechanism. If you install the AD DS role on one of the nodes, you must install it on all nodes.

- The infrastructure requirements for a failover implementation include:
  - The nodes in the cluster must use DNS for name resolution
  - All servers in the cluster must be in the same Active Directory domain
  - The user account that creates the cluster must have administrator rights and permissions on all servers, and the Create Computer Objects permission in the domain
- Failover cluster infrastructure recommendations include:
  - The same roles should be installed on each cluster node
  - The AD DS role should not be installed on any of the cluster nodes

You must have the following network infrastructure for a failover cluster:

- *Network settings and IP addresses.* When you use identical network adapters for a network, also use identical communication settings on those adapters such as speed, duplex mode, flow control, and media type. Also, compare the settings between the network adapter and the switch it connects to, and ensure that no settings are in conflict. Otherwise, network congestion or frame loss might occur which could adversely affect how the cluster nodes communicate among themselves, with clients or with storage systems.
- *Unique subnets.* If you have private networks that are not routed to the rest of the network infrastructure, ensure that each of these private networks uses a unique subnet. This is necessary even if you give each network adapter a unique IP address. For example, if you have a cluster node in a central office that uses one physical network, and another node in a branch office that uses a separate physical network; do not specify 10.0.0.0/24 for both networks, even if you give each adapter a unique IP address. This avoids routing loops and other network communications problems if, for example, the segments are accidentally configured into the same collision domain because of incorrect vLAN assignments.
- *DNS.* The servers in the cluster typically use DNS for name resolution. DNS dynamic update protocol is a supported configuration.
- *Domain role.* All servers in the cluster must be in the same Active Directory domain. As a best practice, all clustered servers should have the same domain role (either member server or domain controller). The recommended role is member server because AD DS inherently includes its own failover protection mechanism.
- *Account for administering the cluster.* When you first create a cluster or add servers to it, you must be logged on to the domain with an account that has administrator rights and permissions on all servers in that cluster. The account does not have to be a Domain Admins account, but can be a Domain Users account that is in the Administrators group on each clustered server. In addition, if the account is not a Domain Admins account, the account (or the group that the account is a member of) must be given the Create Computer Objects permission in the domain.

In Windows Server 2012, there is no cluster service account. Instead, the Cluster service the Cluster service automatically runs in a special context that provides the specific permissions and credentials that are necessary for the service (similar to the local system context, but with reduced credentials). When a failover cluster is created and a corresponding computer object is created in AD DS, that object is configured to prevent accidental deletion. Also, the cluster Network Name resource has additional health check logic, which periodically checks the health and properties of the computer object that represents the Network Name resource.

## Software Requirements for Failover Cluster Implementation

Failover clusters require that each cluster node must run the same edition of Windows Server 2012. The edition can be either Windows Server 2012 Enterprise or Windows Server 2012 Datacenter. The nodes should also have the same software updates and service packs. Depending on the role that will be clustered, a Server Core installation may also meet the software requirements. However, you cannot install Server Core and full editions in the same cluster.

The software requirements for a failover implementation include:

- All nodes must run the same edition of Windows Server 2012, which can be any of the following:
  - Windows Server 2012 Enterprise, Full or Server Core installation
  - Windows Server 2012 Datacenter, Full or Server Core installation
- All nodes must run the same processor architecture (32-bit, x64-based, or Itanium architecture-based)
- All nodes should have the same service pack and updates

It is also very important that the same version of service packs or any operating system updates, exist on all nodes that are parts of a cluster.



**Note:** Windows Server 2012 provides Cluster-Aware Updating technology that can help you maintain updates on cluster nodes. This feature will be discussed in more detail in *Lesson 4: Maintaining a Failover Cluster*.

Each node must run the same processor architecture. This means that each node must have the same processor family, which might be the Intel Xeon processor family with Extended Memory 64Technology, the AMD Opteron AMD64 family, or the Intel Itanium-based processor family.

## Demonstration: Validating and Configuring a Failover Cluster

The Validate a Configuration Wizard runs tests that confirm if the hardware and hardware settings are compatible with Failover Clustering. Using the wizard, you can run the complete set of configuration tests or a subset of the tests. We recommend that you run the tests on servers and storage devices before you configure the failover cluster, and again after any major changes are made to the cluster. You can access the test results in the %windir%\cluster\Reports directory.

### Demonstration Steps

1. Start Failover Cluster Manager on the **LON-SVR3** machine.
2. Start the Validate Configuration Wizard. Add LON-SVR3 and LON-SVR4 as cluster nodes.
3. Review the report.
4. Create a new cluster. Add LON-SVR3 and LON-SVR4 as cluster nodes.
5. Name the cluster as **Cluster1**.
6. Use **172.16.0.125** as **IP address**.

## Lesson 3

# Configuring Highly-Available Applications and Services on a Failover Cluster

After you have configured clustering infrastructure, you should configure specific role or service to be highly available. Not all roles can be clustered. Therefore, you should first identify the resource that you want to put in a cluster and check whether it is supported. In this lesson, you will learn about configuring roles and applications in clusters as well as about configuring cluster settings.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe and identify cluster resources and services.
- Describe the process for clustering server roles.
- Configure a cluster role.
- Describe how to configure cluster properties.
- Describe how to manage cluster nodes.
- Describe how to configure application failover settings.

### Identifying Cluster Resources and Services

A clustered service that contains an IP address resource and a network name resource (and other resources) is published to a client on the network under a unique server name. Because this group of resources is displayed as a single logical server to clients, it is called a cluster instance.

Users access applications or services on an instance in the same manner they would if the applications or services were on a nonclustered server. Usually, applications or users do not know that they are connecting to a cluster and the node they are connected to.

- Clustered services:
  - Are services or applications that are made highly available by installing them on a failover cluster
  - Are active on one node, but can be moved to another node
- Resources:
  - Are the components that make up a clustered service
  - Are moved to another node when one node fails
  - Can only run on one node at a time
  - Include components such as shared disks, names, and IP addresses

Resources are physical or logical entities, such as a file share, disk, or IP address that the failover cluster manages. Resources may provide a service to clients or may be an important part of the cluster. Resources are the most basic and smallest configurable unit. At any time, a resource can run only on a single node in a cluster, and it is online on a node when it provides its service to that specific node.

### Server Cluster Resources

A cluster resource is any physical or logical component that has the following characteristics:

- It can be brought online and taken offline.
- It can be managed in a server cluster.
- It can be hosted (owned) by only one node at a time.



To manage resources, the Cluster service communicates to a resource DLL through a resource monitor. When the Cluster service makes a request of a resource, the resource monitor calls the appropriate entry-point function in the resource DLL to check and control the resource state.

### Dependent Resources

A dependent resource is one that requires another resource to operate. For example, a network name must be associated with an IP address. Because of this requirement, a network name resource depends on an IP address resource. Dependent resources are taken offline before the resources upon which they depend are taken offline; similarly, they are brought online after the resources on which they depend are brought online. A resource can specify one or more resources on which it is dependent. Resource dependencies also determine bindings. For example, clients will be bound to the particular IP address that a network name resource depends on.

When you create resource dependencies, consider the fact that, although some dependencies are strictly required, others are not required but are recommended. For example, a file share that is not a Distributed File System (DFS) root has no required dependencies. However, if the disk resource that holds the file share fails, the file share will be inaccessible to users. Therefore, it is logical to make the file share dependent on the disk resource.

A resource can also specify a list of nodes on which it can run. Possible nodes and dependencies are important considerations when administrators organize resources into groups.

### The Process for Clustering Server Roles

Failover clustering supports the clustering of several Windows Server roles, such as File Services, DHCP, and Hyper-V. To implement clustering for a server role, or for external applications such as SQL Server or Exchange Server, perform the following procedure:

1. Install the Failover Clustering feature. Use Server Manager or Ocsetup to install the Failover Clustering feature on all computers that will be cluster members.
2. Verify configuration and create a cluster with the appropriate nodes. Use the Failover Cluster Management snap-in to first validate a configuration, and then create a cluster with selected nodes.
3. Install the role on all cluster nodes. Use Server Manager to install the server role that you want to use in the cluster.
4. Create a clustered application by using the Failover Clustering Management snap-in.
5. Configure the application. Configure options on the application that is being used in the cluster.
6. Test failover. Use the Failover Cluster Management snap-in to test failover by intentionally moving the service from one node to another.

After the cluster is created, you can monitor its status by using the Failover Cluster Management console, and manage available options.

1. Install the failover clustering feature
2. Verify the configuration and create a cluster
3. Install the role on all cluster nodes, using Server Manager
4. Create a clustered application by using the Failover Clustering Management snap-in
5. Configure the application
6. Test the failover

## Demonstration: Clustering a File Server Role

### Demonstration Steps

1. Open Failover Cluster Manager and verify that three Cluster Disks are available.
2. Start the Configure Role Wizard and Configure the **File Server** as clustered role.
3. For the Client Access Point, use the name **AdatumFS** and the IP address of **172.16.0.130**.
4. Select **Cluster Disk 2** as the storage for the File Server role.

## Failover Cluster Management Tasks

You can perform several failover cluster management tasks. These tasks range from adding and removing cluster nodes to modifying the quorum settings. Some of the most frequently used configuration tasks include:

- Managing cluster nodes – for each node in a cluster, you can stop cluster service temporary, pause it, initiate remote desktop to the node or evict node from the cluster
- Managing cluster networks – You can add or remove cluster networks and you can also configure networks that will be dedicated just for inter-cluster communication
- Managing permissions – By managing permission you delegate rights to administer cluster
- Configuring cluster quorum settings – By configuring quorum settings you determine the way how quorum is achieved as well as who can have vote in a cluster
- Migrating services and applications to a cluster – You can implement existing services to the cluster and make them highly available
- Configuring new services and applications to work in a cluster – You can implement new services to the cluster
- Removing a cluster

The common management tasks include:

- Managing nodes
- Managing networks
- Managing permissions
- Configuring cluster quorum settings
- Migrating services and applications to a cluster
- Configuring new services and applications
- Removing the cluster

You can perform most of these administrative tasks by using the Failover Cluster Management console.



## Managing Cluster Nodes

Cluster nodes are mandatory for each cluster. After you create a cluster and put it into production, you might have to manage cluster nodes occasionally.

There are three aspects to managing cluster nodes:

- You can add a node to an established failover cluster by selecting Add Node in the Failover Cluster Management Actions pane. The Add Node Wizard prompts you for information about the additional node.
- You can pause a node to prevent resources from being failed over or moved to the node. You typically pause a node when a node is undergoing maintenance or troubleshooting.
- You can evict a node, which is an irreversible process for a cluster node. After you evict the node, it must be re-added to the cluster. You evict nodes when a node is damaged beyond repair or is no longer needed in the cluster. If you evict a damaged node, you can repair or rebuild it, and then add it back to the cluster by using the Add Node Wizard.

You can manage cluster nodes by using the Failover Cluster Management console.

To manage cluster nodes, you can:



Add nodes after you create a cluster



Pause nodes, which prevents resources from running on that node



Evict nodes from a cluster, which removes the node from the cluster configuration

All of these actions are available in the Failover Cluster Management Actions pane

## Configuring Application Failover Settings

You can adjust the failover settings, including preferred owners and failback settings, to control how the cluster responds when the application or service fails. You can configure these settings on the property sheet for the clustered service or application (on the General tab or on the Failover tab). The following table provides examples that show how these settings work.

The considerations for using preferred owners include:

- Preferred owners are set on the clustered application
- Multiple preferred owners can be set in an ordered list
- Setting preferred owners gives control over:
  - The order in which an application will select a node to run on
  - The applications that can be run on the same nodes in an Active/Active configuration

The options to modify failover and failback settings include:

- Setting the number of times the cluster service will restart a clustered application in a set period of time
- Setting or preventing failback of the clustered application to the preferred node when it becomes available

Setting	Result
Example 1: General tab, Preferred owner: Node1 Failover tab, Failback setting: Allow failback (Immediately)	If the service or application fails over from Node1 to Node2, when Node1 is again available, the service or application will fail back to Node1.
Example 2: Failover tab, Maximum failures in the specified period: 2 Failover tab, Period (hours): 6	In a six-hour period, if the application or service fails no more than two times, it will be restarted or failed over every time. If the application or service fails a third time in the six-hour period, it will be left in the failed state.  The default value for the maximum number of failures is n-1, where n is the number of nodes. You can change the value, but we recommend a fairly low value so that if multiple node failures occur, the application or service will not be moved between nodes indefinitely.

## Lesson 4

# Maintaining a Failover Cluster

When cluster infrastructure is up and running, it is very important to establish monitoring to prevent possible failures. Also, it is important to have backup and restore procedures for cluster configuration. In Windows Server 2012, there is a new technology that lets you update cluster nodes without downtime. In this lesson, you will learn about monitoring, backup, and restore and about updating cluster nodes.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe how to monitor failover clusters.
- Describe how to back up and restore cluster configuration.
- Describe how to troubleshoot failover clusters.
- Describe Cluster-Aware Updating.
- Configure Cluster-Aware Updating.

### Monitoring Failover Clusters

Many tools are available to help you monitor failover clusters. You can use standard Windows Server tools, such as the Event Viewer and the Performance and Reliability Monitor snap-in, to review cluster event logs, and performance metrics. You can also use Cluster.exe and Tracerpt.exe to export data for analysis. Additionally, you can use the MHTML-formatted cluster configuration reports and the Validate a Configuration Wizard to troubleshoot problems with the cluster configuration and hardware changes.

Some of the tools you can use to monitor clusters include:

- Event Viewer
- Tracerpt.exe
- Performance and Reliability Monitor snap-in
- MHTML-formatted cluster configuration reports
- Validate a Configuration Wizard

### Event Viewer

When problems arise in the cluster, use the Event Viewer to view events with a Critical, Error, or Warning severity level. Additionally, informational level events are logged to the Failover Clustering Operations log, which can be found in the Event Viewer in the Applications and Services Logs\Microsoft\Windows folder. Informational-level events are usually common cluster operations, such as cluster nodes leaving and joining the cluster, or resources going offline or coming online.

In previous Windows Server versions, event logs were replicated to each node in the cluster. This simplified cluster troubleshooting, because you could review all event logs on a single cluster node. Windows Server 2012 does not replicate the event logs between nodes. However, the Failover Cluster Management snap-in has a Cluster Events option that enables you to view and filter events across all cluster nodes. This feature is helpful in correlating events across cluster nodes.

The Failover Cluster Management snap-in also provides a Recent Cluster Events option that will query all the Error and Warning events from all the cluster nodes in the last 24 hours.

You can access additional logs, such as the Debug and Analytic logs, in the Event Viewer. To display these logs, modify the view on the top menu by selecting the Show Analytic and Debug Logs options.

## Windows Event Tracing

Windows event tracing is a kernel component that is available early after startup, and late into shutdown. It is designed to allow for fast tracing and delivery of events to trace files and to consumers. Because it is designed to be fast, it enables only basic in-process filtering of events based on event attributes.

The event trace log contains a comprehensive accounting of the failover cluster actions. Depending on how you want to view the data, use either Cluster.exe or Tracerpt.exe to access the information in the event trace log.

Tracerpt.exe will parse the event trace logs only on the node on which it is run. All the individual logs are collected in a central location. To transform the XML file into a text file or an HTML file that can be opened in Internet Explorer®, you can parse the XML-based file by using the Microsoft XSL parsing command prompt utility msxsl.exe, and an XSL style sheet.

## Performance and Reliability Monitor Snap-In

The Performance and Reliability Monitor snap-in lets you:

- Trend application performance on each node. To determine how an application is performing, you can view and trend specific information on system resources that are being used on each node.
- Trend application failures and stability on each node. You can pinpoint when application failures occur and match the application failures with other events on the node.
- Modify trace log settings. You can start, stop, and adjust trace logs, including their size and location.

## Backing Up and Restoring Failover Cluster Configuration

Cluster configuration can be a time-consuming process with many details, and so backup of cluster configuration is very important. You can perform backup and restore of cluster configuration with Windows Server Backup or a third-party backup tool.

When you back up the cluster configuration, be aware of the following:

- You must test your backup and recovery process, before putting a cluster into production.
- You must first add the Windows Server Backup feature, if you decide to use it. You can do this by using Server Manager.

Windows Server Backup is the built-in backup and recovery software for Windows Server 2012. To complete a successful backup, consider the following:

- For a backup to succeed in a failover cluster, the cluster must be running and must have quorum. In other words, enough nodes must be running and communicating (perhaps with a witness disk or witness file share, depending on the quorum configuration,) that the cluster has achieved quorum.
- You must back up all clustered applications. If you cluster a Microsoft SQL Server® database, you must have a backup plan for the databases and configuration outside the cluster configuration.
- If application data must be backed up, the disks that you store the data on must be made available to the backup software. You can achieve this by running the backup software from the cluster node that owns the disk resource, or by running a backup against the clustered resource over the network.

When backing up failover clusters, keep in mind that:

- Windows Server Backup is a Windows Server 2008 feature
- You install Windows Server Backup as a feature
- Backup and restore operations involve the Volume Shadow Copy Service (VSS)
- Third-party tools are available to perform backups and restores
- You must perform system-state backups

A non-authoritative restore completely restores a single node in the cluster

An authoritative restore restores the entire cluster configuration to a point in time

- The cluster service keeps track of which cluster configuration is the most recent, and it replicates that configuration to all cluster nodes. If the cluster has a witness disk, the Cluster service the Cluster service also replicates the configuration to the witness disk.

## Restoring a Cluster

There are two types of restore:

- *Non-authoritative restore.* Use a non-authoritative restore when a single node in the cluster is damaged or rebuilt, and the rest of the cluster is operating correctly. Perform a non-authoritative restore by restoring the system recovery (system state) information to the damaged node. When you restart that node, it will join the cluster and receive the latest cluster configuration automatically.
- *Authoritative restore.* Use an authoritative restore when the cluster configuration must be rolled back to a previous point in time. For example, you would use an authoritative restore if an administrator accidentally removed clustered resources or modified other cluster settings. Perform the authoritative restore by stopping the cluster resource on each node, and then performing a system recovery (system state) on a single node by using the command-line Windows Server Backup interface. After the restored node restarts the cluster service, the remaining cluster nodes can also start the cluster service.

## Troubleshooting Failover Clusters

Although cluster validation implemented in Windows Server 2012 Failover Clustering prevents misconfigurations and non-working clusters, in some cases, you have to perform cluster troubleshooting.

To troubleshoot a failover cluster, follow these guidelines:

- Use the Validate a Configuration Wizard to highlight configuration issues that might cause cluster problems.
- Review cluster events and trace logs to identify application or hardware issues that might cause an unstable cluster.
- Review hardware events and logs to help pinpoint specific hardware components that might cause an unstable cluster.
- Review SAN components, switches, adapters, and storage controllers to help identify any potential problems.

When troubleshooting failover clusters, you must:

- Identify the perceived problem by collecting and documenting the symptoms of the problem.
- Identify the scope of the problem so that you can understand what is being affected by the problem, and what impact that effect has on the application and the clients.
- Collect information so that you can accurately understand and pinpoint the possible problem. After you identify a list of possible problems, you can prioritize them by probability, or the impact of a repair. If the problem cannot be pinpointed, you should attempt to re-create the problem.

The failover cluster troubleshooting techniques include:

- Reviewing events in logs (cluster, hardware, storage)
- Using the Validate a Configuration Wizard
- Defining a process for troubleshooting failover clusters
- Reviewing storage configuration
- Checking for group and resource failures

- Create a schedule for repairing the problem. For example, if the problem only affects a small subset of users, you can delay the repair to an off-peak time so that you can schedule downtime.
- Complete and test each repair one at a time so that you can identify the fix.

To troubleshoot SAN issues, start by checking physical connections and each of the hardware component logs. Additionally, run the Validate a Configuration Wizard to verify that the current cluster configuration is still supportable. When you run the Validate a Configuration Wizard, ensure that the storage tests that you select can be run on an online failover cluster. Several of the storage tests cause loss of service on the clustered disk when the tests are run.

### Troubleshooting Group and Resource Failures

To troubleshoot group and resource failures:

- Use the Dependency Viewer in the Failover Cluster Management snap-in to identify dependent resources.
- Check the Event Viewer and trace logs for errors from the dependent resources.
- Determine whether the problem only happens on a specific node, or nodes, by trying to re-create the problem on different nodes.

### What Is Cluster-Aware Updating?

Applying operating system updates to nodes in a cluster requires special attention. If you want to provide zero downtime for a clustered role, you must manually update cluster nodes one after another, and you must manually move resources from the node being updated to another node. This procedure can be very time-consuming. In Windows Server 2012, Microsoft has implemented a new feature for automatic update of cluster nodes.

Cluster-Aware Updating (CAU) is a feature that lets administrators automatically update cluster nodes with little or no loss in availability during the update process. During an update procedure, CAU transparently takes each cluster node offline, installs the updates and any dependent updates, performs a restart if necessary, brings the node back online, and then moves to update the next node in a cluster.

For many clustered roles, this automatic update process triggers a planned failover, and it can cause a transient service interruption for connected clients. However, for continuously available workloads in Windows Server 2012, such as Hyper-V with live migration or file server with SMB Transparent Failover, CAU can orchestrate cluster updates with no effect on the service availability.

Cluster-Aware Updating is an automated feature, specific to Windows Server 2012, that updates nodes in a cluster with minimal or zero downtime

Cluster-Aware Updating can work in two modes:

- Remote-Updating Mode
- Self-Updating Mode

## Cluster Updating Modes

CAU can orchestrate the complete cluster updating operation in two modes:

- *Remote-updating mode.* In this mode, a computer that is running Windows Server 2012 or Windows 8, is called and configured as an orchestrator. To configure a computer as a CAU orchestrator, you must install Failover Clustering administrative tools on it. The orchestrator computer is not a member of the cluster that is updated during the procedure. From the orchestrator computer, the administrator triggers on-demand updating by using a default or custom Updating Run profile. Remote-updating mode is useful for monitoring real-time progress during the Updating Run, and for clusters that are running on Server Core installations of Windows Server 2012.
- *Self-updating mode.* In this mode, the CAU clustered role is configured as a workload on the failover cluster that is to be updated, and an associated update schedule is defined. In this scenario, CAU does not have a dedicated orchestrator computer. The cluster updates itself at scheduled times by using a default or custom Updating Run profile. During the Updating Run, the CAU orchestrator process starts on the node that currently owns the CAU clustered role, and the process sequentially performs updates on each cluster node. In the self-updating mode, CAU can update the failover cluster by using a fully automated, end-to-end updating process. An administrator can also trigger updates on-demand in this mode, or use the remote-updating approach if desired. In the self-updating mode, an administrator can access summary information about an Updating Run in progress by connecting to the cluster and running the Get-CauRun Windows PowerShell cmdlet.

To use CAU, you must install the Failover Clustering feature in Windows Server 2012 and create a failover cluster. The components that support CAU functionality are automatically installed on each cluster node.

You must also install the CAU tools, which are included in the Failover Clustering Tools (which are also part of the Remote Server Administration Tools, or RSAT). The CAU tools consist of the CAU UI and the CAU Windows PowerShell cmdlets. The Failover Clustering Tools are installed by default on each cluster node when you install the Failover Clustering feature. You can also install these tools on a local or a remote computer that is running Windows Server 2012 or Windows 8 and that has network connectivity to the failover cluster.

## Demonstration: Configuring Cluster-Aware Updating

### Demonstration Steps

1. Make sure that the cluster is configured and running on LON-SVR3 and LON-SVR4.
2. Add the **Failover Clustering Feature** to LON-DC1.
3. Run **Cluster-Aware Updating** on LON-DC1 and configure it to connect to **Cluster1**.
4. Preview updates that are available for nodes LON-SVR3 and LON-SVR4.
5. Review available options for the Updating Run Profile.
6. Apply available updates to **Cluster1** from LON-DC1.
7. After updates are applied, configure **Cluster self-updating options** on LON-SVR3.



## Lesson 5

# Implementing a Multi-Site Failover Cluster

In some scenarios, you have to deploy cluster nodes on different sites. Usually, you do this when you build disaster-recovery solutions. In this lesson, you will learn about deploying multi-site clusters.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe a multi-site cluster.
- Describe synchronous and asynchronous replication.
- Describe how to choose a quorum mode for multi-site clusters.
- Describe the challenges for implementing multi-site clusters.
- Describe the considerations for deploying multi-site clusters.

### What Is a Multi-Site Cluster?

A multi-site cluster provides highly-available services in more than one location. Multi-site clusters can solve several specific problems. However, they also present specific challenges.

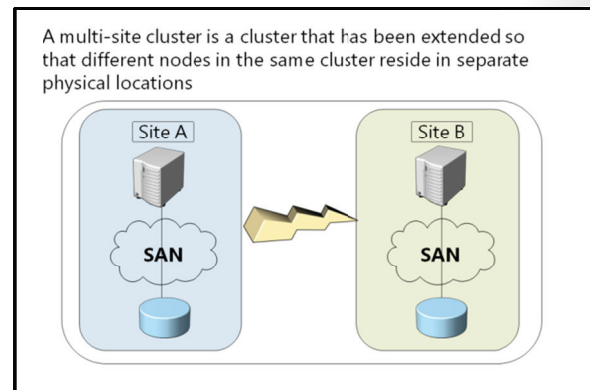
In a multi-site cluster, each site usually has a separate storage system with replication between the sites. Multi-site cluster storage replication enables each site to be independent, and provides fast access to the local disk. With separate storage systems, you cannot share a single disk between sites.

A multi-site cluster has three main advantages in a failover site compared to a remote server:

- When a site fails, a multi-site cluster automatically fails over the clustered service or application to another site.
- Because the cluster configuration is automatically replicated to each cluster node in a multi-site cluster, there is less administrative overhead than a cold standby server, which requires you to manually replicate changes.
- The automated processes in a multi-site cluster reduce the possibility of human error, which is present in manual processes.

Because of increased cost and complexity of a multi-site failover cluster, it might not be an ideal solution for every application or business. When you are considering whether to deploy a multi-site cluster, you should evaluate the importance of the applications to the business, the type of applications, and any alternative solutions. Some applications can provide multi-site redundancy easily with log shipping or other processes, and can still achieve sufficient availability with only a modest increase in cost and complexity.

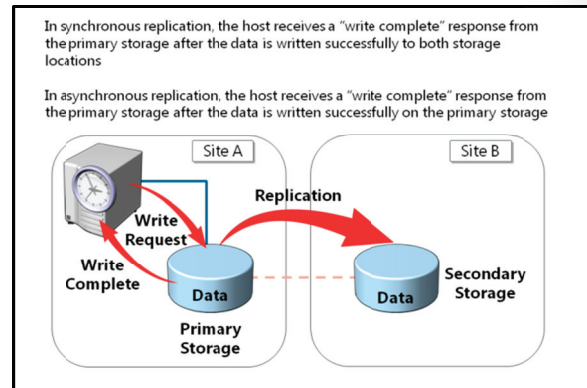
The complexity of a multi-site cluster requires better architectural and hardware planning. It also requires you to develop business processes to routinely test the cluster functionality.



## Synchronous and Asynchronous Replication

It is not possible for a geographically-dispersed failover cluster to use shared storage between physical locations. Wide area network (WAN) links are too slow and have too much latency to support shared storage. Geographically-dispersed failover clusters must synchronize data between locations by using specialized hardware. Multi-site data replication can be either synchronous or asynchronous:

- When you use synchronous replication, the host receives a "write complete" response from the primary storage after the data is written successfully on both storage systems. If the data is not written successfully to both storage systems, the application must attempt to write to the disk again. With synchronous replication, both storage systems are identical.
- When you use asynchronous replication, the node receives a write complete response from the storage after the data is written successfully on the primary storage. The data is written to the secondary storage on a different schedule, depending on the hardware or software vendor's implementation. Asynchronous replication can be storage-based, host-based, or even application-based. However, not all forms of asynchronous replication are sufficient for a multi-site cluster. For example, Distributed File System Replications (DFS-R) provides file-level asynchronous replication. However, it does not support multi-site Failover Clustering replication. This is because DFS-R replicates smaller documents that are not held open continuously. Therefore, it was not designed for high-speed, open-file replication.



### When to Use Synchronous or Asynchronous Replication

Use synchronous replication when data loss cannot be tolerated. Synchronous replication solutions require low-disk write latency, because the application waits for both storage solutions to acknowledge the data writes. The requirement for low latency disk writes also limits the distance between the storage systems because increased distance can cause higher latency. If the disk latency is high, the performance and even the stability of the application can be affected.

Asynchronous replication overcomes latency and distance limitations by acknowledging local disk writes only, and by reproducing the disk write on the remote storage system in a separate transaction. Because asynchronous replication writes to the remote storage system after it writes to the local storage system, the possibility of data loss during a failure is increased.



## Choosing a Quorum Mode for Multi-Site Clusters

For a geographically-dispersed cluster, you cannot use quorum configurations that require a shared disk, because geographically-dispersed clusters do not use shared disks. Both the Node and Disk Majority, and No Majority: Disk Only quorum modes require a shared witness disk to provide a vote for determining quorum. You should only use these two quorum modes if the hardware vendor specifically recommends and supports them.

To use the Node and Disk Majority and No Majority: Disk Only modes in a multi-site cluster, the shared disk requires that:

- You preserve the semantics of the SCSI commands across the sites, even if a complete communication failure occurs between sites.
- You replicate the witness disk in real-time synchronous mode across all sites.

Because multi-site clusters can have WAN failures in addition to node and local network failures, Node Majority and Node and File Share Majority are better solutions for multi-site clusters. If there is a WAN failure that causes the primary and secondary sites to lose communication, a majority must still be available to continue operations.

If there are an odd number of nodes, then use the Node Majority quorum. If there is an even number of nodes, which is typical in a geographically-dispersed cluster, you can use the Node Majority with File Share quorum.

If you are using Node Majority and the sites lose communication, you need a mechanism to determine which nodes stay up, and which nodes drop out of cluster membership. The second site requires another vote to obtain quorum after a failure. To obtain another vote for quorum, you must join another node to the cluster, or create a file share witness.

The Node and File Share Majority mode can help maintain quorum without adding another node to the cluster. To provide for a single-site failure and enable automatic failover, the file share witness might have to exist at a third site. In a multi-site cluster, a single server can host the file share witness. However, you must create a separate file share for each cluster.

You must use three locations to enable automatic failover of a highly-available service or application. Locate one node in the primary location that runs the highly-available service or application. Locate a second node in a disaster-recovery site, and locate the third node for the file share witness in another location.

There must be direct network connectivity between all three locations. In this manner, if one site becomes unavailable, the two remaining sites can still communicate and have enough nodes for a quorum.

When designing automatic failover for geographically-dispersed clusters:

- Use Node Majority or Node Majority with File Share quorum
- Use three locations to allow automatic failover of a single virtual server:
  - All three locations must be linked directly to each other
  - One location is only a file-share witness



**Note:** In Windows Server 2008 R2, administrators could configure the quorum to include nodes. However, if the quorum configuration included nodes, all nodes were treated equally according to their votes. In Windows Server 2012, cluster quorum settings can be adjusted so that when the cluster determines whether it has quorum, some nodes have a vote and some do not. This adjustment, can be useful, when solutions are implemented across multiple sites.

## Challenges for Implementing a Multi-Site Cluster

Implementation of multi-site clusters is more complex than implementation of single-site clusters, and can also present several challenges to the administrator. Most important challenges when you implement multi-site clusters are related to storage and network.

In a multi-site cluster, there is no shared storage that the cluster node uses. This means that nodes on each site must have its own storage instance. On the other hand, Failover Clustering does not include any built-in functionality to replicate data between sites. There are three options for replicating data: block level hardware-based replication, software-based file replication installed on the host, or application-based replication.

Storage Challenge	Description
Requires a separate or third-party data replication solution	<ul style="list-style-type: none"> <li>• Hardware (block level) storage-based replication</li> <li>• Software (file system level) host-based replication</li> <li>• Application-based replication (such as Exchange 2007 Cluster Continuous Replication)</li> </ul>
Can be either synchronous or asynchronous replication	<ul style="list-style-type: none"> <li>• Synchronous. No acknowledgement of data changes made in Site A until the data is successfully written to Site B</li> <li>• Asynchronous. Data changes made in Site A will eventually be written to the storage in Site B</li> </ul>

- Inter-node communications are time sensitive; you might need to configure these thresholds to meet the higher WAN latency
- DNS replication might impact client reconnect times when failover is based on hostname
- Active Directory replication latency might effect application data availability
- Some applications might require all of the nodes to be in the same Active Directory site

Multi-site data replication can be either synchronous or asynchronous. Synchronous replication does not acknowledge data changes that are made in, for example, Site A until the data is successfully written to Site B. With asynchronous replication, data changes that are made in Site A are eventually written to Site B.

When you deploy a multi-site cluster and run the Validate a Configuration Wizard, the disk tests will not find any shared storage, and will therefore not run. However, you can still create a cluster. If you follow the hardware manufacturer's recommendations for Windows Server Failover Clustering hardware, Microsoft will support the solution.

Windows Server 2012 enables cluster nodes to exist on different IP subnets, which enables a clustered application or service to change its IP address based on the IP subnet. DNS updates the clustered application's DNS record so that clients can locate the IP address change. Because clients rely on DNS to find a service or application after a failover, you might have to adjust the DNS records' Time to Live, and the speed at which DNS data is replicated. Additionally, when cluster nodes are in multiple sites, network latency might require you to modify the inter-node communication (heartbeat) delay and time-out thresholds.

## Deploying Considerations for a Multi-Site Cluster

Multi-site clusters are not appropriate for every application or every business. When you design a multi-site solution with a hardware vendor, clearly identify the business requirements and expectations. Not every scenario that involves more than one location is appropriate for multi-site cluster.

Multi-site clustering is a high-availability strategy that primarily focuses on hardware platform availability. However, specific multi-site cluster configuration and deployment have availability ramifications, ranging from the ability of users to connect to the application to the quality of performance of the application. Multi-site clustering can be a powerful solution in dealing with planned and unplanned downtime, but its benefits must be examined against all the dimensions of application availability.

When deploying multi-site clusters:

- Ensure that the business requirements are met by the solution
- Use a hardware vendor to create a full solution for multi-site clusters
- Choose the correct quorum mode to properly maintain functionality in the event of failures
- Choose the correct storage replication solution to meet these needs

Multi-site clusters do require some more overhead than local clusters. Instead of a local cluster, in which each node of the cluster is attached to the mass storage device, each site of a multi-site cluster must have comparable storage. In addition, you will also have to consider vendors to set up your data replication schemes between cluster sites, possibly pay for additional network bandwidth between sites, and develop the management resources within your organization to efficiently administer your multi-site cluster.

Additionally, carefully consider the quorum mode that you will use, and the location of the available cluster votes.

## Lab: Implementing Failover Clustering

### Scenario

As A. Datum's business grows, it is becoming increasingly important that many of the applications and services on the network are available at all times. A. Datum has many services and applications that have to be available to internal and external users who work in different time zones around the world. Many of these applications cannot be made highly available by using Network Load Balancing. Therefore, you have to use a different technology to make these applications highly available.

As one of the senior network administrators at A. Datum, you will be responsible for implementing Failover Clustering on the Windows Server 2012 servers in order to provide high availability for network services and applications. You will also be responsible for planning the Failover Cluster configuration, and deploying applications and services on the Failover Cluster.

### Objectives

After completing this lab, you will be able to:

- Configure a failover cluster.
- Deploy and configure a highly-available file server.
- Validate the deployment of the highly-available file server.
- Configure Cluster-Aware Updating on the failover cluster.

### Lab Setup

Estimated time: **90 minutes**

Virtual Machine(s)	20417A-LON-DC1 20417A-LON-SVR1 20417A-LON-SVR3 20417A-LON-SVR4
User Name	Adatum\Administrator
Password	Pa\$\$w0rd

Virtual Machine(s)	MSL-TMG1
User Name	Administrator
Password	Pa\$\$w0rd

For this lab, you will use the available virtual machine environment. Before you begin the lab, you must complete the following steps:

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, click **20417A-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.

4. Log on using the following credentials:
  - a. User name: **Adatum\Administrator**
  - b. Password: **Pa\$\$w0rd**
5. Repeat steps 2-4 for **20417A-LON-SVR1**, **20417A-LON-SVR3**, and **20417A-LON-SVR4**.
6. Repeat steps 2-3 for **MSL-TMG1**. Log on as **Administrator** with the password of **Pa\$\$w0rd**.

## Exercise 1: Configuring a Failover Cluster

### Scenario

A. Datum has important applications and services that they want to make highly available. Some of these services cannot use Network Load Balancing. Therefore, you decided to implement Failover clustering. Because iSCSI storage is set up, you decided to use the iSCSI storage for Failover Clustering. First, you will implement the core components for Failover Clustering, validate the cluster, and then create the failover cluster.

The main tasks for this exercise are as follows:

1. Connect clients to the iSCSI targets.
2. Install the Failover Clustering feature.
3. Validate the servers for Failover Clustering.
4. Create the Failover Cluster.

#### ► Task 1: Connect clients to the iSCSI targets

1. On LON-SVR3, start **iSCSI Initiator**, and configure **Discover Portal** with IP address **172.16.0.21**.
2. Connect to the discovered target in the **Targets** list.
3. Repeat steps 1 and 2 on LON-SVR4.
4. Open **Disk Management** on LON-SVR3.
5. Bring online and initialize the three new disks.
6. Make a simple volume on each disk and format it with NTFS.
7. On LON-SVR4, open **Disk Management**, and bring online and initialize the three new disks.

#### ► Task 2: Install the Failover Clustering feature

1. On LON-SVR3, install the Failover Clustering feature by using Server Manager.
2. On LON-SVR4, install the Failover Clustering feature by using Server Manager.

#### ► Task 3: Validate the servers for Failover Clustering

1. On LON-SVR3, open the Failover Cluster Manager console.
2. Start the Validate a Configuration Wizard.
3. Use LON-SVR3 and LON-SVR4 as nodes for test.
4. Review report.

#### ► Task 4: Create the Failover Cluster

1. On LON-SVR3, in the Failover Cluster Manager, start the Create Cluster Wizard.
2. Use LON-SVR3 and LON-SVR4 as cluster nodes.

3. Specify **Cluster1** as the **Access Point name**.
4. Specify the **IP address** as **172.16.0.125**.

**Results:** After this exercise, you will have installed and configured the Failover Clustering feature.

## Exercise 2: Deploying and Configuring a Highly-Available File Server

### Scenario

In A. Datum, File Services is one of the important services that must be highly available. After you have created a cluster infrastructure, you decided to configure a highly-available file server and implement settings for failover and failback.

The main tasks for this exercise are as follows:

1. Add the File Server application to the failover cluster.
2. Add a shared folder to a highly-available file server.
3. Configure failover and failback settings.

#### ► Task 1: Add the File Server application to the failover cluster

1. Add the File Server role service to LON-SVR3 and LON-SVR4.
2. On LON-SVR3, open the Failover Cluster Manager console.
3. In the **Storage** node, click **Disks** and verify that three cluster disks are online.
4. Add **File Server** as a cluster role.
5. Specify **AdatumFS** as **Client Access Name**.
6. Specify **172.16.0.130** as the **IP address** for the cluster role.
7. Select **Cluster Disk 2** as the storage disk for AdatumFS role.

#### ► Task 2: Add a shared folder to a highly-available file server

1. On LON-SVR4, open Failover Cluster Manager.
2. Start the New Share Wizard and add a new shared folder to the AdatumFS cluster role.
3. Specify the File share profile as **SMB Share – Quick**.
4. Name the shared folder as **Docs**.

#### ► Task 3: Configure failover and failback settings

1. On LON-SVR4, in the Failover Cluster Manager, open the **Properties** for the **AdatumFS** cluster role.
2. Enable failback between **4 and 5** hours.
3. Select both **LON-SVR3** and **LON-SVR4** as the preferred owners.
4. Move **LON-SVR4** to be first in the Preferred Owners list.

**Results:** After this exercise, you will have configured a highly-available file server.

## Exercise 3: Validate the Deployment of the Highly-Available File Server

### Scenario

In the process of implementing failover cluster, you want to perform failover and failback tests.

The main tasks for this exercise are as follows:

1. Validate the highly-available file server deployment.
2. Validate the failover and quorum configuration for the File Server role.

#### ► Task 1: Validate the highly-available file server deployment

1. On LON-DC1, open Windows Explorer, and attempt to access the **\\AdatumFS\** location. Make sure that you can access the **Docs** folder.
2. Create a test text document inside this folder.
3. On LON-SVR3, in the Failover Cluster Manager, move **AdatumFS** to the second node.
4. On LON-DC1, in Windows Explorer, verify that you can still access **\\AdatumFS\** location.

#### ► Task 2: Validate the failover and quorum configuration for the File Server role

1. On LON-SVR3, determine the current owner for the **AdatumFS** role.
2. Stop the Cluster service on the node that is the current owner of the **AdatumFS** role.
3. Verify that **AdatumFS** has moved to another node and that the **\\AdatumFS\** location is still available.
4. Start the Cluster service on the node in which you stopped it in step 2.
5. Browse to the Disks node, and take the disk witness offline.
6. Verify that **AdatumFS** is still available.
7. Bring the disk witness online.

**Results:** After this exercise, you will have tested the failover scenarios.

## Exercise 4: Configuring Cluster-Aware Updating on the Failover Cluster

### Scenario

Earlier, implementing updates to servers with critical service was causing unwanted downtime. To enable seamless and zero downtime cluster updating, you want to implement the Cluster-Aware Updating feature and test updates for cluster nodes.

The main tasks for this exercise are as follows:

1. Configure Cluster-Aware Updating.
2. Update the failover cluster and configure self-updating.

#### ► Task 1: Configure Cluster-Aware Updating

1. On LON-DC1, install the Failover Clustering feature.
2. From Server Manager, open Cluster-Aware Updating.
3. Connect to **Cluster1**.
4. Preview the updates available for nodes in **Cluster1**.

► **Task 2: Update the failover cluster and configure self-updating**

1. On LON-DC1, start the update process for **Cluster1**.
2. After the process is complete, configure **self-updating for Cluster1**, to be performed **weekly**, on **Sundays** at **4A.M.**

**Results:** After this exercise, you will have configured Cluster-Aware Updating.

► **To prepare for next module**

When you finish the lab, revert the virtual machines to their initial state. To do this, perform the following steps:

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20417A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20417A-LON-SVR1**, **20417A-LON-SVR3**, **MSL-TMG1** and **20417A-LON-SVR4**.



## Module Review and Takeaways

### Review Questions

**Question:** Why is using a Disk-Only quorum configuration generally not a good idea?

**Question:** What is the purpose of Cluster-Aware Updating?

**Question:** What is the main difference between synchronous and asynchronous replication in a multi-site cluster scenario?

**Question:** What is an enhanced feature in multi-site clusters in Windows Server 2012?

### Best Practices

- Try to avoid using quorum model that depends just on disk
- Use Cluster Shared Volumes for Hyper-V high availability or Scale Out File server
- Do regular backups of cluster configuration
- Be sure that, in case of one node failure, other nodes can handle the load
- Carefully plan multi-site clusters

### Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Cluster Validation wizard reports and error	
Create cluster wizard reports that not all nodes support desired clustered role	
You can't create Print Server cluster	

### Real-world Issues and Scenarios

Your organization is considering the use of a geographically-dispersed cluster that includes an alternative data center. Your organization has only a single physical location together with an alternative data center. Can you provide an automatic failover in this configuration?

### Tools

The tools for implementing fail-over clustering include:

- Failover Cluster Manager console
- Cluster-Aware Updating console
- Windows PowerShell
- Server Manager
- iSCSI initiator
- Disk Management

**MCT USE ONLY. STUDENT USE PROHIBITED**

# Module 8

## Implementing Hyper-V

### Contents:

Module Overview	8-1
<b>Lesson 1:</b> Configuring Hyper-V Servers	8-2
<b>Lesson 2:</b> Configuring Hyper-V Storage	8-8
<b>Lesson 3:</b> Configuring Hyper-V Networking	8-16
<b>Lesson 4:</b> Configuring Hyper-V Virtual Machines	8-21
<b>Lab:</b> Implementing Server Virtualization with Hyper-V	8-27
Module Review and Takeaways	8-33

## Module Overview

Although server virtualization was deployed rarely on corporate networks only a decade ago, today it is a core networking technology. Server administrators must be able to distinguish which server workloads might run effectively in virtual machines and which need to remain in a traditional, physical deployment.

This module introduces you to the new features of the Hyper-V® role, the components of the role, and the best practices for deploying the role.

### Objectives

After completing this module, you will be able to:

- Configure Hyper-V servers.
- Configure Hyper-V storage.
- Configure Hyper-V networking.
- Configure Hyper-V virtual machines.

## Lesson 1

# Configuring Hyper-V Servers

The Hyper-V role has undergone a substantial change in Windows Server® 2012. New features, such as network utilization and Resource Metering, provide you with the ability to manage virtual machines effectively with Hyper-V version 3.0. In this lesson, you will learn about the new features in Hyper-V, as well as Hyper-V Integration Services and the factors that you need to consider when you are configuring Hyper-V hosts.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe the new features in Hyper-V 3.0.
- Describe the hardware requirements for Hyper-V 3.0.
- Configure Hyper-V settings.
- Describe Hyper-V Integration services.
- Describe the best practices for configuring Hyper-V hosts.

### What's New in Hyper-V 3.0?

The Hyper-V role first became available after the release of Windows Server 2008. New features were added to the role, both in Windows Server 2008 R2 and Windows Server 2008 R2 Service Pack 1 (SP1).

Hyper-V in Windows Server 2012, also known as Hyper-V 3.0, includes the following major improvements:

- Virtual machine replication
- Hyper-V PowerShell support
- Quality of Service (QoS) bandwidth management
- Non-Uniform Memory Access (NUMA) support
- Memory improvements

The new features in Hyper-V include:

- Virtual machine replication
- Hyper-V PowerShell support
- QoS bandwidth management
- NUMA support
- Memory improvements
- Resource Metering
- Virtual Fibre Channel
- Live migration without shared storage
- New virtual hard disk format
- SMB 3.0 storage
- Network virtualization

### Virtual Machine Replication

You can use Hyper-V replica to perform continuous replication of important virtual machines from a host server to a replica server. In the event that the host server fails, you can configure failover to the replica server. For more information on Hyper-V replicas, visit Module 9: Implementing Failover Clustering with Hyper-V.

### Hyper-V PowerShell support

Windows Server 2012 introduces extensive Windows PowerShell® support for Hyper-V through the Hyper-V PowerShell module. You can manage all aspects of Hyper-V, including creating virtual hard disks, virtual switches, and virtual machines.

## Quality of Service (QoS) Bandwidth Management

Hyper-V administrators can use Quality of Service (QoS) bandwidth management to converge multiple traffic types through a virtual-machine network adapter, which allows a predictable service level for each traffic type. You also can allocate minimum and maximum bandwidth allocations on a per-virtual machine basis.

## Non-Uniform Memory Access (NUMA) Support

Hyper-V 3.0 includes NUMA support. NUMA is a multiprocessor architecture that automatically groups RAM and processors. This leads to performance improvements for virtual machines that are hosted on servers that have multiple processors and large amounts of random access memory (RAM).

## Memory Improvements

Dynamic memory is a feature that lets virtual machine memory to be allocated as necessary, rather than as a fixed amount. For example, rather than setting a virtual machine with a fixed 4 gigabytes (GB) of memory, which Hyper-V allocates to the virtual machine, an administrator can use dynamic memory to allocate a minimum and maximum amount. In this scenario, the virtual machines requests only what it needs. Although Windows Server 2008 R2 SP1 included the ability for virtual machines to use dynamic memory, you had to make any adjustments to these settings after you shut down the server. Hyper-V 3.0 enables administrators to adjust dynamic memory settings on virtual machines that are running. You can use smart paging to configure startup memory, which differs from the minimum and maximum memory allocations. When you use smart paging, the Hyper-V host uses memory paging to ensure that a virtual machine can start when there is not enough memory resources available to support startup, but enough to support the virtual machine's minimum memory allocation.

Other improvements to Hyper-V include:

- *Resource Metering.* Resource Metering allows administrators to track resource utilization of individual virtual machines. You can enable resource metering on a per-virtual machine basis. Use PowerShell to perform resource-metering operations.
- *Virtual Fibre Channel.* Virtual Fibre Channel enables virtual machines to use a virtual Fibre Channel host bus adapter (HBA) to connect to Fibre Channel resources on storage area networks (SANs). To use Virtual Fibre Channel, the host Hyper-V server must have a compatible Fibre Channel HBA.
- *Live migration without shared storage.* Hyper-V 3.0 supports live migration of virtual machines between Hyper-V hosts, without requiring access to shared storage. For more information on live migration, visit Module 9: Implementing Failover Clustering with Hyper-V.
- *New virtual hard disk format.* Hyper-V 3.0 introduces the VHDX format. This disk format supports larger virtual hard disks. It also includes a format that minimizes the chances of data loss during unexpected power outages.
- *Server message block 3.0 (SMB 3.0) storage.* Hyper-V 3.0 virtual machines can use virtual hard disks stored on normal shared folders, as long as the folders are hosted on a server that supports the SMB 3.0 protocol.
- *Network virtualization.* Network virtualization enables virtual machines to retain a static IP address configuration when migrated to different Hyper-V hosts.

## Prerequisites for Installing Hyper-V

Hyper-V on Windows Server 2012 requires that the host computer has an x64 processor, which supports Second Level Address Translation (SLAT). SLAT is a special technology that allows a processor to address memory more efficiently. The server that hosts the Hyper-V role needs a minimum of 4 GB of RAM. A virtual machine hosted on Hyper-V in Windows Server 2012 can support a maximum of 1 terabyte of RAM and up to 32 virtual processors.

When deciding on the server hardware in which you plan to install the Hyper-V role, you need to ensure the following:

- The server must have enough memory to support the memory requirements of all of the virtual machines that must run concurrently. The server also must have enough memory to run the host Windows Server 2012 operating system.
- The storage subsystem performance must meet the I/O needs of the guest virtual machines. It may be necessary to place different virtual machines on separate physical disks to deploy a high performance redundant array of independent disks (RAID), Solid State Drives (SSD), hybrid-SSD, or a combination of all three.
- The CPU capacity of the host server must meet the requirements of the guest virtual machines.
- The host server's network adapters must be able to support the network throughput requirements of the guest virtual machines. This may require installing multiple network adapters and using multiple network interface card (NIC) teams for virtual machines that have high network-use requirements.

Server hardware must support:

- Hardware assisted virtualization
- Data execution prevention
- SLAT



Hardware must be adequate to support the needs of virtual machines with respect to:

- Memory
- Disk I/O
- Processing capability
- Network throughput (typically multiple NICs)

## Demonstration: Configuring Hyper-V Settings

It is necessary to start a traditionally deployed server to run this demonstration because you cannot run Hyper-V from within a virtual machine.

### Demonstration Steps

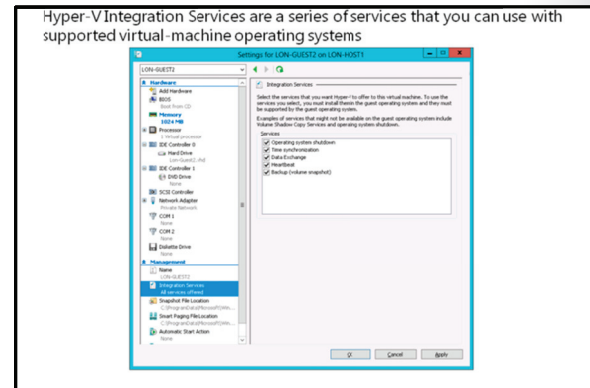
1. Log on to LON-HOST1.
2. Open the Hyper-V Manager console.
3. In the **Hyper-V Settings** dialog box, review the following settings:
  - Virtual Hard Disks
  - Virtual Machines
  - Physical GPUs
  - NUMA Spanning


## Hyper-V Integration Services

Hyper-V Integration Services are a series of services that you can use with supported virtual-machine guest operating systems. Supported operating systems can use Integration Services components and functionality like Small Computer System Interface (SCSI) adapters and synthetic network adapters.

The virtual-machine guest operating systems that Hyper-V supports include:

- Windows Server 2012
- Windows Server 2008 R2 with SP1
- Windows Server 2008 with Service Pack 2 (SP2)
- Windows Server 2003 R2 with SP2
- Windows Home Server 2011
- Windows MultiPoint Server 2011
- Windows Small Business Server 2011
- Windows Server 2003 with Service Pack 2
- CentOS 6.0-6.2
- CentOS 5.5-5.7
- Red Hat Enterprise Linux 6.0-6.2
- Red Hat Enterprise Linux 5.5-5.7
- SUSE Linux Enterprise Server 11 with Service Pack 1 or Service Pack 2
- SUSE Linux Enterprise Server 10 with Service Pack 4
- Windows 7 with Service Pack 1
- Windows Vista® with Service Pack 2
- Windows XP with Service Pack 3



 **Additional Reading:** Note that the Hyper-V support for the Windows XP operating system ends in April 2014, and support for Windows Server 2003 and Windows Server 2003 R2 expires in July 2015. When available, a link will be provided here to the list of supported Hyper-V virtual-machine guest operating systems on Windows Server 2012.

You can install the Integration Services components on an operating system by clicking the Insert Integration Services Setup Disk item on the Action menu in the Virtual Machine Connection window. After this is done, you can install the relevant operating-system drivers either manually or automatically.

You can enable the following virtual-machine integration components:

- *Operating system shutdown.* The Hyper-V server uses this component to initiate a graceful shutdown of the guest virtual machine.
- *Time synchronization.* The virtual machine uses this component to use the host server's processor to conduct time synchronization.
- *Data Exchange.* The Hyper-V host uses this component to write data to the virtual machine's registry.
- *Heartbeat.* Hyper-V uses this component to determine if the virtual machine has become unresponsive.
- *Backup (volume snapshot).* The provider of the Volume Shadow Copy Service (VSS) uses this component to create virtual-machine snapshots for backup operations, without interrupting the virtual machines' normal operation.

## Best Practices for Configuring Hyper-V Hosts

There are several best practices that you should consider when provisioning Windows Server 2012 to function as a Hyper-V host:

- Provision the host with adequate hardware
- Deploy virtual machines on separate disks
- Do not collocate other server roles
- Manage Hyper-V remotely
- Run Hyper-V by using the Server Core configuration
- Run the Best Practices Analyzer and Resource Metering

The best practices when configuring Hyper-V hosts include:

- Provision the Hyper-V host with adequate hardware resources
- Deploy virtual machines on separate disks
- Do not collocate other Windows Server 2012 roles on the Hyper-V host
- Manage Hyper-V remotely
- Run Hyper-V by using the Server Core configuration
- Use Resource Metering and Best Practices Analyzer

### Provision the Host with Adequate Hardware

Perhaps the most important best practice is to ensure that the Hyper-V host is provisioned with adequate hardware. You should ensure that there is appropriate processing capacity, an appropriate amount of RAM, and fast and redundant storage. You should ensure that the Hyper-V host is provisioned with multiple network cards that you configure as a team. If the Hyper-V host is not provisioned adequately with hardware, this has an effect on the performance of all virtual machines that are hosted on the server.

### Deploy Virtual Machines on Separate Disks

You should use separate disks to host virtual-machine files rather than having virtual-machine files stored on the same disk as the host operating-system files. This minimizes contention and ensures that read/write operations occurring on virtual machine files do not conflict with read/write operations occurring at the host operating-system level. It also minimizes the chance that the virtual-machine hard disks will grow to consume all available space on the operating-system volume. Performance considerations are lessened if you deploy to a disk that uses striping, such as a RAID 1+0 array. If you are using shared storage, you can provision multiple virtual machines on the same Logical Unit Number (LUN) if you utilize Cluster Shared Volumes. However, choosing between separate LUNs for each virtual machine or a shared LUN depends heavily on virtual machine workload and SAN hardware.



## Do Not Colocate Other Server Roles

You should ensure that Hyper-V is the only server role deployed on the server. You should not colocate the Hyper-V role with other roles, such as the Domain Controller or File Server role. Each role that you deploy on a server requires resources, and when deploying Hyper-V, you want to ensure that the virtual machines have access to as much of a host server's resources as possible. If it is necessary to locate these roles on the same hardware, deploy these roles as virtual machines rather than installing them on the physical host.

## Manage Hyper-V Remotely

When you log on locally to a server, your logon session consumes server resources. By configuring a Hyper-V server to be managed remotely and not performing administrative tasks by logging on locally, you ensure that all possible resources on the Hyper-V host are available to the hosted virtual machines. You also should restrict access to the Hyper-V server, so that only administrators responsible for the management of virtual machines can make connections. A configuration error on a Hyper-V host can cause downtime to all hosted virtual machines.

## Run Hyper-V by Using the Server Core Configuration

There are two main reasons to run Hyper-V using the Server Core configuration. The first reason is that running Windows Server 2012 in the server core configuration minimizes hardware-resource utilization for the host operating-system. Running the server in server core configuration means that there are more hardware resources for the hosted virtual machines.

The second reason to run the Hyper-V server in server core configuration is that server core requires fewer software updates, which in turn means fewer reboots. When you restart a Hyper-V host, all virtual machines that the server hosts become unavailable when it is unavailable. Because a Hyper-V host can host many critical servers as virtual machines, you want to ensure that you minimize downtime.

## Run the Best Practices Analyzer and Use Resource Metering

If you have enabled performance counters on the Hyper-V host, you can use the Best Practices Analyzer to determine if there are any specific configuration issues that you should address. Enabling performance counters does incur a slight cost to performance, so you should enable these only during periods when you want to monitor server performance, rather than leaving them on permanently.

You can use Resource Metering, a new feature of Hyper-V 3.0, to monitor how hosted virtual machines utilize server resources. You can use Resource Metering to determine if specific virtual machines are using a disproportionate amount of a host server's resources. If the performance characteristics of one virtual machine are having a deleterious effect on the performance of other virtual machines hosted on the same server, you should consider migrating that virtual machine to another Hyper-V host.



**Additional Reading:** 7 Best Practices for Physical Servers Hosting Hyper-V Roles  
<http://technet.microsoft.com/en-us/magazine/dd744830.aspx>

## Lesson 2

# Configuring Hyper-V Storage

Hyper-V provides many different virtual machine storage options. If you know which option is appropriate for a given situation, you can ensure that a virtual machine performs well. But if you do not understand the different virtual-machine storage options, you may end up deploying virtual hard disks that consume unnecessary space or that place an unnecessary performance burden on the host Hyper-V server.

This lesson describes about different virtual hard disk types, different virtual hard disk formats, and the benefits and limitations of using virtual machine snapshots.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe the properties of virtual hard disks in Hyper-V 3.0.
- Select a virtual hard disk type.
- Convert between virtual hard disk types.
- Maintain virtual hard disks.
- Determine where to deploy virtual hard disks.
- Describe the requirements for storing Hyper-V data on SMB file shares.
- Implement virtual machine snapshots.
- Describe the requirements of providing Fibre Channel support within virtual machines.

### Virtual Hard Disks in Hyper-V 3.0

A virtual hard disk is a special file format that represents a traditional hard-disk drive. You can configure a virtual hard disk with partitions and an operating system. Additionally, you can use virtual hard disks with virtual machines and you also can mount virtual hard disks by using the Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, and Windows® 8 and Windows 7 operating systems. Windows Server 2012 supports booting to virtual hard disks. You can use this feature to configure the computer to start into a Windows Server 2012 operating system or some editions of the Windows Server 8 operating system that are deployed on a virtual hard disk. You can create a virtual hard disk by using:

- The Hyper-V manger console.
- The Disk Management console.
- The diskpart command-line utility.
- The **New-VHD** Windows PowerShell cmdlet.



**Note:** Some editions of Windows 7 and the Windows Server 2008 R2 operating system also support booting to virtual hard disk.

Windows Server 2012 introduces the new VHDX format for virtual hard disks, which:

- Supports bigger disks
- Is less likely to lose data during unexpected outages
- Supports better alignment when deployed to a large sector disk
- Allows larger block size for dynamic and differencing disks, which provides better performance



## Comparing VHDX and VHD

Virtual hard disks use the .vhd extension. Windows Server 2012 introduces the new VHDX format for virtual hard disks. In comparison to the VHD format that was used in Hyper-V on Windows Server 2008 and Windows Server 2008 R2, the VHDX format has the following benefits:

- VHDX virtual hard disks can be as large as 64 terabytes. VHD virtual hard disks were limited to 2 TB.
- The VHDX virtual hard disk file structure minimizes the chance that the disk will become corrupt if the host server suffers an unexpected power outage.
- VHDX virtual hard disk format supports better alignment when deployed to large sector disk.
- VHDX allows larger block size for dynamic and differencing disks, which provides better performance for these workloads.

If you have upgraded a Windows Server 2008 or Windows Server 2008 R2 Hyper-V server to Windows Server 2012, you can convert an existing VHD file to VHDX format by using the Edit Disk tool. It also is possible to convert from VHDX format to VHD.



**Additional Reading:** Hyper-V Virtual Hard Disk Format Overview

<http://technet.microsoft.com/en-us/library/hh831446.aspx>

## Disk Types

When you configure a virtual hard disk, you can choose one of the following disk types:

- Fixed
- Dynamic
- Pass-through
- Differencing

### Fixed Virtual Hard Disk

When you create a fixed virtual hard disk, all of the hard-disk space is allocated during the creation process. This has the advantage of minimizing fragmentation, which improves virtual hard disk performance when they are hosted on traditional storage devices. However, a disadvantage is that it requires all of the space that the virtual hard disk potentially can use to be allocated on the host partition. In many situations, you will not know precisely how much disk space a virtual machine needs. If you use fixed hard disks, you may end up allocating space to storage that is not actually required.

To create a fixed virtual hard disk, perform the following steps:

1. Open the Hyper-V Manager console.
2. In the Actions pane, click **New**, and then click **Hard Disk**.
3. On the **Before You Begin** page of the New Virtual Hard Disk Wizard, click **Next**.
4. On the **Choose Disk Format** page, select **VHD** or **VHDX**, and then click **Next**.
5. On the **Choose Disk Type** page, click **Fixed size**, and then click **Next**.
6. On the **Specify Name and Location** page, enter a name for the virtual hard disk, and then specify a folder to host the virtual hard-disk file.

Type of disk	Description
Fixed	All of the hard disk space is allocated during the creation process
Dynamic	The disk itself only uses the amount of space that needs to be allocated and grows as necessary
Pass Through	Virtual machines access a physical disk drive rather than use a virtual hard disk
Differencing	The amount of hard disk space consumed by virtual hard disks is reduced at the cost of disk performance

7. On the **Configure Disk** page, select one of the following options:
  - Create a new blank virtual hard disk of the specified size.
  - Copy the contents of a specified physical disk. You can use this option to replicate an existing physical disk on the server as a virtual hard disk. The fixed hard disk will be the same size as the disk that you have replicated. Replicating an existing physical hard disk does not alter data on the existing disk.
  - Copy the contents of a specified virtual hard disk. You can use this option to create a new fixed hard disk based on the contents of an existing virtual hard disk.

You can create a new fixed hard disk by using the **New-VHD** Windows PowerShell cmdlet with the *-Fixed* parameter.



**Note:** Disk fragmentation is less of an issue when virtual hard disks are hosted on RAID volumes or on SSDs. Hyper-V improvements, since it was first introduced in Windows Server 2008, also minimize performance differences between dynamic and fixed virtual hard disks.

## Dynamic Disks

When you create a dynamic virtual hard disk, you specify a maximum size for the file. The disk itself only uses the amount of space that needs to be allocated, and it grows as necessary. For example, if you create a new virtual machine, and specify a dynamic disk, only a small amount of disk space is allocated to the new disk.

This space is as follows:

- Approximately 260 kilobytes (KB) for a VHD format virtual hard disk
- Approximately 4096 KB for a VHDX format virtual hard disk

As storage is allocated, such as when you deploy the operating system, the dynamic hard disk grows. If you delete files from a dynamically expanding virtual hard disk, the virtual hard-disk file does not shrink. You can only shrink a dynamically expanding virtual hard-disk file by performing a shrink operation.

Creating a dynamically expanding virtual hard disk is similar to creating a fixed disk. In the New Virtual Hard Disk Wizard, on the **Choose Disk Type** page, select **Dynamically expanding size** instead of **Fixed**.

You can create a new dynamic hard disk by using the **New-VHD** Windows PowerShell cmdlet with the *-Dynamic* parameter.

## Pass-Through Disks

Virtual machines use the pass-through disks to access a physical disk drive, rather than use a virtual hard disk. You can use pass-through disks to connect a virtual machine directly to an Internet SCSI (iSCSI) LUN. When you use pass-through disks, the virtual machine must have exclusive access to the target disk. To do this, you must use the host's disk management console to take the disk offline. After the disk is offline, you can connect it to one of the virtual machine's disk controllers.

You can attach a pass-through disk by performing the following steps:

1. Ensure that the target hard disk is offline.
2. Use the Hyper-V Manager console to edit an existing virtual machine's properties.
3. Click an Integrated Drive Electronics (IDE) or SCSI controller, click **Add**, and then click **Hard Drive**.
4. In the **Hard Drive** dialog box, select **Physical Hard Disk**. In the drop-down list, select the disk that you want to use as the pass-through disk.



**Note:** You do not have to shut down a virtual machine if you connect the pass-through disk to a virtual machine's SCSI controller. However, if you want to connect to a virtual machine's IDE controller, it is necessary to shut down the virtual machine.

## Differencing disks

Differencing disks record the changes made to a parent disk. You can use differencing disks to reduce the amount of hard disk space that virtual hard disks consume, but that comes at the cost of disk performance. Differencing disks work well with SSD where there is limited space available on the drive and the performance of the disk compensates for the performance drawbacks of using a differencing disk.

Differencing disks have the following properties:

- You can link multiple differencing disks to a single parent disk.
- When you modify the parent disk, all linked differencing disks fail.

You can reconnect a differencing disk to the parent by using the Inspect Disk tool, available in the actions pane of the Hyper-V Manager console. You also can use the Inspect Disk tool to locate a differencing disk's parent disk.

To create a differencing disk, follow these steps:

1. Open the Hyper-V Manager console.
2. In the Actions pane, click **New**, and then click **Hard Disk**.
3. On the **Before You Begin** page of the New Virtual Hard Disk Wizard, click **Next**.
4. On the **Choose Disk Format** page, select **VHD**, and then click **Next**.
5. On the **Choose Disk Type** page, select **Differencing**, and then click **Next**.
6. On the **Specify Name and Location** page, provide the location of the parent hard disk, and then click **Finish**.

You can create a differencing hard disk by using the **New-VHD** Windows PowerShell cmdlet. For example, to create a new differencing disk named `c:\diff-disk.vhd` that uses the virtual hard disk `c:\parent.vhd`, run the following Windows PowerShell command:

```
New-VHD c:\diff-disk.vhd -ParentPath C:\parent.vhd
```

## Converting Disks

From time to time, it is necessary to perform maintenance operations on virtual hard disks. You can perform the following maintenance operations on virtual hard disks:

- Convert the disk from fixed to dynamic.
- Convert the disk from dynamic to fixed.
- Convert a virtual hard disk in VHD format to VHDX.
- Convert a virtual hard disk in VHDX format to VHD.

- You can perform the following maintenance operations on virtual hard disks:
  - Convert the disk from fixed to dynamic
  - Convert the disk from dynamic to fixed
  - Convert a virtual hard disk in VHD format to VHDX
  - Convert a virtual hard disk in VHDX format to VHD

When you convert a hard disk, the contents of the existing virtual hard disk are copied to a new virtual hard disk that has the properties that you have chosen. To convert a virtual hard disk, perform the following steps:

1. In the Actions pane of the Hyper-V Manager console, click **Edit Disk**.
2. On the **Before You Begin** page of the Edit Virtual Hard Disk Wizard, click **Next**.
3. On the **Local Virtual Hard Disk** page, click **Browse**. Select the virtual hard disk that you wish to convert.
4. On the **Choose Action** page, select **Convert**, and then click **Next**.
5. On the **Convert Virtual Hard Disk** page, select **VHD** or **VHDX** format. By default, the current disk format is selected. Click **Next**.
6. If you want to convert the disk from fixed to dynamic or dynamic to fixed, on the **Convert Virtual Hard Disk** page, select **Fixed Size** or **Dynamically Expanding**. If you want to convert the hard disk type, choose the appropriate type, and then click **Next**.
7. On the **Configure Disk** page, select the destination location for the disk, click **Next**, and then click **Finish**.

You can shrink a dynamic virtual hard disk that is not taking up all the space that is allocated to it. For example, a dynamic virtual hard disk might be 60 GB on the parent volume, but only use 20 GB of that space. You shrink a virtual hard disk by choosing the Compact option in the Edit Virtual Hard Disk Wizard.

You cannot shrink fixed virtual hard disks. You must convert a fixed virtual hard disk to dynamic before you can compact the disk. You can use the **resize-partition** and the **resize-vhd** Windows PowerShell cmdlets to compact a dynamically expanding virtual hard disk.

You also can use the Edit Virtual Hard Disk Wizard to expand a disk. You can expand both dynamically expanding and fixed virtual hard disks.

## Demonstration: Managing Virtual Hard Disks in Hyper-V

In this demonstration, you create a differencing disk based on an existing disk by using both Hyper-V Manager and PowerShell.

### Demonstration Steps

1. Use Windows Explorer to create the following folders on the physical host drive:
  - o **E:\Program Files\Microsoft Learning\Base \LON-GUEST1**
  - o **E:\Program Files\Microsoft Learning\Base \LON-GUEST2**



**Note:** The drive letter may depend upon the number of drives on the physical host machine)

2. In the Hyper-V Manager console, create a virtual hard disk with the following properties:
  - o Disk Format: **VHD**
  - o Disk Type: **Differencing**
  - o Name: **LON-GUEST1.vhd**
  - o Location: **E:\Program Files\Microsoft Learning\Base\LON-GUEST1\**
  - o Parent Location: **E:\Program Files\Microsoft Learning\Base\Base12A-WS2012-RC.vhd**

3. Open Windows PowerShell, import the Hyper-V module, and then run the following command:

```
New-VHD "E:\Program Files\Microsoft Learning\Base\LON-GUEST2\LON-GUEST2.vhd"
-ParentPath "E:\Program Files\Microsoft Learning\Base\Base12A-WS2012-RC.vhd"
```

4. Inspect disk **E:\Program Files\Microsoft Learning\Base\LON-GUEST2\LON-GUEST2.vhd**.
5. Verify that **LON-GUEST2.vhd** is configured as a differencing virtual hard disk with **E:\Program Files\Microsoft Learning\Base\Base12A-WS2012-RC.vhd** as a parent.

## Location Considerations of Virtual Hard Disks

A key factor when provisioning virtual machines is ensuring that virtual hard disks are placed correctly. Virtual hard-disk performance can affect virtual machine performance dramatically. Servers that are otherwise well provisioned with RAM and processor capacity can still experience bad performance if the storage system is overwhelmed.

Consider the following factors when planning the location of virtual hard-disk files:

- **High-performance connection to the storage**

You can locate virtual hard-disk files on local or remote storage. When you locate them on remote storage, you need to ensure that there is adequate bandwidth and minimal latency between the host and the remote storage. Slow network connections to storage, or connections where there is latency, result in poor virtual-machine performance.

- **Redundant storage**

The volume that the virtual hard-disk files are stored on should be fault-tolerant. This should apply if the virtual hard disk is stored on a local disk or a remote SAN device. It is not uncommon for hard disks to fail. Therefore, the virtual machine and the Hyper-V host should remain in operation after a disk failure. Replacement of failed disks also should not affect the operation of the Hyper-V host or virtual machines.

- **High-performance storage**

The storage device on which you store virtual hard-disk files should have excellent I/O characteristics. Many enterprises use SSD hybrid drives in RAID 1+0 arrays to achieve maximum performance and redundancy. Multiple virtual machines that are running simultaneously on the same storage can place a tremendous I/O burden on a disk subsystem. Therefore, you need to ensure that you choose high-performance storage. If you do not, virtual machine performance suffers.

- **Adequate growth space**

If you have configured virtual hard disks to grow automatically, ensure that there is adequate space into which the files can grow. Also, carefully monitor growth so that you are not shocked when a virtual hard disk fills the volume that you allocated to host it. If you configure virtual hard disks to grow automatically, place each virtual machine's virtual hard disk on a separate volume. This way, the virtual hard disks of multiple virtual machines are not affected if the volume's capacity is exceeded.

When planning the location of virtual hard disks, ensure the following:

- Virtual hard disk files are stored on disks that can be accessed quickly from the Hyper-V host
- Virtual hard disk files are stored on a volume that is configured for redundancy
- Virtual hard disk files are stored on high-performance storage
- Virtual hard disk files configured for growth should be placed on volumes with adequate space



## Storage on SMB 3 File Shares

Hyper-V supports storing virtual machine data, such as virtual-machine configuration files, snapshots, and virtual hard-disk files, on SMB 3 file shares.

The file share must support SMB 3. This limits placement of virtual hard disks on file shares that are hosted on file servers that are running Windows Server 2012. Earlier Windows Server versions do not support SMB 3.

You must ensure that network connectivity to the file share is 1 GB or more.

SMB 3 is available in Windows Server 2012 only, and not in earlier Windows Server versions

Hyper-V 3.0 can store the following on SMB 3 file shares:

- Configuration files
- Virtual hard disk files (in VHD or VHDX format)
- Snapshot files

SMB file share provides an alternative to storing virtual-machine files on iSCSI or Fibre Channel SAN devices. When creating a virtual machine in Hyper-V on Windows Server 2012, you can specify a network share when choosing the virtual machine location and the virtual hard-disk location. You also can attach disks stored on SMB 3 file shares. You can use both VHD and VHDX disks with SMB file shares.



**Additional Reading:** Server Message Block overview  
<http://technet.microsoft.com/en-us/library/hh831795.aspx>

## Snapshot Management in Hyper-V

Snapshot is an important technology that provides administrators with the ability to make a replica of a virtual machine at a specific time. You can take snapshots when a virtual machine is shut down or running. However, when you take a snapshot of a virtual machine that is running, the snapshot includes the contents of the virtual machine's memory.

### Taking a Snapshot

You can take a snapshot on the Actions pane of the Virtual Machine Connection window or in the Hyper-V Manager console. Each virtual machine can have a maximum of 50 snapshots.

When taking snapshots of multiple virtual machines, you should take them at the same time. This ensures synchronization of items such as computer-account passwords. Remember that when you revert to a snapshot, you are reverting to a computer's state at that specific time. If you take a computer back to a point before it performed a computer-password change with a domain controller, you will need to rejoin that computer to the domain.

Snapshots provide administrators with the ability to make a replica of a virtual machine at a particular point in time

Snapshots do not replace backups

- Snapshots are written as avhdx files, which merge back in the previous snapshot when the snapshot is deleted
- Snapshot of running virtual machine includes the contents of memory





## Snapshots Do Not Replace Backups

Snapshots are not a replacement for backups. Snapshot data is stored on the same volume as the virtual hard disks. If the volume hosting these files fails, both the snapshot and the virtual hard disk files are lost. You can perform a virtual machine export of a snapshot. When you export the snapshot, Hyper-V creates full virtual hard disks that represent the state of the virtual machine at the time that you took the snapshot. If you choose to export an entire virtual machine, all snapshots associated with the virtual machine also are exported.

### Avhd files

When you create a snapshot, Hyper-V writes avhd files that store the data that differentiates the snapshot from either the previous snapshot or the parent virtual hard disk. When you delete snapshots, this data is discarded or merged into the previous snapshot or parent virtual hard disk. For example, if you delete the most recent snapshot of a virtual machine, the data is discarded. If you delete the second to last snapshot taken of a virtual machine, the data is merged so that the earlier and latter snapshot states of the virtual machine retain their integrity.

### Managing Snapshots

When you apply a snapshot, the virtual machine reverts to the configuration as it existed at the time that the snapshot was taken. Reverting to a snapshot does not delete any existing snapshots. If you revert to a snapshot after making a configuration change, you are prompted to take a snapshot. It only is necessary to create a new snapshot if you want to return to that current configuration.

It is possible to create snapshot trees that have different branches. For example, if you took a snapshot of a virtual machine on Monday, Tuesday, and Wednesday, applied the Tuesday snapshot, and then made changes to the virtual machine's configuration, you create a new branch that diverts from the original Tuesday snapshot. You can have multiple branches as long as you do not exceed the 50-snapshot limit per virtual machine.

## Fibre Channel Support in Hyper-V

Hyper-V virtual Fibre Channel is a virtual hardware component that you can add to a virtual machine, and which enables the virtual machine to access Fibre Channel storage on SANs. To deploy a virtual Fibre Channel:

- You must configure the Hyper-V host with a Fibre Channel HBA.
- The Fibre Channel HBA must have a driver that supports virtual Fibre Channel.
- The virtual machine must support virtual machine extensions.

#### The Fibre Channel adapter:

- Allows a virtual machine to directly connect to a Fibre Channel SAN
- Requires that the Hyper-V host has a Fibre Channel HBA
- Requires that the Fibre Channel HBA driver supports virtual Fibre Channel

Virtual Fibre Channel adapters support port virtualization by exposing HBA ports in the guest operating system. This allows the virtual machine to access the SAN by using a standard World Wide Name (WWN) associated with the virtual machine.

You can deploy up to four virtual Fibre Channel adapters to each virtual machine.



**Additional Reading:** Hyper-V Virtual Fibre Channel Overview

<http://technet.microsoft.com/en-us/library/hh831413.aspx>

## Lesson 3

# Configuring Hyper-V Networking

Hyper-V provides several different options for allowing network communication between virtual machines. You can use Hyper-V to configure virtual machines that communicate with an external network in a manner similar to physical hosts that you deploy traditionally. You also can use Hyper-V to configure virtual machines that are able to communicate only with a limited number of other virtual machines hosted on the same Windows Server 2012 Hyper-V host. This lesson describes the various options available for Hyper-V virtual networks, which you can leverage to best meet your organization's needs.

### Lesson Objectives

After completing this lesson, you will be able to:


- Describe the new features in Hyper-V networking.
- Describe virtual switches.
- Configure a public and private switch.
- Describe network virtualization.
- Describe the best practices for configuring virtual networks.

### What's New in Hyper-V Networking?

There are several new features in Hyper-V 3.0 networking that improve the network performance of a large number of virtual machines in private and public cloud environments. In most cases, you should use the default settings in small scale deployments.

The new features in Hyper-V 3.0 networking include:

- *Network virtualization.* This feature enables IP addresses to be virtualized in hosting environments so that virtual machines migrated to the host can keep their original IP address rather than being allocated an IP address on the Hyper-V server's network.
- *Bandwidth management.* You can use this feature to specify a minimum and a maximum bandwidth to be allocated to the adapter by Hyper-V. Hyper-V reserves the minimum bandwidth allocation for the network adapter, even when other virtual network adapters on virtual machines hosted on the Hyper-V host are functioning at capacity.
- *Dynamic Host Configuration Protocol (DHCP) guard.* This feature drops DHCP messages from virtual machines that are functioning as unauthorized DHCP servers. This may be necessary in scenarios where you are managing a Hyper-V server that hosts virtual machines for others, but in which you do not have direct control over the virtual machines' configuration.
- *Router guard.* This feature drops router advertisement and redirection messages from virtual machines configured as unauthorized routers. This may be necessary in scenarios where you do not have direct control over the configuration of virtual machines.

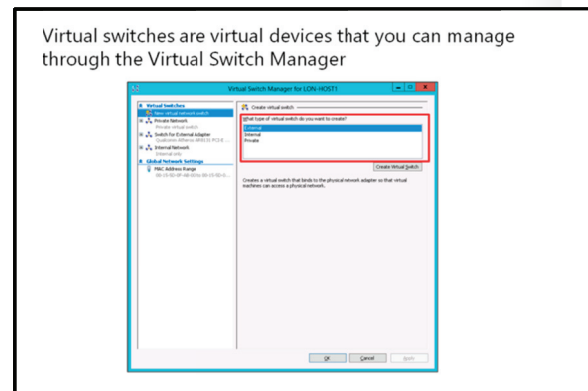


- *Port mirroring.* You can use this feature to copy incoming and outgoing packets from a network adapter to another virtual machine that you have configured for monitoring.
- *NIC teaming.* You can use this feature to add the virtual network adapter to an existing team on the host Hyper-V server.
- *Virtual Machine Queue.* This feature requires that the host computer has a network adapter that supports the feature. Virtual Machine Queue uses hardware packet filtering to deliver network traffic directly to the guest. This improves performance because the packet does not need to be copied from the host operating system to the virtual machine. Only synthetic network adapters support these feature.
- *IP security (IPsec) task offloading.* This feature requires that the guest operating system and network adapter are supported. This feature enables the host's network adapter to perform calculation-intensive security-association tasks. If sufficient hardware resources are not available, the guest operating system performs these tasks. You can configure a maximum number of offloaded security associations between a range of one and 4,096. This feature is supported only on synthetic network adapters.
- *Single-root I/O virtualization (SR-IOV).* This feature requires specific hardware and special drivers to be installed on the guest operating system. SR-IOV enables multiple virtual machines to share the same Peripheral Component Interconnect Express (PCIe) physical hardware resources. If sufficient resources are not available, network connectivity falls back so that the virtual switch provides it. This feature is only supported on synthetic network adapters.

## What Is a Hyper-V Virtual Switch?

Virtual switches are virtual devices that you can manage through the Virtual Switch Manager, which enables you to create three types of virtual switches. The virtual switches control how the network traffic flows between virtual machines hosted on the Hyper-V server, as well as how the network traffic flows between virtual machines and the rest of the organizational network.

Hyper-V on Windows Server 2012 supports the three types of virtual switches that the following table details.



Type	Description
External	You use this type of switch to map a network to a specific network adapter or network-adapter team. Windows Server 2012 supports mapping an external network to a wireless network adapter, if you have installed the Wireless LAN Service on the host Hyper-V server, and the Hyper-V server has a compatible adapter.
Internal	You use internal virtual switches to communicate between the virtual machines on the Hyper-V host and to communicate between the virtual machines and the Hyper-V host itself.
Private	You use private switches only to communicate between virtual machines on the Hyper-V host. You cannot use private switches to communicate between the virtual machines and the Hyper-V host.

When configuring a virtual network, you can also configure a virtual LAN (VLAN) ID to be associated with the network. You can use this to extend existing VLANs on the external network to VLANs within the Hyper-V host's network switch. You can use VLANs to partition network traffic. VLANs function as separate logical networks. Traffic can pass only from one VLAN to another if it passes through a router.

You can configure the following extensions for each virtual switch type:

- *Microsoft Network Driver Interface Specification (NDIS) Capture.* This extension allows the capture of data travelling across the virtual switch.
- *Microsoft Windows Filtering Platform.* This extension allows filtering of data travelling across the virtual switch.



**Additional Reading:** Hyper-V Virtual Switch Overview  
<http://technet.microsoft.com/en-us/library/hh831452.aspx>

## Demonstration: Configuring Hyper-V Networking


In this demonstration, you will see how to create two types of virtual network switches.

### Demonstration Steps

1. In Hyper-V Manager, use the **Virtual Switch Manager** to create a new **External** virtual network switch with the following properties:
  - o Name: **Corporate Network**
  - o External Network: Mapped to the host computer's physical network adapter. Will vary depending on host computer
2. In Hyper-V Manager, use the **Virtual Switch Manager** to create a new virtual switch with the following properties.
  - o Name: **Private Network**
  - o Connection type: **Private network**

## What Is Network Virtualization?

You can use network virtualization to isolate virtual machines from different organizations, even if they share the same Hyper-V host. For example, you might be providing an Infrastructure as a Service (IaaS) to competing businesses. You can use network virtualization to go beyond assigning these virtual machines to separate VLANs as a way of isolating network traffic. Network virtualization is a technology that you would deploy primarily in scenarios where you use Hyper-V to host virtual machines for third-party organizations. Network virtualization has the advantage that you can configure all network isolation on the Hyper-V host. With VLANs, it also is necessary to configure switches with the appropriate VLAN IDs.



When you configure network virtualization, each guest virtual machine has two IP addresses, which work as follows:

- *Customer IP address.* The customer assigns this IP address to the virtual machine. You can configure this IP address so that communication with the customer's internal network can occur even though the virtual machine might be hosted on a Hyper-V server that is connected to a separate public IP network. Using the **ipconfig** command on the virtual machine shows the customer IP address.
- *Provider IP address.* The hosting provider assigns this IP address, which is visible to the hosting provider and to other hosts on the physical network. This IP address is not visible from the virtual machine.

You can use network virtualization to host multiple machines that use the same customer address, such as 192.168.15.101, on the same Hyper-V host. When you do this, the virtual machines are assigned different IP addresses by the hosting provider, though this address will not be apparent from within the virtual machine.

You manage network virtualization by using PowerShell cmdlets. All Network Virtualization cmdlets are in the NetWNV PowerShell module. Tenants gain access to virtual machines that take advantage of network virtualization through routing and remote access. They make a tunneled connection from their network through to the virtualized network on the Hyper-V server.



**Additional Reading:** Hyper-V Network Virtualization Overview  
<http://technet.microsoft.com/en-us/library/hh831395.aspx>

## Best Practices for Configuring Virtual Networks

Best practices with respect to configuring virtual networks typically revolve around ensuring that virtual machines are provisioned with adequate bandwidth. You do not want to have the performance on all virtual machines affected if a bandwidth-intensive operation, such as a large file copy or website traffic spike, occurs on one virtual machine on the same host.

The following general best practices apply to configuring virtual networks:

- Considerations for NIC teaming. You should deploy multiple network adapters to the Hyper-V host, and then configure those adapters as part of a team. This ensures that network connectivity will be retained if the individual network cards fail. Configure multiple teams connected to different switches to ensure that connectivity remains if a hardware switch fails.
- Considerations for bandwidth management. You can use bandwidth management to allocate a minimum and a maximum bandwidth allocation on a per-virtual-network adapter basis. You should configure bandwidth allocation to guarantee that each virtual machine has a minimum bandwidth allocation. This ensures that if another virtual machine hosted on the same Hyper-V server experiences a traffic spike, other virtual machines are able to communicate with the network normally.

When configuring virtual networks:

- Use NIC teaming on the Hyper-V host to ensure connectivity to virtual machines if an adapter fails
- Enable bandwidth management to ensure that no single virtual machine is able to monopolize the network interface
- Use network adapters that support a Virtual Machine Queue
- Use network virtualization when you have to ensure that virtual machines keep their original IP addresses after migrating to a new host

- Considerations for Virtual Machine Queue. You should provision the Hyper-V host with an adapter that supports Virtual Machine Queue. Virtual Machine Queue uses hardware-packet filtering to deliver network traffic directly to the virtual machine. This improves performance because the packet does not need to be copied from the host operating system to the virtual machine. When you do not configure virtual machines to support Virtual Machine Queue, the host operating system can become a bottleneck when it processes large amounts of network traffic.
- Considerations for network virtualization. Network virtualization is complicated to configure, but has an advantage over VLAN. That is, it is not necessary to configure VLANs on all of the switches that are connected to the Hyper-V host. You can perform all necessary configurations when you need to isolate servers on the Hyper-V host without needing to involve the network team. If you are hosting large numbers of virtual machines, and need to isolate them, use Network Virtualization rather than VLANs.

## Lesson 4

# Configuring Hyper-V Virtual Machines

When planning a server-virtualization strategy, you need to know what you can and cannot accomplish when you are using Windows Server 2012 as a virtual machine host.

In this lesson, you will learn about Hyper-V, the hardware requirements required for deploying Hyper-V on a computer running Windows Server 2012, the different components of a virtual machine, and the benefits of virtual machine Integration Services. You also will learn how to measure virtual machine resource use with Windows PowerShell cmdlets.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe the hardware and management options in virtual machine settings.
- Describe how dynamic memory works in Hyper-V.
- Create a virtual machine.
- Import, export, and move virtual machines in Hyper-V.
- Describe the best practices for configuring virtual networks.

## Overview of Virtual Machine Settings

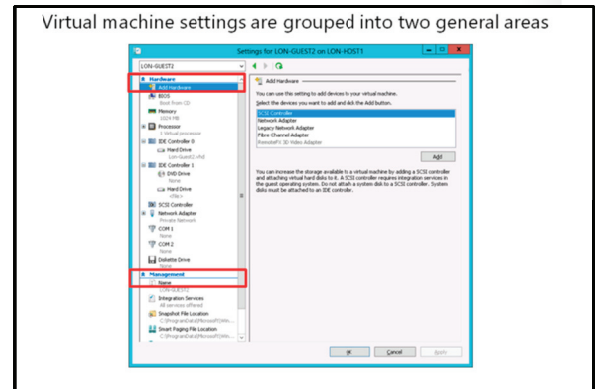
Virtual machine settings are grouped into two general areas: Hardware and Management.

### Hardware

Virtual machines use simulated hardware. The hypervisor uses this virtual hardware to mediate access to actual hardware. For example, you can map a virtual network adapter to a virtual network that, in turn, maps to an actual network interface.

Virtual machines have the following hardware, by default:

- **BIOS.** This virtual hardware simulates the computer's BIOS. You can configure the virtual machine so that Num Lock is switched on or off. You also can choose the boot order for the virtual machine's virtual hardware. You can start a machine from a DVD drive, integrated device electronics (IDE) device, legacy network adapter, or a floppy disk.
- **Memory.** You can allocate memory resources to the virtual machine. An individual virtual machine can allocate as much as 1 terabyte of memory.
- **Processor.** You can allocate processor resources to the virtual machine. You can allocate up to 32 virtual processors to a single virtual machine.
- **IDE Controller.** A virtual machine can support only two IDE controllers. By default, two IDE controllers are allocated to the virtual machine. These are: IDE Controller 0 and IDE Controller 1. Each IDE controller can support two devices. You can connect virtual disks or virtual DVD drives to an IDE controller. If starting from a hard disk drive or DVD-ROM, the boot device must be connected to an IDE controller. Use IDE controllers to connect virtual hard disks and DVD-ROMS to virtual machines that use operating systems that do not support Integration Services.





- *SCSI Controller.* You can use SCSI controllers only on virtual machines that you deploy with operating systems that support Integration Services.
- *Synthetic Network Adapter.* Synthetic network adapters represent computer network adapters. You can only use synthetic network adapters with supported virtual-machine guest operating systems.
- *COM port.* Com port enables connections to a simulated serial port on the virtual machine.
- *Diskette Drive.* You can map a .vhd floppy disk image to a virtual diskette drive.

You can add the following hardware to a virtual machine by editing the virtual machine's properties, and clicking on **Add Hardware**:

- *SCSI Controller.* You can add up to four virtual SCSI devices. Each controller supports up to 64 disks.
- *Network Adapter.* A single virtual machine can have a maximum of eight synthetic network adapters.
- *Legacy network adapter.* Legacy network adapters allow network adapters to be used with operating systems that do not support Integration Services. You also can use legacy network adapters to allow network deployment of operating-system images. A single virtual machine can have up to four legacy network adapters.
- *Fibre Channel Adapter.* Allows a virtual machine to connect directly to a Fibre Channel SAN. This requires that the Hyper-V host have a Fibre Channel HBA that also has a Windows Server 2012 driver that supports Virtual Fibre Channel.
- *RemoteFX 3D Adapter.* The RemoteFX 3D Adapter allows virtual machines to take advantage of DirectX and graphics processing power on the host Windows Server 2012 server to display high performance graphics.

## Management

You can use Management settings to configure how the virtual machine behaves on the Hyper-V host. You can configure the following virtual-machine management settings:

- *Name.* You can use this setting to configure the virtual machine's name on the Hyper-V host. This does not alter the virtual machine's hostname.
- *Integration Services.* You can use this setting to configure which virtual-machine integration settings are enabled.
- *Snapshot File Location.* You can use this setting to specify a location for storing virtual-machine snapshots.
- *Smart Paging File Location.* The location used when smart paging is required to start the virtual machine.
- *Automatic Start Action.* You can use this setting to handle how the virtual machine responds when the Hyper-V host is powered on.
- *Automatic Stop Action.* You can use this setting to handle how the virtual machine responds when the Hyper-V host is gracefully shut down.



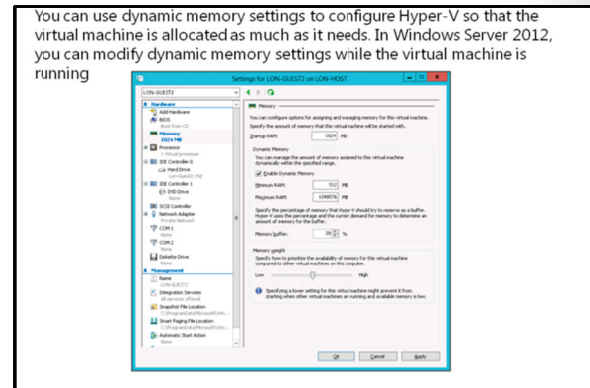
## How Dynamic Memory Works in Hyper-V

In the first release of Hyper-V with Windows Server 2008, virtual machines only could be assigned a static amount of memory. Unless you took special precautions to measure the precise amount of memory that a virtual machine requires, you were likely to under-allocate or over-allocate memory.

Windows Server 2008 R2 SP1 introduced dynamic memory, which you can use to allocate a minimum amount of memory to a virtual machine. You then can allow the virtual machine to use request additional memory, as necessary.

Rather than attempting to guess how much memory a virtual machine requires, dynamic memory allows you to configure Hyper-V so that the virtual machine is allocated as much as it needs. You can choose a minimum value, which will always be allocated to the virtual machine. You can choose a maximum value, which the virtual machine will not exceed, even if more memory is requested. Virtual machines must support Hyper-V Integration Services to be able to use dynamic memory.

With Windows Server 2012, you can modify dynamic memory settings while the virtual machine is running. This was not possible in Windows Server 2008 R2 SP1.



### Smart Paging

Another new memory feature available in Windows Server 2012 is smart paging. Smart paging provides a solution to the problem of minimum memory allocation, as it relates to virtual machine startup. Virtual machines can require more memory during startup than they would require during normal operation. In the past, it was necessary to allocate the minimum required for startup to ensure that startup occurred even though that value could be more than the virtual machine needed during normal operation. Smart paging uses disk paging for additional temporary memory when additional memory beyond the minimum allocated is required to restart a virtual machine. This provides you with the ability to allocate a minimum amount of memory based on the amount needed when the virtual machine is operating normally, rather than the amount required during startup. One drawback of smart paging is a decrease in performance during virtual-machine restarts.

You can configure virtual machine memory by using the **Set-VMemory** Windows PowerShell cmdlet.



**Additional Reading:** Hyper-V Dynamic Memory  
<http://technet.microsoft.com/en-us/library/hh831766.aspx>

## Demonstration: Creating a Virtual Machine

In this demonstration, you will see how to create a virtual machine by using the traditional method of using the Hyper-V Manager console. You also will see how you can automate the process by using Windows PowerShell.

## Demonstration Steps

1. Use the Hyper-V Manager console to create a virtual machine with the following properties:
  - o Name: **LON-GUEST1**
  - o Location: **E:\Program Files\Microsoft Learning\Base\LON-GUEST1\**
  - o Memory: **1024 MB**
  - o Use Dynamic Memory: **Yes**
  - o Networking: **Private Network**
  - o Connect Virtual Hard Disk: **E:\Program Files\Microsoft Learning\Base\LON-GUEST1\lon-guest1.vhd**
2. Open Windows PowerShell, import the Hyper-V module, and then run the following command:
 

```
New-VM -Name LON-GUEST2 -MemoryStartupBytes 1024MB -VHDPATH "E:\Program Files\Microsoft Learning\Base\LON-GUEST2\LON-GUEST2.vhd" -SwitchName "Private Network"
```
3. Use the Hyper-V Manager console and edit the settings of LON-GUEST2. Configure the following:
  - o Automatic Start Action: **Nothing**
  - o Automatic Stop Action: **Shut down the guest operating system**

## Importing, Exporting, and Moving Virtual Machines in Hyper-V

You can use the import and export functionalities in Hyper-V to transfer virtual machines between Hyper-V hosts and create point-in-time backups of virtual machines.

### Importing Virtual Machines

The virtual machine import feature in Windows Server 2012 provides more detailed information than previous Hyper-V versions featured. You can use this information to identify configuration problems such as missing hard disks or virtual switches. This was more difficult to determine in Windows Server 2008 and Windows Server 2008 R2.

In Hyper-V 3.0, you can import virtual machines from copies of virtual machine configuration, snapshot, and virtual hard-disk files rather than specially exported virtual machines. This is beneficial in recovery situations where the operating-system volume might have failed but the virtual machine files remain intact.

To import a virtual machine by using Hyper-V Manager, perform the following general steps:

1. In the Actions pane of the Hyper-V Manager console, click **Import Virtual Machine**.
2. On the **Before You Begin** page of the Import Virtual Machine wizard, click **Next**.
3. On the **Locate Folder** page, specify the folder that hosts the virtual machine files, and then click **Next**.

When importing virtual machines in Hyper-V:

- You can get access to a more detailed diagnostic information
- You can import copied virtual machine files

When exporting virtual machines, there are two options:

- Export snapshot for point in time export
- Export virtual machine to export all snapshots

When moving virtual machines:

- You can relocate virtual-machine files while the virtual machine is online
- You can perform a live migration

4. On the **Select Virtual Machine** page, select the virtual machine that you want to import, and then click **Next**.
5. On the **Choose Import Type** page, choose from the following options:
  - Register the virtual machine in-place (use the existing unique ID)
  - Restore the virtual machine (use the existing unique ID)
  - Copy the virtual machine (create a new unique ID)

You can import virtual machines by using the **Import-VM** cmdlet.

## Exporting Virtual Machines

When performing an export, you can select one of the following options:

- Export a snapshot. You can do this by right-clicking the snapshot in the Hyper-V manager console, and then selecting **Export**. This enables you to create an exported virtual machine as it existed at the point that the snapshot was created. The exported virtual machine will have no snapshots.
- Export Virtual Machine with Snapshot. You can do this by selecting the virtual machine, and then clicking **Export**. This exports the virtual machine and all snapshots associated with the virtual machine.

Exporting a virtual machine does not affect the existing virtual machine. However, you cannot import the virtual machine again unless you use the **Copy the Virtual Machine** option, which creates a new unique ID.

You can export virtual machines by using the **Export-VM** cmdlet.

## Moving Virtual Machines

You can perform two types of moves by using the Hyper-V move function: a live migration and a move of the actual virtual machine.

You can move virtual machines from one Hyper-V 3.0 server to another if you have enabled live migrations. Live migration of virtual machines occurs when you move a virtual machine from one host to another while keeping the virtual machine online and available to clients. For more information on migrating virtual machines, visit Module 9: Implementing Failover Clustering with Hyper-V.

You can use the move functionality to move some or all of the virtual-machine files to a different location. For example, if you want to move the virtual machines from one volume to an SMB share, while keeping the virtual machine hosted in the same location, you have the following options:

- *Move all the virtual machine's data to a single location.* This moves all configuration files, snapshots, and virtual hard-disk files to the destination location.
- *Move the virtual machine's data to different locations.* This moves the virtual machine's configuration files, snapshots, and virtual hard disks to separate locations.
- *Move the virtual machine's virtual hard disks.* This moves the hard disks to a separate location, while keeping the snapshot and configuration files in the same location.

You can move virtual machines in PowerShell by using the **Move-VM** cmdlet.

## Best Practices for Configuring Virtual Machines

When creating new virtual machines, keep the following best practices in mind:

- *Use dynamic memory.* The only time you should avoid dynamic memory is if you have an application that does not support it. For example, some Microsoft Exchange 2010 roles keep requesting memory, if it is available. In such cases, set static memory limits. You should monitor memory utilization, and set the minimum memory to the server's minimum memory utilization. Also, set a maximum amount of memory. The default maximum is more memory than most host servers have available.
- *Avoid differencing disks.* Differencing disks reduce the amount of space required, but decrease performance as multiple virtual machines access the same parent virtual hard disk file.
- *Use multiple synthetic network adapters connected to different external virtual switches.* Configure virtual machines to use multiple virtual network adapters that are connected to host NICs, which in turn are connected to separate physical switches. This means that network connectivity is retained if a NIC fails or a switch fails.
- *Store virtual machine files on its own volume.* This minimizes the chance that one virtual machine's virtual hard disk growth affects the other virtual machines on the same server.

The best practices for configuring virtual machines are:

- Use dynamic memory unless an application does not support it
- Avoid using differencing disks
- Configure multiple synthetic network adapters
- Store each virtual machine's files on a separate volume

# Lab: Implementing Server Virtualization with Hyper-V

## Scenario

IT management at A. Datum is concerned about the low utilization for many of the physical servers deployed in the London data center. Also, A. Datum is exploring options for expanding into multiple branch offices, and deploying servers in public and private clouds. For this purpose, the company is exploring the use of virtual machines.

As one of the senior network administrators at A. Datum, you are responsible for implementing Hyper-V in the London data center. You will deploy the Hyper-V server role, configure virtual machine storage and networking, and deploy the virtual machines.

## Objectives

After performing this lab you will be able to:

- Install the Hyper-V Server role.
- Configure virtual networking.
- Configure a virtual machine.

## Lab Setup

Estimated time: **60 minutes**

Virtual Machine(s)	20417A-LON-HOST1 Or 20417A-LON-HOST2
User Name	Adatum\Administrator
Password	Pa\$\$w0rd

## Lab Setup Instructions

1. Restart the classroom computer and in **Windows Boot Manager**, select **20417A-LON-HOST1** or **20417A-LON-HOST2**. Your instructor will specify which host to log on to.
2. Log on to **LON-HOST1** or **LON-HOST2** server with the following credentials:
  - Account: **Adatum\Administrator**
  - Password: **Pa\$\$w0rd**

## Exercise 1: Install the Hyper-V Server Role

### Scenario

The first step in migrating to a virtualized environment is to install the Hyper-V server role on a new server.

The main tasks for this exercise are as follows:

1. Configure network settings on LON-HOST1 and LON-HOST2.
2. Install the Hyper-V server role.
3. Complete Hyper-V role installation and verify settings.

► **Task 1: Configure network settings on LON-HOST1 and LON-HOST2**

1. Restart the classroom computer, and in the **Windows Boot Manager**, select either **20417A-LON-HOST1** or **20417A-LON-HOST2**.

If you start LON-HOST1, your partner must start LON-HOST2.

2. Log on to the server by using the following credentials:
  - Account: **Adatum\Administrator**
  - Password: **Pa\$\$w0rd**
3. In Server Manager, click **Local Server**, and then configure the following network settings:
  - LON-HOST1: **172.16.0.31**
  - LON-HOST2: **172.16.0.32**
  - Subnet mask: **255.255.0.0**
  - Default gateway: **172.16.0.1**
  - Preferred DNS server: 172.16.0.10

► **Task 2: Install the Hyper-V server role**

1. In Server Manager, use the **Add Roles and Features Wizard** to add the Hyper-V role to LON-HOST1 or LON-HOST2 with the following options:
  - Do not create a virtual switch
  - Use the Default stores locations
  - Allow the server to restart automatically if required.
2. After a few minutes, the server will automatically restart. Ensure that you restart the machine by using the **Boot** menu, and then selecting **20417-LON-HOST1** or **20417-LON-HOST2**. The computer will restart several times.

► **Task 3: Complete Hyper-V role installation and verify settings**

1. Log on to **LON-HOST1** or **LON-HOST2** by using **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. When the installation of the Hyper-V tools completes, click **Close**.
3. Open the Hyper-V Manager console, and then click **LON-HOST1** or **LON-HOST2**.
4. Open the Hyper-V settings, and then configure or verify the following settings:
  - Keyboard: **Use on the virtual machine**
  - Virtual Hard Disks: **C:\Users\Public\Documents\Hyper-V\Virtual Hard Disks**
5. Question: What additional features are required to support the Hyper-V role?

**Results:** After completing this exercise, you will have deployed the Hyper-V role to a physical server.

## Exercise 2: Configuring Virtual Networking

### Scenario

After installing the Hyper-V server role on the new server, you need to configure the virtual networks you are your manager specifies. You need to create a network that connects to the physical network and a private network that you can use only for communication between virtual machines. The private network is used when virtual machines are configured for high availability. You also need to configure a specific range of media access control (MAC) addresses for the virtual machines.

The main tasks for this exercise are as follows:

1. Configure the external network.
2. Create a private network.
3. Create an internal network.

#### ► Task 1: Configure the external network

1. In Hyper-V Manager, use the **Virtual Switch Manager** to create a new **External** virtual network switch with the following properties:
  - Name: **Corporate Network**
2. External Network: Mapped to the host computer's physical network adapter. Will vary depending on host computer.

#### ► Task 2: Create a private network

- In Hyper-V Manager, use the **Virtual Switch Manager** to create a new virtual switch with the following properties.
  - Name: **Private Network**
  - Connection type: **Private network**

#### ► Task 3: Create an internal network

- In Hyper-V Manager, use the **Virtual Switch Manager** to create a new virtual switch with the following properties:
  - Name: **Internal Network**
  - Connection type: **Internal network**

**Results:** After completing this exercise, you will have configured virtual switch options on a physically deployed Windows Server 2012 server that is running the Hyper-V role.

## Exercise 3: Creating and Configuring a Virtual Machine

### Scenario

You have been asked to deploy two virtual machines and to import a third virtual machine. You have copied a sysprepped VHD file that hosts a Windows Server 2012 Hyper-V host.

To minimize disk space use at the cost of performance, you are going to create two differencing files based on the sysprepped VHD. You use these differencing files as the hard-disk files for the new virtual machines.

You also will import a specially prepared virtual machine.

The main tasks for this exercise are as follows:

1. Configure virtual machine storage.
2. Create virtual machines.
3. Configure VLANs and network bandwidth settings.
4. Import a virtual machine.
5. Configure virtual machine dynamic memory.
6. Configure and test virtual machine snapshots.

### ► Task 1: Configure virtual machine storage

1. Use Windows Explorer to create the following folders on the physical host drive:
  - **E:\Program Files\Microsoft Learning\Base \LON-GUEST1**
  - **E:\Program Files\Microsoft Learning\Base \LON-GUEST2**



**Note:** The drive letter may depend upon the number of drives on the physical host machine)

2. In the Hyper-V Manager console, create a virtual hard disk with the following properties:
  - Disk Format: **VHD**
  - Disk Type: **Differencing**
  - Name: **LON-GUEST1.vhd**
  - Location: **E:\Program Files\Microsoft Learning\Base\LON-GUEST1\**
  - Parent Location: **E:\Program Files\Microsoft Learning\Base\Base12A-WS2012-RC.vhd**
3. Open Windows PowerShell, import the Hyper-V module, and then run the following command:

```
New-VHD "E:\Program Files\Microsoft Learning\Base\LON-GUEST2\LON-GUEST2.vhd"
-ParentPath "E:\Program Files\Microsoft Learning\Base\Base12A-WS2012-RC.vhd"
```

4. Inspect disk **E:\Program Files\Microsoft Learning\Base\LON-GUEST2\LON-GUEST2.vhd**.
5. Verify that **LON-GUEST2.vhd** is configured as a differencing virtual hard disk with **E:\Program Files\Microsoft Learning\Base\Base12A-WS2012-RC.vhd** as a parent.

### ► Task 2: Create virtual machines

1. Use the Hyper-V Manager console to create a virtual machine with the following properties:
  - Name: **LON-GUEST1**
  - Location: **E:\Program Files\Microsoft Learning\Base\LON-GUEST1\**
  - Memory: **1024 MB**
  - Use Dynamic Memory: **Yes**
  - Networking: **Private Network**
  - Connect Virtual Hard Disk: **E:\Program Files\Microsoft Learning\Base\LON-GUEST1\lon-guest1.vhd**



2. Open Windows PowerShell, import the Hyper-V module, and then run the following command:

```
New-VM -Name LON-GUEST2 -MemoryStartupBytes 1024MB -VHDPATH "E:\Program
Files\Microsoft Learning\Base\LON-GUEST2\LON-GUEST2.vhd" -SwitchName "Private
Network"
```

3. Use the Hyper-V Manager console, and then edit the settings of LON-GUEST2. Configure the following:
  - Automatic Start Action: **Nothing**
  - Automatic Stop Action: **Shut down the guest operating system**

### ► Task 3: Configure VLANs and network bandwidth settings

1. In Hyper-V Manager, use **Virtual Switch Manager** to configure the **Internal Network** virtual switch to use a **VLAN ID** of **4**.
2. Configure the following properties for the network adapter on **LON-GUEST2**:
  - Virtual Switch: **Internal Network**
  - VLAN ID: **4**
  - Enable DHCP guard
  - Enable router advertisement guard

**Question:** What kind of switch would you create if you added a new physical network adapter to the Hyper-V host and wanted to keep this separate from the existing networks you create during this exercise?

### ► Task 4: Import a virtual machine

1. Perform the following task:
  - If you are using LON-HOST1, use the Hyper-V Manager console to import the virtual machine **E:\Program Files\Microsoft Learning\20417\Drives\20417A-LON-DC1-B**.
  - If you are using LON-HOST2, use the Hyper-V Manager console to import the virtual machine **E:\Program Files\Microsoft Learning\20417\Drives\20417A-LON-SVR1-B**.
2. When importing, select the **Register the virtual machine in-place** option.

### ► Task 5: Configure virtual machine dynamic memory.

- Edit the properties of virtual machine **LON-GUEST2**, and then configure the following settings:
  - Startup RAM: **1024 MB**
  - Enable Dynamic Memory
  - Minimum RAM: **512 MB**
  - Maximum RAM: **2048 MB**

### ► Task 6: Configure and test virtual machine snapshots

1. If you are using **LON-HOST1**, start and then log on to **20417A-LON-DC1-B**. If you are using **LON-HOST2**, log on to virtual machine **20417A-LON-SVR1-B**.
2. On the desktop of the virtual machine, create the following folders:
  - Sydney
  - Melbourne
  - Brisbane

3. Create a snapshot of the virtual machine named **Before Change**.
4. Delete the following folders on the desktop:
  - Sydney
  - Brisbane
5. Revert the virtual machine.
6. Verify that the following folders are present on the desktop:
  - Sydney
  - Melbourne
  - Brisbane
7. Delete all three folders from the desktop.

**Question:** What state must the virtual machine be in to configure dynamic memory when using Windows Server 2008 R2 as a host? How is this different to Windows Server 2012 as a host?

**Results:** After completing this exercise, you will have deployed two separate virtual machines by using a sysprepped virtual hard-disk file to act as a parent disk for two differencing disks. You also will have imported a specially prepared virtual machine.

► **To prepare for the next module**

- When you are finished the lab, leave the virtual machines running, as they are needed for the lab in Module 9.

## Module Review and Takeaways

### Review Questions

**Question:** In which situations, should you use a fixed-memory allocation rather than dynamic memory?

**Question:** In which situations must you use virtual hard disks in VHDX format as opposed to virtual hard disks in VHD format?

**Question:** You want to deploy a Windows Server 2012 Hyper-V virtual machine's virtual hard disk on a file share. What operating system must the file server be running to support this configuration?

### Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Cannot deploy Hyper-V on x64 processor	
Virtual machine does not use dynamic memory	

### Real-world Issues and Scenarios

You have 10 servers that run Windows Server 2008 with Hyper-V. You are planning to upgrade these servers to Windows Server 2012 and want them to continue to run the Hyper-V role. What technology should you verify that the processor supports before performing the upgrade?

### Tools

Tool	Used for	Where to find it?
The Sysinternals disk2vhd tool	Convert physical hard disks to VHD format	Microsoft TechNet website <a href="http://technet.microsoft.com/en-us/sysinternals/bb842062">http://technet.microsoft.com/en-us/sysinternals/bb842062</a>
Virtual Machine Manager 2012	<ul style="list-style-type: none"> <li>Manage virtual machines across multiple Hyper-V servers</li> <li>Perform online physical to virtual conversions</li> </ul>	Microsoft TechNet website <a href="http://technet.microsoft.com/en-us/library/gg610610.aspx">http://technet.microsoft.com/en-us/library/gg610610.aspx</a>

**MCT USE ONLY. STUDENT USE PROHIBITED**

# Module 9

## Implementing Failover Clustering with Hyper-V

### Contents:

Module Overview	9-1
<b>Lesson 1:</b> Overview of the Integration of Hyper-V with Failover Clustering	9-2
<b>Lesson 2:</b> Implementing Hyper-V Virtual Machines on Failover Clusters	9-7
<b>Lesson 3:</b> Implementing Hyper-V Virtual Machine Movement	9-14
<b>Lesson 4:</b> Managing Hyper-V Virtual Environments by Using System Center Virtual Machine Manager	9-19
<b>Lab:</b> Implementing Failover Clustering with Hyper-V	9-29
Module Review and Takeaways	9-33

## Module Overview

One benefit of implementing server virtualization is the opportunity to provide high availability, both for applications or services that have built-in high availability functionality, and for applications or services that do not provide high availability in any other way. With the Windows Server® 2012 Hyper-V® technology, failover clustering, and Microsoft® System Center 2012 Virtual Machine Manager (VMM), you can configure high availability by using several different options.

In this module, you will learn about how to implement failover clustering in a Hyper-V scenario to achieve high availability for virtual environment. You will also learn about basic features of virtual machine.

### Objectives

After completing this module, you will be able to:

- Describe how Hyper-V integrates with failover clustering.
- Implement Hyper-V virtual machines on failover clusters.
- Implement Hyper-V virtual machine movement.
- Manage a Hyper-V virtual environment by using VMM.

## Lesson 1

# Overview of the Integration of Hyper-V with Failover Clustering

Failover clustering is a Windows Server 2012 feature that enables you to make applications or services highly available. To make virtual machines highly available in Hyper-V environment, you must implement failover clustering on the Hyper-V host computers.

This lesson summarizes the high availability options for Hyper-V based virtual machines, and then focuses on how failover clustering works, and how to design and implement failover clustering for Hyper-V.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe options for making virtual machines highly available.
- Describe how failover clustering works with Hyper-V nodes.
- Describe new features of failover clustering for Hyper-V.
- Describe best practices for implementing high availability in a virtual environment.

## Options for Making Virtual Machines Highly Available

Most organizations have some applications that are business critical and must be highly available. To make an application highly available, you must deploy it in an environment that provides redundancy for all components that the application requires. For virtual machines to be highly available, you can choose between several options. You can implement virtual machine as a clustered role (host clustering), you can implement clustering inside virtual machines (guest clustering) or you can use Network Load Balancing (NLB) inside virtual machines.

High availability options	Description
Host clustering	<ul style="list-style-type: none"> <li>• Virtual machines are highly available</li> <li>• Does not require virtual machine operating system or application to be cluster aware</li> </ul>
Guest clustering	<ul style="list-style-type: none"> <li>• Virtual machines are failover cluster nodes</li> <li>• Virtual machine applications must be cluster aware</li> <li>• Requires iSCSI or virtual fiber channel interface for shared storage connections</li> </ul>
NLB	<ul style="list-style-type: none"> <li>• Virtual machines are NLB cluster nodes</li> <li>• Use for web-based applications</li> </ul>

### Host Clustering

Host clustering enables you to configure a failover cluster by using the Hyper-V host servers. When you configure host clustering for Hyper-V, you configure the virtual machine as a highly available resource. Failover protection is implemented at the host server level. This means that the guest operating system and applications that are running within the virtual machine do not have to be cluster-aware. However, the virtual machine is still highly available. Some examples of non-cluster-aware applications are a File Server or Print Server, or perhaps a proprietary network-based application, such as an accounting application. Should the host node that controls the virtual machine unexpectedly become unavailable, the secondary host node takes control and restarts the virtual machine as quickly as possible. You can also move the virtual machine from one node in the cluster to another in a controlled manner. For example, you could move the virtual machine from one node to another while patching the Host operating system, and the applications or services that are running in the virtual machine, do not have to be compatible with failover clustering nor are they aware that virtual machine is clustered. Because the failover is at the virtual machine level, there are no dependencies on software that is installed inside the virtual machine.

## Guest Clustering

Guest failover clustering is configured very similarly to physical server failover clustering, except that the cluster nodes must include multiple virtual machines. In this scenario, you create two or more virtual machines, and enable failover clustering within the guest operating system. The application or service is then enabled for high availability between the virtual machines by using failover clustering in each virtual machine. Because failover clustering is implemented within each virtual machine node's guest operating system, you can locate the virtual machines on a single host. This can be a quick and cost-effective configuration in a test or staging environment.

For production environments however, you can more robustly protect the application or service if you deploy the virtual machines on separate failover clustering enabled Hyper-V host computers. With failover clustering implemented both at the host and virtual machine levels, the resource can be restarted regardless of whether the node that fails is a virtual machine or a host. This configuration is also known as a "Guest Cluster Across Hosts." It is considered an optimal high availability configuration for virtual machines running mission-critical applications in a production environment.

You should consider several factors when you implement guest clustering:

- The application or service must be failover cluster-aware. This includes any of the Windows Server 2012 services that are cluster-aware, and any applications, such as clustered Microsoft SQL Server and Microsoft Exchange Server.
- Hyper-V virtual machines can use fiber channel-based connections to shared storage (this is specific only to Microsoft Hyper-V Server 2012), or you can implement iSCSI connections from the virtual machines to the shared storage.

You should deploy multiple network adapters on the host computers and the virtual machines. Ideally, you should dedicate a network connection to the iSCSI connection (if you are using this method to connect to storage), to the private network between the hosts, and to the network connection that the client computers use.

## Network Load Balancing

NLB works with virtual machines in the same manner that it works with physical hosts. It distributes IP traffic to multiple instances of a TCP/IP service, such as a web server that is running on a host within the NLB cluster. NLB transparently distributes client requests among the hosts, and it enables the clients to access the cluster by using a virtual Host Name or a virtual IP addresses. From the client computer's point of view, the cluster seems to be a single server that answers these client requests. As enterprise traffic increases, you can add another server into the cluster.

Therefore, NLB is an appropriate solution for resources that do not have to accommodate exclusive read or write requests. Examples of NLB-appropriate applications would be web-based front ends to database applications or Exchange Server Client Access Servers.

When you configure an NLB cluster, you must install and configure the application on all virtual machines. After you configure the application, you install the network load balancing feature in Windows Server 2012 within each virtual machine's guest operating system (**not** on the Hyper-V hosts), and then configure an NLB cluster for the application. Earlier versions of Windows Server also support NLB, so that the Guest operating system is not limited to only Windows Server 2012. Similar to a "Guest Cluster Across Hosts", the NLB resource typically benefits from overall increased I/O performance when the virtual machine nodes are located on different Hyper-V hosts.



**Note:** As with earlier versions of Windows Server, you should not implement NLB and failover clustering within the same operating system because the two technologies conflict with one another.

## How Does a Failover Cluster Work with Hyper-V Nodes?

When you implement failover clustering and configure virtual machines as highly available resources, the failover cluster treats the virtual machines like any other application or service. Namely, if there is host failure, failover clustering will act to restore access to the virtual machine as quickly as possible on another host in the cluster. Only one node at a time runs the virtual machine. However, you can also move the virtual machine to any other node in the same cluster.

The failover process transfers the responsibility of providing access to resources in a cluster from one node to another. Failover can occur when an administrator intentionally moves resources to another node for maintenance or other reasons, or when unplanned downtime of one node occurs because of hardware failure or other reasons.

The failover process consists of the following steps:

1. The node where the virtual machine is running owns the clustered instance of the virtual machine, controls access to the shared bus or iSCSI connection to the cluster storage, and has ownership of any disks, or Logical Unit Numbers (LUNs), assigned to the virtual machine. All the nodes in the cluster use a private network to send regular signals, known as heartbeat signals, to one another. The heartbeat signals that a node is functioning and communicating on the network. The default heartbeat configuration specifies that each node send a heartbeat over TCP/UDP port 3343 each second (or 1000 milliseconds).
2. Failover starts when the node hosting the virtual machine does not send regular heartbeat signals over the network to the other nodes. By default, this is five consecutively missed heartbeats (or 5000 milliseconds elapses). Failover may occur because of a node failure or network failure.
3. When heartbeat signals stop arriving from the failed node, one of the other nodes in the cluster begins taking over the resources that the virtual machines use. You define the node(s) that could take over by configuring the **Preferred and Possible Owners** properties. The Preferred Owner specifies the hierarchy of ownership if there is more than one possible failover node for a resource. By default all nodes are members of Possible Owners. Therefore, removing a node as a Possible Owner absolutely excludes it from taking over the resource in a failure situation. Suppose that a failover cluster is implemented by using four nodes. However, only two nodes are configured as Possible Owners. In a failover event, the resource might still be taken over by the third node if neither of the Preferred Owners is online. Although the fourth node is not configured as a Preferred Owner, as long as it remains a member of Possible Owners, the failover cluster uses it to restore access to the resource if necessary. Resources are brought online in order of dependency. For example, if the virtual machine references an iSCSI LUN, access to the appropriate host bus adapters (HBAs), network(s) and LUNs will be stored in that order. Failover is complete when all the resources are online on the new node. For clients interacting with the resource, there is a short service interruption, which most users might not notice.
4. You can also configure the cluster service to fail back to the offline node after it again becomes active. When the cluster service fails back, it uses the same procedures that it performs during failover. This means that the cluster service takes all the resources associated with that instance offline, moves the instance, and then brings all the resources in the instance back online.



## What's New in Failover Clustering for Hyper-V in Windows Server 2012?

In Windows Server 2012, failover clustering is much improved with respect to Hyper-V clusters. Some of the most important improvements are:

- Failover clustering now supports up to 4,000 virtual machines, and the improved Failover Cluster Manager snap-in simplifies managing many virtual machines.
- Administrators can now perform multiselect actions to queue live migrations of multiple virtual machines, instead of doing it one by one, as in earlier versions.
- Administrators can also configure virtual machine priority attribute to control the order in which virtual machines are started. Priority is also used to ensure that lower-priority virtual machines automatically release resources if they are needed by higher priority virtual machines.
- The Cluster Shared Volume (CSV) feature, which simplifies the configuration and operation of virtual machines, is improved for more security and performance. It now supports scalable file-based server application storage, increased backup and restore and single consistent file namespace. Also, you can now protect CSV volumes by using BitLocker® Drive Encryption and configuring them to make storage visible to only a subset of nodes.
- Virtual machine application monitoring. You can now monitor services running on clustered virtual machines. In clusters running Windows Server 2012, administrators can configure monitoring of services on clustered virtual machines that are also running Windows Server 2012. This functionality extends the high-level monitoring of virtual machines that is implemented in Windows Server 2008 R2 failover clusters.
- It is now possible to store virtual machines on SMB file shares in a file server cluster. This is a new way to provide high availability for virtual machines. Instead of making a cluster between Hyper-V nodes, you can now have Hyper-V nodes out of cluster but with virtual machine files on a highly available file share. To make this work, you should deploy a file server cluster in a scale-out file server mode. Scale-out file servers can also use Cluster Shared Volumes for storage.

- Support for up to 4000 virtual machines per cluster
- Multi select virtual machines for Live Migration
- Virtual machine priority attribute
- CSV improvements
- Virtual machine application monitoring
- Storing virtual machines on highly available SMB file share

## Best Practices for Implementing High Availability in a Virtual Environment

After you determine which applications are deployed on highly available failover clusters, you plan and deploy the failover clustering environment. Apply the following recommendations when you implement the failover cluster:

- Use Windows Server 2012 as the Hyper-V host. Windows Server 2012 provides enhancements such as Hyper-V 3.0, improved CSVs, virtual machine migrations, and other features that improve flexibility and performance when you implement host failover clustering.

- Use Windows Server 2012 as the Hyper-V host
- Plan for failover scenarios
- Plan the network design for failover clustering
- Plan the shared storage for failover clustering
- Use the recommended failover cluster quorum mode
- Deploy standardized Hyper-V hosts
- Develop standard management practices

- Plan for failover scenarios. When you design the hardware requirements for the Hyper-V hosts, make sure that you include the hardware capacity required when hosts fail. For example, if you deploy a six-node cluster, you must determine the number of host failures that you want to accommodate. If you decide that the cluster must sustain the failure of two nodes, then the four remaining nodes must have the capacity to run all the virtual machines in the cluster.
- Plan the network design for failover clustering. To optimize the failover cluster performance and failover, you should dedicate a fast network connection for internode communication. As with earlier versions, this network should be logically and physically separate from the network segment(s) used for clients to communicate with the cluster. You can also use this network connection to transfer virtual machine memory during a Live Migration. If you are using iSCSI for any virtual machines, dedicate a network connection to the iSCSI network connection also.
- Plan the shared storage for failover clustering. When you implement failover clustering for Hyper-V, the shared storage must be highly available. If the shared storage fails, the virtual machines will all fail, even if the physical nodes are functional. To ensure the storage availability, plan for redundant connections to the shared storage and redundant array of independent disks (RAID) redundancy on the storage device.
- Use the recommended failover cluster quorum mode. If you deploy a cluster with an even number of nodes, and shared storage is available to the cluster, the Failover Cluster Manager automatically selects Node and Disk Majority quorum mode. If you deploy a cluster with an odd number of nodes, the Failover Cluster Manager selects the Node Majority quorum mode. You should not modify the default configuration unless you understand the implications of doing this.
- Deploy standardized Hyper-V hosts. To simplify the deployment and management of the failover cluster and Hyper-V nodes, develop a standard server hardware and software platform for all nodes.
- Develop standard management practices. When you deploy multiple virtual machines in a failover cluster, you increase the risk that a single mistake may shut down a large part of the server deployment. For example, if an administrator accidentally configures the failover cluster incorrectly, and the cluster fails, all virtual machines in the cluster will be offline. To avoid this, develop and thoroughly test standardized instructions for all administrative tasks.

## Lesson 2

# Implementing Hyper-V Virtual Machines on Failover Clusters

Implementation of highly available virtual machines is somewhat different from implementing other roles in a failover cluster. Failover clustering in Windows Server 2012 provides many features for Hyper-V clustering in addition to tools for virtual machine high availability management. In this lesson, you will learn about how to implement highly available virtual machines.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe components of Hyper-V cluster.
- Describe prerequisites for Hyper-V failover cluster implementation.
- Implement Hyper-V virtual machines in a cluster.
- Configure CSVs.
- Implement highly available virtual machines on SMB 3.0 file shares
- Describe considerations for implementing Hyper-V virtual machines in a cluster.

### Components of Hyper-V Clusters

Hyper-V as a role has some specific requirements for cluster components. To form a Hyper-V cluster, you must have at least two physical nodes. Whereas other clustered roles (such as DHCP, file server, and so on) allow for nodes to be virtual machines, Hyper-V nodes must be composed of physical hosts. You cannot run Hyper-V as a virtual machine on a Hyper-V host.

In addition to having nodes, you must also have physical and virtual networks. Failover clustering requires a network for internal cluster communication, and also a network for clients.

You can also implement a storage network separately, depending of type of storage being used. Again, specific to Hyper-V role, you should also consider virtual networks for clustered virtual machines. It is very important to create the same virtual networks on all physical hosts that participate in one cluster. Failing to do this causes a virtual machine to lose network connectivity when moved from one host to another.

Storage is an important component of virtual machine clustering. You can use any type of storage that is supported by Windows Server 2012 failover clustering. We recommended that you configure storage as a CSV. This is discussed in a following topic.

Virtual machines are components of a Hyper-V cluster. In Failover Cluster Manager you can create new highly available virtual machines, or you can make existing virtual machines highly available. In both cases, the virtual machine storage location must be on shared storage that can be accessed to both nodes. You might not want to make all virtual machines highly available. In Failover Cluster Manager you can select which virtual machines are part of a cluster configuration.

#### Hyper-V cluster components:

- Cluster nodes (must be physical computers)
- Cluster networks
- Virtual networks
- Storage for virtual machines
- Virtual machines

## Prerequisites for Implementing Hyper-V Clusters

To deploy Hyper-V on a failover cluster, you must make sure that you meet the hardware, software, account, and network infrastructure requirements that the following sections detail.

### Hardware Requirements for Failover Clustering with Hyper-V

You must have the following hardware for a two-node failover cluster:

- **Server hardware:** Hyper-V requires an x64-based processor, hardware-assisted virtualization, and hardware-enforced Data Execution Prevention (DEP). As a best practice, the servers should have very similar hardware. If you are using Windows Server 2008, the processors on the servers must be the same version. If you are using Windows Server 2008 R2 or Windows Server 2012, the processors must use the same architecture.



**Note:** Microsoft supports a failover cluster solution only if all the hardware features are marked as “Certified for Windows Server.” Additionally, the complete configuration (servers, network, and storage) must pass all tests in the Validate This Configuration wizard, which is included in the Failover Cluster Manager snap-in.

- **Network adapters:** The network hardware, just as other features in the failover cluster solution, must be marked as “Certified for Windows Server”. To provide network redundancy, you can connect cluster nodes to multiple, distinct networks, or you can connect the nodes to one network that uses teamed network adapters, redundant switches, redundant routers, or similar hardware to remove single points of failure. We recommend that you configure multiple network adapters on the host computer that you configure as a cluster node. One network adapter should be connected to the private network that the inter-host communications uses.
- **Storage adapters:** If you use Serial Attached SCSI (SAS) or fiber channel, the mass-storage device controllers in all clustered servers should be identical and should use the same firmware version. If you are using iSCSI, each clustered server should have one or more network adapters that are dedicated to the cluster storage. The network adapters that you use to connect to the iSCSI storage target should be identical, and you should use Gigabit Ethernet or a faster network adapter.
- **Storage:** You must use shared storage that is compatible with Windows Server 2008 R2. If you deploy a failover cluster that uses a witness disk, the storage must contain at least two separate volumes (LUNs). One volume functions as the witness disk, and additional volumes contain the virtual machine files that are shared between the cluster nodes. Storage considerations and recommendations include the following:
  - Use basic disks, not dynamic disks. Format the disks with the NTFS file system.
  - Use either master boot record (MBR) or GUID partition table (GPT).
  - If you are using a storage area network (SAN), the miniport driver that the storage uses must work with the Microsoft Storport storage driver.
  - Consider using multipath input/output (I/O) software: If your SAN uses a highly available network design with redundant components, you can deploy failover clusters with multiple host bus adapters by using multipath I/O software. This provides the highest level of redundancy and availability. For Windows Server 2008 R2 and 2012, your multipath solution must be based on Microsoft Multipath I/O (MPIO).

## Software Requirements for Using Hyper-V and Failover Clustering

The following are the software requirements for using Hyper-V and failover clustering:

- All the servers in a failover cluster must run the x64-based version of Windows Server 2012 Enterprise or Datacenter Edition. The nodes in a single failover cluster cannot run different versions.
- All the servers should have the same software updates and service packs.
- All servers must be either a full installation or a Server Core installation. You cannot mix the full installation and Server Core installation.

## Network Infrastructure Requirements

The following network infrastructure is required for a failover cluster and an administrative account with the following domain permissions:

- Network settings and IP addresses. Use identical communication settings on all network adapters, including the speed, duplex mode, flow control, and media type settings. Ensure that all network hardware supports the same settings.
- If you use private networks that are not routed to your whole network infrastructure for communication between cluster nodes, ensure that each of these private networks uses a unique subnet.
- DNS. The servers in the cluster must use Domain Name System (DNS) for name resolution. You should use the DNS dynamic update protocol.
- Domain role. All servers in the cluster must be in the same Active Directory® domain. As a best practice, all clustered servers should have the same domain role (either member server or domain controller). The recommended role is member server.
- Account for administering the cluster. When you first create a cluster or add servers to it, you must be logged on to the domain with an account that has administrator rights and permissions on all the cluster's servers. Additionally, if the account is not a Domain Admins account, the account must have the Create Computer Objects permission in the domain.

## Implementing Hyper-V Virtual Machines on Failover Cluster

To implement failover clustering for Hyper-V, you must complete the following high-level steps:

1. Install and configure the required versions of Windows Server 2012. After you complete the installation, configure the network settings, join the computers to an Active Directory domain, and configure the connection to the shared storage.
2. Configure the shared storage. You must use Disk Manager to create disk partitions on the shared storage.
3. Install the Hyper-V and failover clustering features on the host servers. You can use Server Manager in MMC or Windows PowerShell® for this.

1. Install and configure Windows Server 2012
2. Configure shared storage
3. Install the Hyper-V and failover clustering features
4. Validate the cluster configuration
5. Create the cluster
6. Create a virtual machine on one of the cluster nodes
7. Make the virtual machine highly available

4. Validate the cluster configuration. Validate This Cluster wizard checks all the prerequisite components that are required to create a cluster, and provides warnings or errors if any components do not meet the cluster requirements. Before you continue, resolve any issues that the Validate This Cluster Wizard identifies.
5. Create the cluster. When the components pass the Validate This Cluster wizard, you can create a cluster. When you configure the cluster, assign a cluster name and an IP address. A computer account for the cluster name is created in Active Directory domain and the IP address is registered in DNS.



**Note:** You can enable Clustered Shared Storage for the cluster only after you configure the cluster. If you want to use Cluster Shared Volumes (CSV), you should configure CSV before you move to the next step.

6. Create a virtual machine on one of the cluster nodes. When you create the virtual machine, ensure that all files associated with the virtual machine, including both the virtual hard disk and virtual machine configuration files, are stored on the shared storage. You can create and manage virtual machines in either Hyper-V Manager or Failover Cluster Manager. When you create a virtual machine by using Failover Cluster Manager, the virtual machine is automatically made highly available.
7. Make the virtual machine highly available. To make the virtual machine highly available, in the Failover Cluster Manager, select to make a new service or application highly available. Failover Cluster Manager then presents a list of services and applications that can be made highly available. When you select the option to make virtual machines highly available, you can select the virtual machine that you created on shared storage.



**Note:** When you make a virtual machine highly available, you see a list of all virtual machines hosted on all cluster nodes, including virtual machines that are not stored on the shared storage. If you make a virtual machine that is not located on shared storage highly available, you receive a warning, but Hyper-V adds the virtual machine to the services and applications list. However, when you try to migrate the virtual machine to a different host, the migration will fail.

8. Test virtual machine failover. After you make the virtual machine highly available, you can migrate the computer to another node in the cluster. If you are running Windows Server 2008 R2 or Windows Server 2012, you can select to perform a Quick Migration or a Live Migration.

## Configuring Clustered Shared Volumes

You do not have to configure and use CSV when you implement high availability for virtual machines in Hyper-V. You can cluster Hyper-V by using the regular approach. However, we recommend that you use CSV because of the following advantages:

- Reduced LUNs for the disks. You can use CSV to reduce the number of LUNs that your virtual machines require. When you configure a CSV, you can store multiple virtual machines on a single LUN and multiple host computers can access the same LUN concurrently.

### CSV benefits:

- Fewer LUNs required
- Better use of disk space
- Virtual machine files are in a single logical location
- No special hardware required
- Increased resiliency

### To implement CSV:

1. Create and format volumes on shared storage.
2. Add the disks to failover cluster storage.
3. Add the storage to the CSV.



- Better use of disk space. Instead of placing each .vhd file on a separate disk with empty space so that the .vhd file can expand, you can oversubscribe disk space by storing multiple .vhd files on the same LUN.
- Virtual machine files are in a single logical location. You can track the paths of .vhd files and other files that virtual machines use. Instead of using drive letters or Globally Unique Identifiers (GUIDs) to identify disks, you can specify the path names. When you implement CSV, all added storage appears in the \ClusterStorage folder. The \ClusterStorage folder is created on the cluster node's system folder, and you cannot move it. This means that all Hyper-V hosts that are members of the cluster must use the same drive letter as their system drive, or virtual machine failovers will fail.
- No specific hardware requirements. There are no specific hardware requirements to implement CSV. You can implement CSV on any supported disk configuration, and on either fiber channel or iSCSI SANs.
- Increased resiliency. CSV increases resiliency because the cluster can respond correctly even if connectivity between one node and the SAN is interrupted, or part of a network is down. The cluster reroutes the CSV traffic through an intact part of the SAN or network.

### Implementing CSV

You can configure CSV only when you create a failover cluster that hosts highly available virtual machines. After you create the failover cluster, you can enable CSV for the cluster, and then add storage to the CSV.

Before you can add storage to the CSV, the LUN must be available as shared storage to the cluster. When you create a failover cluster, all the shared disks configured in Server Manager are added to the cluster, and you can add them to a CSV. If you add more LUNs to the shared storage, you must first create volumes on the LUN, add the storage to the cluster, and then add the storage to the CSV.

As a best practice, you should configure CSV before you make any virtual machines highly available. However, you can convert from regular disk access to CSV after deployment. The following considerations apply:

- The LUN's drive letter or mount point is removed when you convert from regular disk access to CSV. This means that you must re-create all virtual machines that are stored on the shared storage. If you must keep the same virtual machine settings, consider exporting the virtual machines, switching to CSV, and then importing the virtual machines in Hyper-V.
- You cannot add shared storage to CSV if it is used. If you have a running virtual machine that is using a cluster disk, you must shut down the virtual machine, and then add the disk to CSV.

### Implementing Highly Available Virtual Machines on an SMB 3.0 File Share

In Windows Server 2012, it is possible to use one more technique to make virtual machines highly available. Instead of using host or guest clustering, virtual machine files can now be stored on a highly available SMB 3.0 file share. By using this approach, high availability is achieved not by clustering Hyper-V nodes, but by file servers that host virtual machine files on their file shares. With this new capability, Hyper-V can store all virtual machine files, including configuration, virtual hard disk (VHD) files, and snapshots, on highly available SMB file shares.

- In Windows Server 2012 you can store VM files on SMB 3.0 file share
- File servers should be running Windows Server 2012
- File server cluster should be running in scale-out mode
- Hyper-V Manager can be used to create or move virtual machine files to SMB file share

To implement this technology, the following requirements must be met:

- One or more computers running Windows Server 2012 with the Hyper-V role installed.
- One or more computers running Windows Server 2012 with the File and Storage Services role installed.
- A common Active Directory infrastructure. The servers running Active Directory Domain Services (AD DS) do not need to run Windows Server 2012.

Before you implement virtual machines on an SMB file share, you should set up a file server cluster. To do that, you should have at least two cluster nodes with File Services and Failover Clustering installed. In the failover clustering console, you should create a scale-out file server cluster. After you configure the cluster, you deploy the new SMB file share for applications. This share is used to store virtual machine files. When the share is created, you can use Hyper-V Manager console to deploy new virtual machines on the SMB file share, or you can migrate existing VMs to the SMB file share by using the storage migration method.

## Considerations for Implementing Hyper-V Clusters

By implementing host failover clustering, you can make virtual machines highly available. However, implementing host failover clustering also adds significant cost and complexity to a Hyper-V deployment. You must invest in additional server hardware to provide redundancy, and you should implement or have access to a shared storage infrastructure.

Use the following recommendations to ensure that the failover clustering strategy meets the organization's requirements:

- Identify the applications that require high availability
- Identify the application components that must be highly available
- Identify the application characteristics
- Identify the total capacity requirements
- Create the Hyper-V design
  1. Verify basic requirements
  2. Configure a dedicated network adapter for the private virtual network
  3. Use similar host hardware
  4. Verify network configuration
  5. Manage Live Migrations

- Identify the applications or services that require high availability. If you were to ask the people who use the organization's applications, most of them would probably say that they want all applications to be highly available. However, unless you have the option of making all virtual machines highly available, you must develop priorities for which applications will be made highly available.
- Identify the components that must be highly available to make the applications highly available. In some cases, the application might run on a single server, and making that server highly available is all that you have to do. Other applications may require that several servers, and other components, such as storage or the network, be highly available.
- Identify the application characteristics. You must understand several things about the application:
  - Is virtualizing the server that is running the application an option? Some applications are not supported or recommended in a virtual environment.
  - What options are available for making the application highly available? You can make some applications highly available through options other than host clustering. If other options are available, evaluate the benefits and disadvantages of each option.
  - What are the performance requirements for each application? Collect performance information on the servers currently running the applications to gain an understanding of the hardware requirements that are required when you virtualize the server.



- What capacity is required to make the Hyper-V virtual machines highly available? As soon as you identify all the applications that must be highly available by using host clustering, you can start to design the actual Hyper-V deployment. By identifying the performance requirements, and network and storage requirements, for applications, you can define the hardware that you have to implement all the applications in a highly available environment.

Live Migration is one of the most important aspects of Hyper-V clustering. When you implement Live Migration, consider the following:

- Verify basic requirements. The basic requirements for Live Migration are that all hosts must be part of a Windows Server 2008 R2 failover cluster, and host processors must be from the same manufacturer. All hosts in the cluster must have access to shared storage.
- Configure a dedicated network adapter for the private virtual network. When you implement failover clustering, you should configure a private network for the cluster heartbeat traffic. You use this network to transfer the virtual machine memory during a failover. To optimize this configuration, configure a network adapter for this network that has a capacity of one gigabits per second (Gbps) or higher.



**Note:** You must enable the Client for Microsoft Networks and File and Printer Sharing for Microsoft Networks components for the network adapter that you want to use for the private network.

- Use similar host hardware. All failover cluster nodes must use the same hardware for connecting to shared storage, and all cluster nodes must have processors from the same manufacturer. Whereas you can enable failover for virtual machines on a host with different processor versions by configuring processor compatibility settings, the failover experience and performance is more consistent if all servers have very similar hardware.
- Verify network configuration. All nodes in the failover cluster must connect through the same IP subnet so that the virtual machine can keep the same IP address after Live Migration. Also, the IP addresses assigned to the private network on all nodes must be on the same logical subnet, which means that multisite clusters must use a stretched virtual local area network (VLAN), which is a subnet that spans a wide area network (WAN) connection.
- Manage Live Migrations. Each node in the failover cluster can perform only one Live Migration at a time. If you try to start a second Live Migration before the first one finishes, the migration fails. If you start additional Live Migrations from Virtual Machine Manager (VMM), it queues the Live Migration, and retries it for 15 minutes. If the migration cannot be initiated in 15 minutes, the migration is canceled.

## Lesson 3

# Implementing Hyper-V Virtual Machine Movement

Moving virtual machines from one location to another is a fairly common procedure in the administration of Hyper-V environments. Most of the moving techniques in previous Windows Server versions required downtime. Windows Server 2012 introduces new technologies to enable seamless virtual machine movement. In this lesson, you will learn about virtual machine movement and migration options.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe migration options for virtual machines.
- Describe Storage Migration.
- Describe Live Migration.
- Describe and configure a Hyper-V replica.

### Virtual Machine Migration Options

There are several scenarios where you would want to migrate virtual machine from one location to another. For example, you might want to move a virtual machine virtual hard disk from one physical drive to another on the same host. Another example is moving a virtual machine from one node in a cluster to another, or just moving a computer from one host server to another host server without the hosts being members of a cluster. Compared with Windows Server 2008 R2, Windows Server 2012 provides significant enhancements in addition to simplified procedures for this process.

Available options for moving virtual machines are:

- Virtual machine and storage migration
- Quick Migration
- Live Migration
- Hyper-V replica
- Export/Import of a virtual machine

In Windows Server 2012, you can perform migration of virtual machines by using these methods:

- **Virtual machine and storage migration.** With this method, you move a powered on virtual machine from one location to another (or from one host to another) by using a wizard in Hyper-V Manager. Virtual machine and storage migration do not require failover clustering or any other high availability technology to work. Shared storage is not required when you move just the virtual machine.
- **Quick Migration.** This method is also available in Windows Server 2008. It requires failover clustering to be installed and configured. It.
- **Live Migration.** This improvement over Quick Migration is also available in Windows Server 2008 R2. It enables you to migrate a virtual machine from one host to another without downtime.
- **Hyper-V replica.** This new feature in Windows Server 2012 enables you to replicate a virtual machine to another host, instead of move the virtual machine, and to synchronize all virtual machine changes from the primary host to the host that holds the replica.
- **Exporting and importing virtual machine.** This is an established method of moving virtual machines without using a cluster. You export a virtual machine on one host, and then physically move exported files to another host by performing an import operation. This is a very time-consuming operation. It requires that a virtual machine is turned off during export and import. In Windows

Server 2012 this migration method is improved. You can import a virtual machine to a Hyper-V host without exporting it before import. Windows Server 2012 Hyper-V is now capable of configuring all the necessary settings during the import operation.

## How Does Virtual Machine and Storage Migration Work?

There are many cases in which an administrator might want to move the virtual machine files to another location. For example, if the disk where a virtual machine hard disk resides runs out of space, you must move the virtual machine to another drive or volume. Also, moving a virtual machine to another host is a very common procedure.

In earlier versions of Windows Server, such as Windows Server 2008 or Windows Server 2008 R2, moving a virtual machine resulted in downtime because it had to be turned off. If you moved a virtual machine between two hosts, then you also had to perform export and import operations for that specific virtual machine. Export operations can be time-consuming, depending on the size of the virtual machine hard disks.

In Windows Server 2012, Virtual Machine and Storage Migration enables you to move a virtual machine to another location on the same host or on another host computer without turning off the virtual machine.

Let's examine how storage migration actually works.

To copy a virtual hard disk, an administrator starts live storage migration by using the Hyper-C console or Windows PowerShell, and completes the wizard (or specifies parameters in Windows PowerShell). A new virtual hard disk is created on destination location and the copy process starts. During the copy process, the virtual machine is fully functional. However, all changes that occur during copying are written to both the source and destination location. Read operations are performed only from the source location. As soon as the disk copy process is complete, Hyper-V switches virtual machines to run on the destination virtual hard disk. Also, if the virtual machine is moved to another host, the computer configuration is copied and the virtual machine is associated with another host. If a failure were to occur on the destination side, there is always a fail back option to run back again on the source directory. After the virtual machine is successfully migrated and associated to a new location, the process deletes the source VHDs.

The time that is required to move a virtual machine depends on the source and destination location, the speed of hard disks or storage, and the size of the virtual hard disks. The moving process is speeded up if source and destination locations are on storage, and storage supports Offloaded Data Transfer (ODX).

When you move a virtual machine's vhd's to another location, a wizard presents three available options:

- Move all the virtual machine's data to a single location: You specify one single destination location, such as disk file, configuration, snapshot, and smart paging.
- Move the virtual machine's data to a different location: You specify individual locations for each virtual machine item.
- Move only the virtual machine's virtual hard disk: You move only the virtual hard disk file.

Storage Migration technology enables you to move a virtual machine and its storage to another location without downtime.

- During migration the virtual machine hard drive is copied from one location to another
- Changes are written to both source and destination drive
- You can move virtual machine storage to same host, another host, or server message block share
- Storage and virtual machine configuration can be in different locations

## How Live Migration Works?

Live Migration enables you to move running virtual machines from one failover cluster node to another node in the same cluster. With Live Migration, users who are connected to the virtual machine should experience almost no server outage.



**Note:** Whereas you can also do live migration of virtual machine by using Virtual Machine and Storage migration described in previous topic, you should be aware that live migration is based on a different technology (failover clustering). Unlike the storage migration scenario, Live Migration can be performed only if a virtual machine is highly available.

You can start a Live Migration through one of the following:

- The Failover Cluster Management console.
- The VMM Administrator console, if you use VMM to manage your physical hosts.
- A Windows Management Instrumentation (WMI) or Windows PowerShell script.



**Note:** Live Migration enables you to reduce the perceived outage of a virtual machine significantly during a planned failover. During a planned failover, you start the failover manually. Live Migration does not apply during an unplanned failover, such as when the node hosting the virtual machine fails.

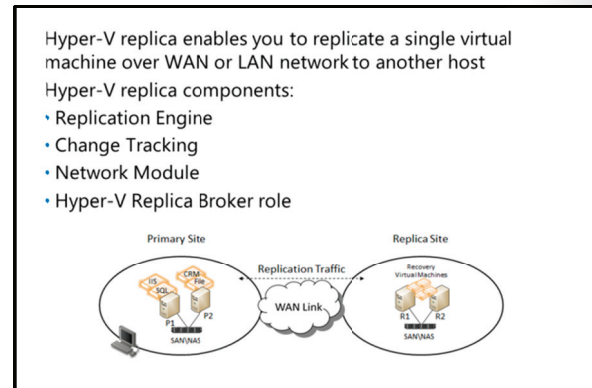
## Live Migration Process

The Live Migration process consists of four steps:

1. **Migration setup.** When the administrator starts the failover of the virtual machine, the source node creates a TCP connection with the target physical host. This connection is used to transfer the virtual machine configuration data to the target physical host. Live Migration creates a temporary virtual machine on the target physical host, and allocates memory to the destination virtual machine. The migration preparation also checks to determine whether a virtual machine can be migrated.
2. **Guest-memory transfer.** The guest memory is transferred iteratively to the target host while the virtual machine is still running on the source host. Hyper-V on the source physical host monitors the pages in the working set. As the system modifies memory pages, it tracks and marks them as being modified. During this phase of the migration, the migrating virtual machine continues to run. Hyper-V iterates the memory copy process several times, and every time that a smaller number of modified pages are copied to the destination physical computer. A final memory copy process copies the remaining modified memory pages to the destination physical host. Copying stops as soon as the number of dirty pages drops below a threshold or after 10 iterations are complete.
3. **State transfer.** To actually migrate the virtual machine to the target host, Hyper-V stops the source partition, transfers the state of the virtual machine (including the remaining dirty memory pages) to the target host, and then restores the virtual machine on the target host. The virtual machine has to be paused during the final state transfer.
4. **Clean up.** The cleanup stage finishes the migration by tearing down the virtual machine on the source host, terminating the worker threads, and signaling the completion of the migration.

## How Does Hyper-V Replica Work?

In some cases, you might want to have a spare copy of one virtual machine that you can run if the original virtual machine fails. By implementing high availability, you have one instance of a virtual machine. High availability does not prevent corruption of software running inside the VM. One way to address the issue of corruption is to copy the VM. You can also back up the virtual machine and its storage. Although this solution achieves the desired result it is resource intensive and time consuming.



To resolve this problem, and to enable administrators to have an up-to-date copy of a single virtual machine, Microsoft has implemented Hyper-V replica technology in Windows Server 2012. This technology enables virtual machines running at a primary site (can also be location or host) to be efficiently replicated to a secondary site (location or host) across a WAN or LAN link. Hyper-V replica enables you to have two instances of a single virtual machine residing on different hosts, one as the primary (live) copy and the other as a replica (offline) copy. These copies are synchronized, and you can failover at any time. In the event of a failure at a primary site (e.g. fire, natural disaster, power outage, server failure etc...), an administrator can use Hyper-V Replica to execute a failover of production workloads to replica servers at a secondary location within minutes, thus incurring minimal downtime.

The site configurations do not have to use the same server or storage hardware. Hyper-V Replica enables an administrator to restore virtualized workloads to a point in time depending on the Recovery History selections for the virtual machine.

Hyper-V replica technology consists of several components:

- **Replication Engine:** This component is the core of Hyper-V Replica. It manages the replication configuration details and handles initial replication, delta replication, failover, and test-failover operations. It also tracks virtual machine and storage mobility events and takes appropriate actions as needed (i.e. it pauses replication events until migration events complete and then resumes where they left off).
- **Change Tracking:** This component tracks changes that are happening on primary copy of virtual machine. It is designed to make the scenario work regardless of where the virtual machine VHD file(s) resides.
- **Network Module:** The Networking Module provides a secure and efficient way to transfer virtual machine replicas between primary host and replica host. Data compression is enabled by default. This communication is also secure as it relies on HTTPS and certification-based authentication.
- **Hyper-V Replica Broker role:** This is new role implemented in Windows Server 2012. It is configured in Failover Clustering, and it enables you to have Hyper-V replica functionality even when the virtual machine being replicated is highly available and can move from one cluster node to another. The Hyper-V Replica Broker redirects all virtual machine specific events to the appropriate node in the replica cluster. The Broker queries the cluster database to determine which node should handle which events. This ensures all events are redirected to the correct node in the cluster in the event that a Quick Migration, Live Migration, or Storage Migration process was executed.

## Configuring Hyper-V Replica

Before you implement Hyper-V replica technology, ensure that these prerequisites are met:

- The server hardware supports the Hyper-V role on Windows Server 2012.
- Sufficient storage exists on both the primary and replica servers to host the files that are used by replicated virtual machines.
- Network connectivity exists between the locations hosting the primary and replica servers. This can be a WAN or LAN link.
- Firewall rules are correctly configured to enable replication between the Primary and Replica sites (default traffic is going over TCP port 80 or 443).
- An X.509v3 certificate exists to support Mutual Authentication with certificates (if you want).

To configure Hyper-V replica you should:

1. Configure authentication options
2. Configure ports
3. Select replica servers
4. Select location for replica files
5. Enable replication on virtual machine

You do not have to install Hyper-V replica separately because it is not a Windows Server role or feature. Hyper-V Replica is implemented as part of the Hyper-V Role. It can be used on Hyper-V servers that are stand-alone or servers that are part of a Failover Cluster (in which case, you should configure Hyper-V Replica Broker). Unlike failover clustering, a Hyper-V role is not dependent on Active Directory Domain Services (AD DS). You can use it with Hyper-V servers that are stand-alone, or that are members of different Active Directory domains (except in case when servers are part of a failover cluster).

To enable Hyper-V replica technology, you should first configure Hyper-V server settings. In the Replication Configuration group of options, you should enable Hyper-V server as a replica server, and you should also select authentication and port options. You should also configure authorization options. You can choose to enable replication from any server that successfully authenticates (which is convenient in scenarios where all servers are part of same domain), or you can type fully qualified domain names (FQDNs) of servers that you accept as replica servers. Also, you must configure the location for replica files. These settings should be configured on each server that will serve as replica server.

After you configure options on server level, you should enable replication on a virtual machine. During this configuration, you must specify replica server name, as well as options for connection. You can select which virtual hard disk drives you replicate (in case when virtual machine has more than one VHD), and you can also configure Recovery History as well as initial replication method. After you have configured these options then you can start replication.



## Lesson 4

# Managing Hyper-V Virtual Environments by Using System Center Virtual Machine Manager

System Center Virtual Machine Manager 2012 is a part of the System Center 2012 family of products. It is a successor of Virtual Machine Manager 2008 R2. Its main purpose is to extend management functionality for Hyper-V hosts and virtual machines and to provide deployment and provisioning for virtual machines and services. In this lesson, you will learn the basics of VMM.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe System Center VMM.
- Describe Prerequisites for Installing VMM.
- Describe private cloud infrastructure components.
- Describe how VMM Manage Hosts and Host Clusters with VMM.
- Describe how to manage Virtual Machines with VMM.
- Describe Services and Service Templates.
- Describe Physical to Virtual and Virtual to Virtual Migrations.
- Describe considerations for deploying a highly available VMM Server.

### What Is VMM?

VMM is a management solution for a virtualized data center. VMM enables you to create and deploy virtual machines and services to private clouds by configuring and managing your virtualization host, networking, and storage resources.

VMM is a component of Microsoft System Center 2012 that discovers, captures, and aggregates knowledge of the virtualization infrastructure. VMM also manages policies, and processes, and best practices with automations by discovering, capturing and aggregating knowledge of virtualization infrastructure.

- VMM provides centralized administration and management of your virtual environment
- VMM is used to:
  - Manage Hyper-V hosts
  - Manage Vmware, XEN hosts
  - Manage and deploy virtual machines
  - Manage and deploy services
  - Perform physical-to-virtual (P2V) and virtual-to-virtual (V2V) conversions

VMM succeeds VMM 2008 R2 and is a key component in enabling private cloud infrastructure, which helps transition enterprise IT from an infrastructure-focused deployment model into a service-oriented, user-centric environment.

VMM architecture consists of several interrelated components. These components are:

- **VMM server.** The VMM server is the computer on which the VMM service runs. The VMM server processes commands and controls communications with the VMM database, the library server, and the virtual machine hosts. The VMM server is the hub of a VMM deployment through which all other VMM components interact and communicate. The VMM server also connects to a Microsoft SQL Server database (VMM database) that stores all VMM configuration information.

- **Database.** VMM uses a SQL Server database to store the information that you view in the VMM management console, such as managed virtual machines, virtual machine hosts, virtual machine libraries, jobs, and other virtual machine-related data.
- **Management console.** The management console is a program that you use to connect to a VMM management server, to view and manage physical and virtual resources, including virtual machine hosts, virtual machines, services, and library resources. Virtual Machine Manager library
- **Library.** A library is a catalog of resources (for example, virtual hard disks, templates, and profiles), that are used to deploy virtual machines and services. A library server also hosts shared folders that store file-based resources. The VMM management server is always the default library server, but you can add additional library servers later.
- **Command shell.** Windows PowerShell is the command-line interface in which you execute cmdlets that perform all available VMM functions. You can use these VMM-specific cmdlets to manage all the actions in a VMM environment.
- **Self-Service Portal.** The Self-Service Portal is a web site that users who are assigned to a self-service user role can use to deploy and manage their own virtual machines.

## Prerequisites for Installing VMM

Before you deploy VMM and its components, you should be certain that your system meets hardware and software requirements. While software requirements do not change based on the number of hosts that VMM manages, hardware prerequisites may vary. In addition, not all VMM components have the same hardware and software requirements. However, Windows Server 2008 R2 and Windows Server 2012 are the only supported operating systems for VMM 2012.

- **Software requirements for the VMM Server:**
  - Windows Server 2008 R2
  - SQL Server 2008 SP2 or SQL Server 2008 R2
  - Microsoft .NET Framework 3.5 SP1 or newer
  - Windows AIK
  - Windows PowerShell 2.0 (if the VMM Management Console will run on the same server)
  - WinRM 2.0
- **Hardware requirements:**
  - The number of hosts determine hardware requirements
  - CPU: Single core CPU 2 GHz
  - RAM: 4 – 8 GB
  - Disk space: 40 GB – 150 GB

### VMM Server

In addition to having Windows Server 2008 R2 or Windows Server 2012 installed, you have to ensure that the following software is installed on the server that will run the VMM server:

- Microsoft .NET Framework 3.5 Service Pack 1 (SP1) or later versions
- Windows Automated Installation Kit (AIK)
- Windows PowerShell 2.0 (if the VMM management console will run on the same server)
- Windows Remote Management 2.0 (this is installed by default in Windows Server 2008 R2, so you should just verify that the service is running)
- SQL Server 2008 SP2 (Standard or Enterprise) or SQL Server 2008 R2 SP1 Standard, Enterprise, or Datacenter. This is necessary only when you install the VMM management server and SQL Server on same computer.

Hardware requirements vary, depending on number of hosts, and have the following limits:

- CPU: Single core CPU 2 gigahertz (GHz), Dual core CPU 2.8 GHz
- Random access memory (RAM): 4 – 8 gigabytes (GB)



- Disk space: 40 GB – 150 GB (depending on whether a SQL Server database is installed on the same server. In addition, if the library is on the same server, then disk space will also depend on library content.)

### VMM Database

The VMM database stores all VMM configuration information, which you can access and modify by using the VMM management console. The VMM database requires SQL Server 2008 SP2 or later. Because of this, the base hardware requirements for the VMM database are equal to the minimum system requirements for installing SQL Server. Additionally, if you are managing more than 150 hosts, you should have at least 4 GB of RAM on the database server. Software requirements for the VMM Database are the same as for SQL Server.

### VMM Library

The VMM library is the server that hosts resources for building virtual machines, services and business unit clouds. In smaller environments, you usually install the VMM library on the VMM Management Server. If this is the case, the hardware and software requirements are the same as for the VMM Management Server. In larger and more complex environments, we recommend that you have VMM library on separate server in highly available configuration. If you want to deploy another VMM library server, the server should fulfill following requirements:

- Supported operating system: Windows Server 2008 or Windows Server 2008 R2
- Hardware management: Windows Remote Management 2.0
- CPU: at least 2.8 GHz
- RAM: at least 2 GB
- Hard disk space: varies based on the number and size of files that are stored

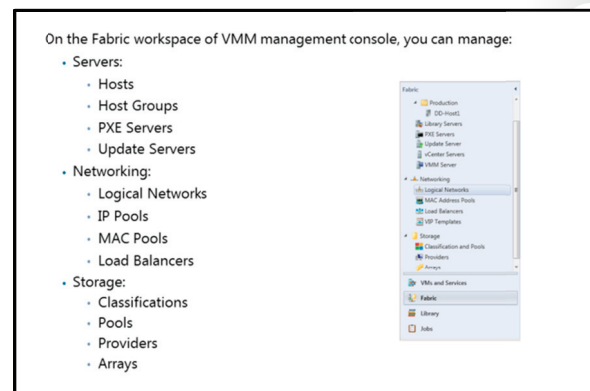
## Private Cloud Infrastructure Components in VMM

The key architectural concept in VMM is private cloud infrastructure. Similar to public cloud solutions, such as in Windows Azure™, private cloud infrastructure in VMM is an abstraction layer that shields the underlying technical complexities, and lets you manage defined resource pools of servers, networking, and storage in the enterprise infrastructure.

This concept is presented explicitly in the VMM management console user interface. With VMM, you can create a private cloud from Hyper-V, VMware ESX, and Citrix XenServer hosts, and benefit from cloud computing attributes, including self-servicing, resource pooling, and elasticity.

You can configure the following resources from the VMM management console Fabric workspace:

- **Servers.** In the Servers node, you can configure and manage several types of servers. Host groups contain virtualization hosts, which are the destinations for where you can deploy virtual machines. Library servers are the repositories of building blocks—such as images, .iso files, and templates—for creating virtual machines.
- **Networking.** In the VMM management console, the Networking node is where you can define logical networks, assign pools of static IPs and media access control (MAC) addresses, and integrate



load balancers. Logical networks are user-defined groupings of IP subnets and virtual local area networks (VLANs) to organize and simplify network assignments. Logical networks provide an abstraction of the underlying physical infrastructure, and enable an administrator to provision and isolate network traffic based on selected criteria such as connectivity properties and service level agreements (SLAs).

- **Storage.** Using the VMM 2012 admin console, an administrator can discover, classify, and provision remote storage on supported storage arrays. VMM 2012 uses the Microsoft Storage Management Service (which is enabled by default during the installation of VMM 2012), to communicate with external arrays.

## Managing Hosts and Host Groups with VMM

In addition to virtual machine management, VMM can also manage and deploy Hyper-V hosts. In VMM you can use technologies such as Windows Deployment Services to deploy Hyper-V hosts on bare metal machines and then manage it with VMM. When hosts are associated with VMM, you can configure several options, such as host reserves, quotas, permissions, cloud membership, and so on VMM can also manage Hyper-V failover clusters.

VMM provides two new features that help optimize power and resource usage on hosts managed by VMM: dynamic optimization and power optimization. Dynamic optimization balances the virtual machine load within a host cluster, while power optimization enables VMM to evacuate balanced cluster hosts, and then turn them off to save power.

The recommended way to organize hosts in VMM is to create host groups. This greatly simplifies management tasks. A host group enables you to apply settings to multiple hosts with a single action. By default, there is a single host group in VMM named All Hosts. However, if necessary, you can create additional groups for your environment.

Host groups are hierarchical. When you create a new child host group, it inherits the settings from the parent host group. When a child host group moves to a new parent host group, the child host group maintains its original settings except for Performance and Resource Optimization (PRO) settings, which are managed separately. When the settings in a parent host group change, you can apply those changes to child host groups.

You would use host groups in the following scenarios:

- Providing basic organization when you are managing lots of hosts and virtual machines. You can create custom views within the Hosts view and Virtual Machines view to provide easy monitoring and access to a host. For example, you might create a host group for each branch office in your organization.
- Reserving resources for use by hosts. Host reserves are useful when placing virtual machines on a host. Host reserves determine the CPU, memory, disk space, disk I/O capacity, and network capacity that are continuously available to the host operating system.
- Use the Host group properties action for the root host group All Hosts, to set default host reserves for all hosts that VMM manages. If you want to use more of the resources on some hosts instead of on other hosts, you can set host reserves differently for each host group.

- VMM can deploy and manage Hyper-V hosts, Hyper-V clusters and host groups
- Host groups simplify management tasks by using a single action to apply settings to multiple hosts
- Host group scenarios:
  - Provide basic organization when managing large numbers of hosts
  - Reserve resources for use by hosts
  - Designate hosts on which a user can create and operate their own virtual machines
  - Create private clouds

- Designating hosts on which users can create and operate their own virtual machines. When a VMM administrator adds self-service user roles, one part of role creation is to identify the hosts on which self-service users or groups in that role can create, operate, and manage their own virtual machines. We recommend that you designate a specific host group for this purpose.

## Deploying Virtual Machines with VMM

One of the advantages of using a virtualized environment that is managed by VMM is the flexibility that it provides to create and deploy new virtual machines quickly.

Using VMM, you can manually create a new virtual machine with new configuration settings and a new hard disk. You can then deploy the new virtual machine from one of following sources:

- An existing virtual hard disk (.vhd) file (blank or preconfigured)
- A virtual machine template
- A VMM library

In VMM there are several ways how you can create and deploy new virtual machine:

- Creating new virtual machine from an empty hard disk
- Creating new virtual machine based on pre-defined template
- Deploying a new virtual machine from VMM Library

You can create new virtual machines either by converting an existing physical computer, or by cloning an existing virtual machine.

### Creating a New Virtual Machine from an Existing VHD

You can create a new virtual machine based on either a blank VHD, or on a preconfigured VHD that contains a guest operating system. VMM provides two blank VHD templates that you can use to create new disks:

- Blank Disk – Small
- Blank Disk – Large

You can also use a blank VHD when you want to use an operating system with a PXE. Or, you can place an ISO image on a virtual DVD-ROM, and then install an operating system from scratch. This is an effective way to build a virtual machine's source image, which you can then use as a future template. To install the operating system on such a virtual machine, you can use an ISO image file from the library or from local disk, then map a physical drive from the host computer, or start the guest operating system setup through a network service boot.

If you have a library of VHDs that you want to use in your VMM environment, you can create a virtual machine from an existing VHD. You can also select existing VHDs when you deploy any operating system from which VMM cannot create a template, such as an operating system that is not Windows based.

When you create a new virtual machine using an existing VHD, you are basically creating a new virtual machine configuration that is associated with the VHD file. VMM will create a copy of the source VHD so that you do not have to move or modify the original.

In this scenario, the source VHD must meet the following requirements:

- Leave the Administrator password blank on the VHD as part of the System Preparation Tool (Sysprep) process.
- Install the Virtual Machine Additions on the virtual machine.
- Use Sysprep to prepare the operating system for duplication.

## Deploying from a Template

This method creates a new virtual machine based on a template from the VMM library. The template is a library resource, which links to a virtual hard disk drives that has a generalized operating system, hardware settings, and guest operating system settings. You use the guest operating system settings to configure operating system settings such as computer name, local administrator password, and domain membership.

The deployment process does not modify the template, which you can reuse multiple times. If you are creating virtual machines in the Self-Service Portal, you must use a template.

The following requirements apply if you want to deploy a new virtual machine from a template:

- You must install a supported operating system on the VHD.
- You must leave the Administrator password blank on the VHD as part of the Sysprep process. However, you do not have to leave blank the Administrator password for the guest operating system profile.
- For customized templates, you must prepare the operating system on the VHD by removing computer identity information. For Windows operating systems, you can prepare the VHD by using Sysprep.

## Deploying from the VMM Library

If you deploy a virtual machine from the library, the virtual machine is removed from the library, and then placed on the selected host. When you use this method, you must provide the following details in the Deploy Virtual Machine wizard:

- The host for deployment. The template that you use provides a list of potential hosts and their ratings.
- The path of the virtual machine files on the host.

The virtual networks used for the virtual machine. You are presented with a list of existing virtual networks on the host.

## What Are Services and Service Templates?

Services are a new concept in VMM. You must understand services fully before you deploy a private cloud infrastructure.

### Traditional Services Scenario

When we think about services, we usually refer to an application or set of applications that provide some service to end-users. For example, we can deploy various types of web-based services, but we can also implement a service such as email. In a non-cloud computing scenario, deployment of any type of service usually requires users, developers, and administrators to work together through the phases of creating a service, deploying a service, testing the service, and maintaining the service.

- In terms of VMM clouds, a service is a set of one or more virtual machines that are deployed together and managed as a single entity
- A service template encapsulates all necessary components required to deploy and run a new instance of an application
  - Service is deployed by administrator or end-user
  - Service can contain several different components
  - Service can be deployed to a private cloud or to host group
  - Administrator creates a service template in VMM
  - Application owner deploys a service based on the service template
  - App Controller or VMM Manager console can be used to deploy service based on template

A service frequently includes several computers that must work together to provide a service to end-users. For example, a web-based service is usually an application that deploys on a web server, connects to a database server (which can be hosted on another computer), and performs authentication on an Active Directory domain controller. Enabling this application requires three roles, and possibly three computers: a web server, a database server, and a domain controller. Deploying a test environment for a service such as this can be time and resource consuming. Ideally, developers work with IT administrators to create an environment where they can deploy and test their web application.

### **Concept of a Service in a Private Cloud Scenario**

With the concept of a private cloud, how you deal with services can change significantly. You can prepare the environment for a service, and then let developers deploy it by using a self-service application such as App Controller.

In VMM, a service is a set of one or more virtual machines that you deploy and manage together as a single entity. You configure these machines to run together to provide a service. In VMM in Windows Server 2008, users were able to deploy new virtual machines by using Self Service Portal. In VMM, end-users can deploy new services. By deploying a service, users are actually deploying the whole infrastructure, including the virtual machines, network connections, and applications that are required to make the service work.

However, you can use services to deploy only a single virtual machine without any specific purpose. Instead of deploying virtual machines in the historic way, you can now create a service that will deploy a virtual machine with—for example—Windows Server 2008 R2, and with several roles and features preinstalled and joined to domain. This simplifies the process of creating and later updating new virtual machines.

Deploying a new service requires a high level of automation and predefined components, and requires management software support. This is why VMM provides service templates. A service template is a template that encapsulates everything required to deploy and run a new instance of an application. Just as a private cloud user can create new virtual machines on demand, the user can also use service templates to install and start new applications on demand.

### **Process for Deploying a New Service**

Follow this procedure when you use service templates in VMM to deploy a new service/application:

1. The system administrator creates and configures service templates in VMM by using Service Template Designer.
2. The end-user application owner (for example, a developer who has to deploy the application environment) opens the App Controller console, and requests a new service deployment based on available service templates that he or she can access. The developer can deploy the service to a private cloud where a user has access. As an alternative to App Controller, the user can also use the VMM Manager console.
3. A request is submitted and evaluated by the VMM Server. VMM searches for available resources in the private cloud, then calculates the user quota and verifies that the cloud is capable for the requested service deployment.
4. Whereas the service is created automatically, the virtual machines and applications (if any) are deployed on the host selected by VMM.
5. The user application owner gains control over service virtual machines through the App Controller console, or by RDP.
6. If you need manual approval for resource creation, you can use Microsoft System Center 2012 - Service Manager to create workflows for this purpose.

## Information Included in the Service Template

The service template includes information about the virtual machines that are deployed as part of the service, which applications to install on the virtual machines, and the networking configuration needed for the service (including the use of a load balancer). The service template can use existing virtual machine templates. You can define the service without using any existing virtual machine templates. However, it is much easier to build a template if you have already created virtual machine templates. After you create the service template, you configure it for deployment using the Configure Deployment option.

## Physical to Virtual and Virtual to Virtual Migrations

Many organizations have physical servers that they do not use fully. VMM can convert existing physical computers into virtual machines through a process known as physical-to-virtual (P2V) conversion. VMM simplifies P2V by providing a task-based wizard to automate much of the conversion process. Because the P2V process is scriptable, you can start large-scale P2V conversions through the Windows PowerShell (Powershell.exe) command line.

VMM converts an operating system that is running on physical hardware to an operating system that is running in a virtual machine in Hyper-V environment. VMM provides a conversion wizard, which automates much of the conversion process.

- P2V process converts an operating system that is running on physical hardware to an operating system running inside a virtual machine
- V2V process converts existing VMware virtual machines to virtual machines running on Hyper-V

During a P2V conversion process, VMM makes disk images of the hard disks on the physical computer. It creates VHD files for the new virtual machine, using the disk images as a basis. Also, it creates a hardware configuration for the virtual machine similar to, or the same as, the hardware in the physical computer.

The new virtual machine has the same computer identity as the physical computer on which it is based. Because of that, we do not recommend that you use both a physical computer and its virtual replica concurrently. After the P2V conversion is finished, you typically disconnect the physical computer from the network and decommission it.

P2V conversion is finished in Online or Offline mode. In Online mode, the source operating system is running during the conversion process. In Offline mode, the operating system is not running, and conversion occurs through the Windows Preinstallation Environment (Windows PE). Later topics in this lesson describe these modes and their specifics.

In addition to converting underused physical computers, VMM supports the management, migration and conversions of other virtual machines that you create in VMware environment. You can convert these virtual machines to Hyper-V virtual machines, place them on Hyper-V hosts, and then manage them under the VMM Administrator Console. Also, VMM and Hyper-V support migrating virtual machines from one host to another with minimal or zero downtime.

VMM 2012 allows you to convert existing VMware virtual machines to virtual machines running on the Hyper-V platform. This process is known as a V2V conversion. With V2V conversion, administrators can easily and quickly consolidate a virtual environment that is running various virtual platforms without rebuilding virtual machines from scratch or moving data.



VMM allows you to copy existing VMware virtual machines and create Hyper-V virtual machines. You can copy VMware virtual machines that are on an ESX Server host, in the VMM library, or on a Windows share. Although V2V is called a conversion, V2V is a read-only operation that does not delete or affect the original source virtual machine. Also, the term conversion is dedicated only to the process of converting VMware virtual machines. The term migration is used for Virtual Server machines.

During the conversion process, the VMM converts the VMware .vmdk files to .vhd files, and makes the operating system on the virtual machine compatible with Microsoft virtualization technologies. The virtual machine that the wizard creates matches VMware virtual machine properties, including name, description, memory, and disk-to-bus assignment.

## Considerations for Deploying a Highly Available VMM Server

VMM now supports a highly available VMM Server. You can use failover clustering to achieve high availability for VMM, because VMM is now a cluster-aware application. However, you should consider several things before you deploy a VMM cluster.

Before you begin the installation of a highly available VMM management server, ensure the following:

- You have installed and configured a failover cluster that is running Windows Server 2008 R2, Windows Server 2008 R2 SP1, or Windows Server 2012.
- All computers on which you install the highly available VMM management server meet the minimum hardware requirements, and all prerequisite software is installed on all computers.
- You have created a domain account to be used by the VMM service. You must use a domain user account for a highly available VMM management server.
- You are prepared to use distributed key management to store encryption keys in AD DS. You must use distributed key management for a highly available VMM management server.
- You have a computer with a supported SQL Server version installed and running. Unlike VMM 2008 R2, VMM will not automatically install a SQL Server Express edition.

### Considerations for deploying highly available:

- VMM is cluster-aware and can be highly available
- When deploying VMM in a cluster, the service account must be the domain account
- Use Distributed Key Management for encryption keys
- Make database and library servers highly available
- Do not install a self-service portal on a clustered VMM server
- Use the Failover Cluster Manager to perform a planned failover

## Highly Available Databases and Library Servers

To achieve full redundancy, we recommend that you use a highly available SQL Server. You should install a highly available SQL Server on a separate failover cluster from the failover cluster on which you are installing the highly available VMM management server. Similarly, we also recommend that you use a highly available file server for hosting your library shares.

## Self Service Portal and Clustered VMM Server

For best practices, do not install the VMM Self-Service Portal on the same computer as the highly available VMM management server. If your VMM Self-Service Portal currently resides on the same computer as the VMM server, we recommend that you uninstall the VMM Self-Service Portal for VMM 2008 R2 SP1 before upgrading to VMM. We also recommend that you install the VMM Self-Service Portal on a highly available web server to achieve redundancy and load balancing.

**Failover Cluster Manager**

You cannot perform a planned failover (for example, to install a security update or do maintenance on a cluster node) by using the VMM console. Instead, to perform a planned failover, use the Failover Cluster Manager console.

During a planned failover, ensure that there are no tasks actively running on the VMM management server. Any tasks that are executing during a failover will be stopped and will not restart automatically. Any connections to a highly available VMM management server from the VMM console or the VMM Self-Service Portal will also be lost during a failover. However, the VMM console can reconnect automatically to the highly available VMM management server after a failover if it was opened before you performed failover to another VMM server.



## Lab: Implementing Failover Clustering with Hyper-V

### Scenario

The initial deployment of virtual machines on Hyper-V is very successful for A. Datum. As a next step in the deployment, A. Datum is now considering ways to ensure that the services and applications deployed on the virtual machines are highly available. As part of the implementation of high availability for most network services and applications, A. Datum is also considering options for making the virtual machines that run on Hyper-V highly available.

As one of the senior network administrators at A. Datum, you are responsible for integrating Hyper-V with failover clustering in order to ensure that the virtual machines deployed on Hyper-V are highly available. You are responsible for planning the virtual machine and storage configuration, and for implementing the virtual machines as highly available services on the Failover Cluster. Also, you are considering some other techniques for virtual machines high availability such as Hyper-V replica.

### Lab Setup

Estimated time: **75 minutes**

Virtual Machines	20417A-LON-DC1 20417A-LON-SVR1
User Name	Adatum\Administrator
Password	Pa\$\$w0rd

This lab should be performed with a partner. To perform this lab, you must boot the host computers to Windows Server 2012. The host computers should be in this state from the previous lab in Module 8. Make sure that you and your partner have booted into different hosts (one should boot to LON-Host1 and the other should boot to LON-Host2). Also, make sure that LON-DC1 is imported on LON-Host1 and LON-SVR1 is imported on LON-Host2, and that these VMs are started.

### Exercise 1: Configuring Hyper-V Replicas

#### Scenario

Before you start with cluster deployment, you decided to evaluate new technology in Hyper-V 3.0, for replicating virtual machines between hosts. You want to be able to manually mount a copy of virtual machine on another host if active copy (or host) fails.

The main tasks for this exercise are as follows:

1. Import LON-CORE virtual machine on LON-HOST1.
2. Configure a replica on both host machines.
3. Configure replication for LON-CORE virtual machine.
4. Validate a planned failover to the replica site.

► **Task 1: Import LON-CORE virtual machine on LON-HOST1**

- On LON-HOST1, open Hyper-V Manager and import the 20417A-LON-CORE virtual machine.
  - Use path **E:\Program Files\Microsoft Learning\20417\Drives\20417A-LON-CORE**
  - Accept default values.



**Note:** The drive letter may be different based upon the number of drives on the physical host machine.

► **Task 2: Configure a replica on both host machines**

1. On LON-HOST1 and LON-HOST2 configure each server to be Hyper-V replica server.
  - Use Kerberos (HTTP) for authentication.
  - Enable replication from any authenticated server.
  - Create and use folder **E:\VMReplica** as a default location to store replica files.
2. Enable the firewall rule named **Hyper-V Replica HTTP Listener (TCP-In)** on both hosts.

► **Task 3: Configure replication for LON-CORE virtual machine**

1. On LON-HOST1 enable replication for the **20417A-LON-CORE** virtual machine.
  - Use Kerberos (HTTP)
  - Select to have only latest recovery point available
  - Start replication immediately.
2. Wait for initial replication to finish and make sure that **20417A-LON-CORE** VM has appeared in Hyper-V Manager console on LON-HOST2.

► **Task 4: Validate a planned failover to the replica site**

1. On LON-HOST2, view replication health for **20417A-LON-CORE**.
2. On LON-HOST1, perform planned failover to LON-HOST2. Verify that 20417A-LON-CORE is running on LON-HOST2.
3. On LON-HOST1, remove replication for 20417A-LON-CORE.
4. On LON-HOST2, shut down 20417A-LON-CORE.

**Results:** After completing this exercise you will have Hyper-V replica configured.

## Exercise 2: Configuring a Failover Cluster for Hyper-V

### Scenario

A. Datum has several virtual machines that are hosting important services that must be highly available. Because these services are not cluster-aware, A. Datum decided to implement Failover cluster on the Hyper-V host level. You plan to use iSCSI drives as storage for these virtual machines.

The main tasks for this exercise are as follows:

1. Connect to iSCSI target from both host machines.
2. Configure failover clustering on both host machines.
3. Configure disks for failover cluster.

### ► Task 1: Connect to iSCSI target from both host machines

1. On LON-HOST1, start **iSCSI initiator**.
2. Use 172.16.0.21 address to discover and connect to iSCSI target.
3. On LON-HOST2, start **iSCSI initiator**.
4. Use 172.16.0.21 address to discover and connect to iSCSI target.
5. On LON-HOST2, open Disk Management and initialize and bring online all iSCSI drives
  - Format the first drive and name it **ClusterDisk**
  - Format the second drive and name it **ClusterVMs**
  - Format the third drive and name it **Quorum**
6. On LON-HOST1, open Disk Management and bring online all three iSCSI drives.

### ► Task 2: Configure failover clustering on both host machines

1. On LON-HOST1 and LON-HOST2, install the failover clustering feature.
2. On LON-HOST1, create a failover cluster:
  - Add **Lon-host1** and **Lon-Host2**
  - Name it **VMCluster**
  - Assign the **172.16.0.126** address
  - Deselect the option to Add all eligible storage to the cluster

### ► Task 3: Configure disks for failover cluster

1. In Failover Cluster Manager on LON-HOST1, add all three iSCSI disks to the cluster.
2. Verify that all three iSCSI disks appear available for cluster storage.
3. Add the disk with the volume name of **ClusterVMs** to **Cluster Shared Volumes**.
4. From the **VMCluster.adatum.com** node, select **More Actions** and then configure the Cluster Quorum Settings to use typical settings.

## Exercise 3: Configuring a Highly Available Virtual Machine

### Scenario

After you have configured the Hyper-V failover cluster, you want to add virtual machines as Highly Available resources. Also, you want to evaluate Live migration as well as test storage migration.

The main tasks for this exercise are as follows:

1. Move Virtual Machine Storage to iSCSI Target.
2. Configure the Virtual Machine as Highly Available.
3. Perform a Live Migration for the Virtual Machine.
4. Perform a Storage Migration for the Virtual Machine.
5. To Prepare for Next Module.

► **Task 1: Move Virtual Machine Storage to iSCSI Target**

1. Make sure that LON-HOST1 is the owner of the ClusterVMs disk. If it is not, move the ClusterVMs disk to LON-HOST1.
2. On LON-HOST1, open Windows Explorer and browse to **E:\Program Files\Microsoft Learning\20417\Drives\20410A-LON-CORE\Virtual Hard Disks** and move the **20417A-LON-CORE.vhd** virtual hard drive file to the **C:\ClusterStorage\Volume1** location.

► **Task 2: Configure the Virtual Machine as Highly Available**

1. In Failover Cluster Manager, click the **Roles** node, and then start the New Virtual Machine wizard.
  - Select **LON-Host2** as the cluster node.
  - Name the computer as **TestClusterVM**.
  - Store the file at **C:\ClusterStorage\Volume1**.
  - Assign **1536MB** of RAM to the **TestClusterVM**.
  - Connect machine to existing virtual hard disk drive **20417A-LON-CORE.vhd** located at **C:\ClusterStorage\Volume1**.
2. From the **Roles** node, start the virtual machine.

► **Task 3: Perform a Live Migration for the Virtual Machine**

1. On LON-HOST2, in Failover Cluster Manager, start Live Migration failover of TestClusterVM from Lon-Host2 to Lon-host1.
2. Connect to TestClusterVM and make sure that you can operate it.

► **Task 4: Perform a Storage Migration for the Virtual Machine**

1. On LON-HOST1, open Hyper-V Manager and start LON-GUEST1.
2. Perform a Move operation on LON-GUEST1. Move the VM from its current location to C:\GUEST1.
3. Check whether machine is operational during move process.
4. When complete, shut down all running virtual machines.

► **To prepare for Next Module**

- Restart both host machines, and select to boot to Windows Server 2008 R2. Log on to the host machines as directed by your instructor.

## Module Review and Takeaways

### Best Practices

- Develop standard configurations before you implement highly available virtual machines. The host computers should be configured as close to identically as possible. To make sure that you have a consistent Hyper-V platform, you should configure standard network names, and use consistent naming standards for CSV volumes.
- Implement VMM. VMM provides a management layer on top of Hyper-V and Failover Cluster Management that can block you from making mistakes when you manage highly available virtual machines. For example, it blocks you from creating virtual machines on storage that is inaccessible from all nodes in the cluster.

### Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Virtual machine failover fails after I implement CSV and migrate the shared storage to CSV.	
A virtual machine fails over to another node in the host cluster, but loses all network connectivity.	
Four hours after restarting a Hyper-V host that is a member of a host cluster, there are still no virtual machines running on the host.	

### Review Question

Do you have to implement CSV in order to provide high availability for virtual machines in VMM in Windows Server 2008 R2?

**MCT USE ONLY. STUDENT USE PROHIBITED**

# Module 10

## Implementing Dynamic Access Control

### Contents:

Module Overview	10-1
<b>Lesson 1:</b> Overview of Dynamic Access Control	10-2
<b>Lesson 2:</b> Planning for a Dynamic Access Control Implementation	10-8
<b>Lesson 3:</b> Implementing and Configuring Dynamic Access Control	10-13
<b>Lab:</b> Implementing Dynamic Access Control	10-22
Module Review and Takeaways	10-31

## Module Overview

Windows Server 2012 introduces Dynamic Access Control for enhancing access control for file- and folder-based resources. Dynamic Access Control extends regular New Technology File System (NTFS)-based access control by enabling administrators to use claims, resource properties, rules and conditional expressions to manage access. In this module you will learn about Dynamic Access Control and how to plan for and implement it.

### Objectives

After completing this module, you will be able to:

- Describe Dynamic Access Control and its components.
- Plan for Dynamic Access Control implementation.
- Configure Dynamic Access Control.

## Lesson 1

# Overview of Dynamic Access Control

Dynamic Access Control is a new technology for access management in Windows Server 2012. It offers a new way of controlling access to resources. Before you implement this technology, you should learn how it works and which components it uses. This lesson presents an overview of Dynamic Access Control.

### Lesson Objectives

After completing this lesson, you will be able to:

- Define Dynamic Access Control.
- Describe the foundation technologies for Dynamic Access Control.
- Compare Dynamic Access Control with alternative or similar technologies, such as NTFS permissions and Active Directory Rights Management Services (AD RMS).
- Define identity.
- Define claim and claim types.
- Define Central Access Policy.

### What Is Dynamic Access Control?

Because most of the data in an organization is stored on file servers, IT administrators must help provide security and access control to file server resources. In previous versions of Windows Server, most access control to file server resources was controlled by using NTFS permissions and access control lists.

Dynamic Access Control in Windows Server 2012 is a new access control mechanism for file-system resources. It enables administrators to define central file-access policies that can apply to every file server in the organization. Dynamic Access Control helps implement security over file servers, in addition to any existing share and NTFS permissions. Dynamic Access Control ensures that regardless of how the share and NTFS permissions might change, this central overriding policy is still enforced. What Dynamic Access Control does is combining multiple criteria into the access decision. This is something that NTFS permissions can't achieve.

Dynamic Access Control provides:

- **Data identification.** You can use automatic and manual classification of files to tag data in file servers across the organization.
- **Access control to files.** Central access policies enable organizations to define (for example, who can access health information within the organization).
- **Auditing of access to files.** Central audit policies for compliance reporting and forensic analysis. For example, you can identify who accessed highly sensitive information.
- **Optional RMS protection integration.** Automatic Rights Management Services (RMS) encryption for sensitive Microsoft® Office documents. For example, you can configure RMS to encrypt all documents containing Health Insurance Portability and Accountability Act (HIPAA) information.

- Dynamic Access Control helps provide secure file server-based resources over all file server based resources
- Dynamic Access Control provides:
  - Data Identification
  - Access Control to files
  - Auditing of access to files
  - Optional RMS protection integration



Dynamic Access Control focuses on four main end-to-end scenarios:

- **Central access policy for access to files.** Enable organizations to set safety net policies that reflect the business and regulatory compliance.
- **Auditing for compliance and analysis.** Enable targeted auditing across file servers for compliance reporting and forensic analysis.
- **Protecting sensitive information.** Identify and protect sensitive information both in a Windows Server 2012 environment and when it leaves the Windows Server 2012 environment.
- **Access denied remediation.** Improve the access denied experience to reduce the helpdesk load and incident time for troubleshooting.

Dynamic Access Control provides a flexible way to apply and manage access and auditing to domain-based file servers. Dynamic Access Control uses claims in the authentication token, resource properties on the resource, and conditional expressions within permission and auditing entries. With this combination of features, you can now grant access to files and folders based on Active Directory attributes.

## Foundation Technologies for Dynamic Access Control

Dynamic Access Control combines many Windows Server 2012 technologies to provide a robust, flexible, and granular authorization and auditing experience. Dynamic Access Control uses these fundamental technologies:

- **Network protocols, such as TCP/IP, Remote Procedure Call (RPC), Server Message Block (SMB), and Lightweight Directory Access Protocol (LDAP).** For network communications between hosts, interaction with file system and directory lookups, respectively.
- **Domain Name System (DNS).** For host name resolution.
- **Active Directory Domain Services (AD DS) and its dependent technologies.** For enterprise network management.
- **The Microsoft Kerberos v5 implementation including FAST Search and Compound Identity.** For secure authentication.
- **Windows Security (local security authority [LSA], Netlogon).** For secure logon transactions.
- **File Classifications.** For file categorization.
- **Auditing.** For secure monitoring and accountability.

Dynamic Access Control relies on many technologies in Windows Server 2012, such as:

- Network protocols
- DNS
- AD DS
- Kerberos
- Windows Security
- File Classifications
- Auditing

Several components and technologies were updated in Windows Server 2012 to support Dynamic Access Control. The most important updates are:

- A new Windows authorization and audit engine that can process conditional expressions and central policies.
- Kerberos authentication support for user claims and device claims.

- Improved File Classification Infrastructure.
- Optional Rights Management Services (RMS) extensibility support so that partners can provide solutions that encrypt non-Office files.

## Dynamic Access Control Versus Alternative Technologies

Dynamic Access Control is a new technology for controlling access to file based resources. It does not overlap with older well-known technologies with similar purpose. Instead, Dynamic Access Control extends the functionality of older technologies for controlling file-based resource access.

In previous versions of Windows Server, the basic mechanism for file and folder access control was NTFS permissions. By using NTFS permissions and their Access Control Lists (ACLs), administrators can control access to resources, based on user name or group membership, and the level of access, such as Read-only, Change, Full Control, etc. However, once you provide someone with, for example, Read-only access to a document, you cannot prevent that person from copying the content of that document into a new document or printing the document. By implementing AD RMS, you can establish an additional level of control. Unlike, NTFS permissions, which are not application aware, AD RMS sets a policy that can control document access inside the application that is being used to open it. By implementing AD RMS, you enable users to additionally protect documents within applications.


However, you cannot set conditional access to files by using NTFS and AD RMS. For example, you cannot set NTFS permissions in a way that users can access a document if they are a member of some specific group and have the attribute EmployeeType set to FTE. Or, you might want to set permissions so that only users that have a department attribute populated with the same value as the department attribute for the resource can access the content. You can accomplish this by using conditional expressions.

For these scenarios, in Windows Server 2012, you can use Dynamic Access Control. In simple terms, Dynamic Access Control enables you to count attribute values on users or resource objects, when providing or denying access.

- NTFS permissions and ACLs provide access control based on user's SID or group membership SID
- AD RMS provides deeper protection for documents by controlling how applications can use them
- Dynamic Access control provides access control based on claims – values of specific attributes

## What Is an Identity?

We usually define identity as a set of data that uniquely describes a person or a thing (sometimes referred to as subject or entity) and contains information about the subject's relationships to other entities. Identity is usually proved by using some trusted source of information. For example, when you go to the airport, you show your passport. Your passport contains your name, address, date of birth, and photograph. Each item of personal information is a claim that is made about you by the country issuing your passport. Your country ensures the information published in



a passport is accurate for the passport owner. Since you usually use the passport outside of your country of residence, other countries must also trust the information in your passport. They must trust the organization that issued your passport and consider it reliable. Based on that trust, other countries grant you access their territory (which can be considered as a resource).

In other words, to access resources in other countries, each person is required to have a document (passport) that is issued by a reliable and trusted source and that contains some critical claims that describe the person.

The Windows operating system uses a similar concept of identity. An administrator creates a user account for person in AD DS. The domain controller publishes user account information, such as a security identifier, and group membership attributes. Windows creates an authorization token when a user accesses a resource.

To continue the analogy, you are the user. The authorization token is the passport. Each unique piece of information in the authorization token is a claim made about your user account. Domain controllers publish these claims. Domain-joined computers and domain users trust domain controllers to provide authoritative information.

We can then say that Identity, with respect to authentication and authorization, is simply information published about an entity from a trusted source. The information is considered authoritative because the source is trusted.

Earlier versions of Windows Server used the security identifier (SID) to represent identity of a user or computer. Users authenticate to the domain with a specific user name and password. The unique logon name translates into the SID. The domain controller validates the password and publishes the SID of the security principal and the SIDs of all the group of which the principal is a member. The domain controller "claims" the user's SID is valid and should be used as the identity of the user. All domain members trust the domain controller; therefore, the response is treated as authoritative.

Identity is not limited to the user's SID. Applications can use any information about the user as a form of identity, provided that the application trusts the source of the information to be authoritative. For example, many applications implement role-based access control. Role-based access control limits access to resources based on whether the user is a member of a specific role. SharePoint Server is good example of software that implements role-based security. Windows Server 2012 can also take advantage of these options to extend and enhance the way identity is determined for a security principal.

## What Is a Claim?

Windows Server 2008 and Windows Server 2003 use claims in Active Directory Federation Services (AD FS). In this context, claims are statements made about users (for example, name, identity, key, group, privilege, or capability), which are understood by both partners in an AD FS federation. AD FS also introduced AD DS-based claims and the ability to convert AD DS-based claim data into Secure Application Markup Language (SAML) format. In previous versions of AD FS, the only attributes that could be retrieved from AD DS and directly incorporated into a claim was SID information for user and group accounts. All other claim information was defined within and referenced from a separate database, known as an attribute store. New in Windows Server 2012 is the

- Claims are statements made by AD DS about specific user of computer object in AD DS
- AD DS in Windows Server 2012 supports:
  - User claims
  - Device claims

capability to read and use any attribute directly from AD DS. It is not necessary to use a separate AD FS attribute store to hold this type of information for Active Directory-based computer or user accounts.

By definition, a claim is something that AD DS states about specific object (usually a user or computer). A claim provides some information from trusted source about an entity. Some examples of claims are the SID of a user or computer, the department classification of a file, and the health state of a computer. All these claims state something about a specific object. In more technical language, claims state the value of a specific attribute of a user or computer object.

An entity can contain more than one claim. When configuring resource access, any combination of those claims can be used to authorize access to resources.

In Windows Server 2012, authorization mechanism is extended so that you can use claims for authorization on files and folders, besides using just NTFS permissions, based on user's SID or group SIDs. By using claims, you can now base your access control not only on SID, but also on other attribute values. Because SID is also an attribute of a user or computer object, we can say that older authorization mechanisms are, in a way, subsets of claims-based authorization.

Windows Server 2012 introduces two new types of claims: user claims and device claims. Windows Server 2012 continues to enable you to use group membership for authorization decisions.

### User Claim

A user claim is information provided by a Windows Server 2012 domain controller about a user. Windows Server 2012 domain controllers can use most AD DS user attributes as claim information. This provides administrators with wide range of possibilities to configure and use claims for access control.

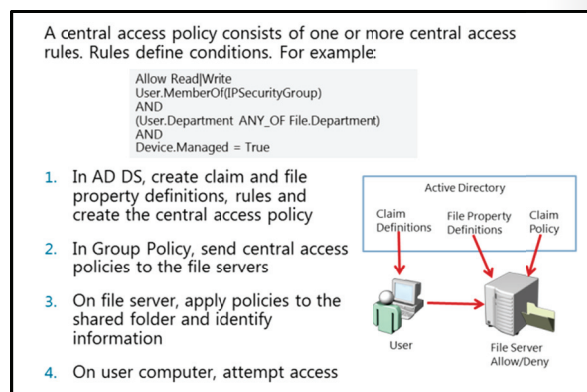
### Device Claim

A device claim is information provided by a Windows Server 2012 domain controller about a device represented by a computer account in AD DS. As with a user claim, a device claim, often called a computer claim, can use most of the AD DS attributes that are applicable to computer objects.

## What is a Central Access Policy?

One of the fundamental components in Dynamic Access Control technology is Central Access Policy. It is a feature in Windows Server 2012 that enables administrators to create a policy that is applied to one or more file servers. This policy is created in Active Directory Administrative Center, stored in AD DS, and applied by using Group Policy. Central Access Policy contains one or more Central Access Policy rules. Each rule contains settings that determine applicability and permissions.

Before you create Central Access Policy, it is mandatory that you create at least one Central Access Rule. Central Access Rule defines all parameters and conditions that control access to specific resource.



A central access rule has three configurable parts:

- **Name:** For each Central Access Rule you should configure descriptive name.
- **Target resources:** A condition that defines which data the policy applies to. This is defined by specifying an attribute and its value. For example, a particular central policy might apply to any data classified as Sensitive. You can also choose to apply rule to all resources where Central Access Policy applies.
- **Permissions:** A list of one or more access control entries (ACEs) that define who can access the data. For example, you can specify Full Control Access to a user with attribute EmployeeType populated with FTE. This is the key component of each Central Access rule. You can combine and group conditions that you place in central access rule. You can set permission as proposed (for staging purposes) or current.

After you configure one or more central access rules, you then place these rules in Central Access Policy which is applied to the resources.

Central Access Policy enhances, but does not replace, the local access policies or discretionary access control lists (DACL) that are applied to files and folders on a specific server. For example, if a DACL on a file allows access to a specific user, but a central policy that is applied to the file restricts access to the same user, the user cannot obtain access to the file. Likewise, if the central access policy allows access but the DACL does not allow access, then the user cannot obtain access to the file.

Before you implement Central Access Policy, you should perform these steps:

1. Create claims and connect it with attributes on user or computer objects.
2. Create file property definitions.
3. Create one or more Central Access Rules
4. Create a Central Access Policy object and place rules in it.
5. Use Group Policy to deploy the policy to file servers. By doing this, you make file servers aware that a Central Access Policy exists in AD DS.

On the file server, apply that policy to a specific shared folder.

## Lesson 2

# Planning for a Dynamic Access Control Implementation

Dynamic Access Control is a technology that requires detailed planning before implementation. You should identify reasons to implement Dynamic Access Control, as well as plan for Central Access Policy, file classifications, auditing and access denied assistance. In this lesson, you will learn about planning Dynamic Access Control.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe reasons for implementing Dynamic Access Control.
- Plan for Central Access Policy.
- Plan for File Classifications.
- Plan for File Access Auditing.
- Plan for Access Denied Assistance.
- Plan for policy changes.

### Reasons for Implementing Dynamic Access Control

Before you implement Dynamic Access Control you should clearly identify the reasons for implementation. This technology should be well designed before implementation, so it is very important to have business case that requires implementation of Dynamic Access Control. An improperly planned implementation can result in some users being denied access to data they need, while other users are inappropriately granted access to data to which they should otherwise be restricted.

Most common reasons for implementing Dynamic Access Control:

- Inability to achieve desired results with NTFS
- Requirement for access control based on attributes

The most common reason to implement Dynamic Access Control is to extend functionality of an existing model for access control management. Most companies use NTFS and share permissions to implement access control for file and folder resources. In most cases, NTFS is sufficient, but in some scenarios it does not work. For example, you cannot use NTFS ACL to protect a resource on a file server so that a user must be member of two groups at the same time to access the resource. This relatively simple scenario requires a new technology.

In general, you must use Dynamic Access Control instead of traditional methods for implementing access control when you want to use more specific information for authorization purposes. NTFS and share permissions use only user or group objects, but if you want to implement more complex access control scenarios, you should use Dynamic Access Control.

## Planning for Central Access Policy

Implementing Central Access Policy is not mandatory for Dynamic Access Control. However, for consistent configuration of access control on all file servers, we recommended implementing Central Access Policy. By doing that, you enable all file servers to use Central Access Policy when protecting content in shared folders.

If you decide to implement Central Access Policy, you should make a detailed plan before implementation. When planning Central Access Policy you must clearly identify and understand the business requirements for implementing Central Access Policy and Dynamic Access Control.

When planning Central Access Policy, you should:

- Identify business case
- Identify resources to be protected
- Understand business requirements
- Translate business requirements to conditional expressions
- Define claim types, resource properties, and rules

You should first identify the resources that you want to protect. If all these resources are on one file server or in just one folder, then you might not have to implement Central Access Policy. Instead, you can configure conditional access on the folder's ACL. If resources are distributed across several servers or folders, then you can benefit from deploying Central Access Policy. Examples of data that might require protecting are payroll records, medical history data, employee personal information, company customer lists, and so on. You can also use targeting within central access rules to identify resources where you want to apply central access policy.

After you identify resources, you should define criteria for protection. This is usually defined by business requirements. Some examples are:

- All documents that have property confidentiality set to high must be available only to managers.
- Marketing documents from each country should be accessible only to marketing people from the same country.
- Only full time employees should be able to access technical documentation from previous projects.

A central access policy is targeted to provide an easy interpretation from a business requirement language to an authorization language.

The next step in the planning process is to translate the policies you require into expressions. In the case of Dynamic Access Control, expressions are attributes associated with both the resources (files and folders) and the user or device that seeks access to the resources. These expressions state additional identification requirements that must be met in order to access protected data. Values associated with any expressions on the resource obligates the user or device to produce the same value

Next, you should break down the expressions that you created and determine what claim types, resource properties, and device claims you must create to deploy your policies. In other words, you must identify the attributes for access filtering.



**Note:** You are not required to use user claims to deploy central access policies. You can use security groups to represent user identities.



## Planning File Classifications

When planning implementation of Dynamic Access Control, you should include File Classifications in complete scenarios. Although file classifications are not mandatory for Dynamic Access Control, they can greatly enhance the automation of the entire process. For example, if you require that all documents with classification Confidentiality: High must be accessible to top management only, regardless of the server on which the documents exist, you should first ask yourself how you identify these documents, and how to classify them appropriately.

When planning for file classification implementation do following:

- Identify classifications
- Determine method for classification
- Determine schedule
- Perform review

File Classification Infrastructure uses classification rules to automatically scan files and classify them according to the contents of the file. Classification properties are defined centrally in AD DS so that these definitions can be shared across file servers in the organization. You can create classification rules that scan files for a standard string or for a string that matches a pattern (regular expression). When a configured classification parameter is found in a file, that file is classified as configured in the classification rule.

When planning for file classifications, you should do following:

- Identify which classification or classifications you want to apply on documents.
- Determine the method to identify documents for classification.
- Determine the schedule for automatic classifications.
- Establish a review of classification success.

You configure file classifications in the File Server Resource Manager console.

When you have a defined the classifications, you can plan the implementation of Dynamic Access Control by defining conditional expressions that enable you to control access to high confidential documents based on particular user attributes.

## Planning File Access Auditing

In Windows Server 2008 R2 and Windows Server 2012, you can use new advanced audit policies to implement more detailed and more precise auditing on file system. In Windows Server 2012, you can also implement auditing together with Dynamic Access Control to take advantage of the new Windows Security auditing capabilities. By using conditional expressions, you can configure auditing to be implemented only in specific cases. For example, you want to audit attempts to open shared folders only by users located in countries other than the country where the shared folder is located.

File access auditing:

- Track changes to user and machine attributes
- Get more information from user logon events
- Provide more information from object access auditing
- Track changes to Central Access Policies, Central Access Rules and Claims
- Track changes to file attributes



With Global Object Access Auditing, administrators can define computer SACLs per object type for either the file system or registry. The specified SACL is then automatically applied to every object of that type. You can use a Global Object Access Audit Policy to enforce the object access audit policy for a computer, file share, or registry without configuring and propagating conventional SACLs. Configuring and propagating SACLs is a more complex administrative task and it is difficult to verify, particularly if you must verify to an auditor that security policy is being enforced.

Auditors can prove that every resource in the system is protected by an audit policy by just viewing the contents of the Global Object Access Auditing policy setting.

Resource SACLs are also useful for diagnostic scenarios. For example, setting a Global Object Access Auditing policy to log all activity for a specific user and enabling the Access Failures audit policies in a resource (file system, registry) can help administrators quickly identify which object in a system is denying a user access.

You should make an audit plan before you implement any auditing. In the auditing plan you should identify resources, users, and activities that you want to track. You can implement auditing for several scenarios, such as:

- Tracking changes to user and machine attributes. As with files, users and machine objects can have attributes, and changes to these can affect whether users can access files. Therefore it can be valuable to track changes to user or machine attributes. Users and machine objects live in AD and therefore changes to their attributes can be tracked using Directory Service Access Auditing.
- Get more information from user logon events. In Windows Server 2012, user logon event (4624) contains information about the attributes of the file that was accessed. You can take advantage of this additional information by using audit log management tools to correlate user logon events with object access events, and enabling event filtering based on both file attributes and user attributes.
- Provide more information from object access auditing. In Windows Server 2008 R2 and Windows Server 2012 File Access events (4656, 4663) now contain information about the attributes of the file that was accessed. This additional information can be used by event log filtering tools to help you identify the most relevant audit events.
- Track changes to Central Access Policies, Central Access Rules and Claims. These objects define the central policy that you can use to control access to critical resources. Tracking changes to these could be important for the organization. Since all of these objects are stored in AD DS you can audit them just as any other securable object in Active Directory by using the Directory Service Access Auditing.
- Tracking changes to file attributes. File attributes determine which Central Access Policy applies to the file. A change to the file attributes can potentially affect the access restrictions on the file. You can track changes to file attributes on any machine by configuring Authorization Policy Change auditing and Object Access auditing for File Systems. Event 4911 has been introduced to differentiate this event from other Authorization policy change events.

## Planning Access Denied Assistance

Access Denied Assistance helps end users to determine the reason why they cannot access a resource. It also helps IT staff to properly diagnose a problem and properly direct the resolution. Windows Server 2012 enables you to customize messages about access denied as well as to provide users with ability to request access without contacting help desk or IT team. In combination with DAC, Access Denied Assistance can inform the file administrator of the user and resource claims, enabling him to make educated decisions to adjust policies or fix user attributes (e.g. if department is written as HR instead of Human Resources).

When planning for Access Denied Assistance, consider:

- Message that users see
- Text of email that users use to request access
- Recipients for request access emails
- Target operating systems

When planning for Access Denied Assistance, you should include the following:

- Plan for message that users see when they try to access resource where they do not have access rights. It is important that the message is informal and easy to understand.
- Create the email text that users use to request access. If you allow users to request access for resources, you can prepare text that is added to the end of their email message.
- Determine the recipients for access request email messages. You can choose that email is sent to folder owners, file server administrators, or any other specified recipient. It is important that messages are always directed to the proper person. If you have a help desk tool or monitoring solution which allows emails, you can also direct those emails to automatically generate user requests in your helpdesk solution.
- Plan the target operating systems. Access Denied Assistance only works with Windows 8 or Windows Server 2012.

## Planning Policy Changes

After you implement a Dynamic Access Control infrastructure you might have to implement changes. For example, you might have to change some conditional expression, or you might want to change claims. You must carefully plan any change to Dynamic Access Control components.

Windows Server 2012 enables you to stage policy changes. A change to Central Access Policy can severely affect access control. For example, a change could potentially grant more access than desired, or, an overly restrictive change in policy could generate an excessive number of helpdesk calls. It is therefore important to test changes before implementation. For this purpose, Windows Server 2012 introduces the concept of staging. Staging enables users to verify their proposed policy changes before enforcing them. To use policy staging, proposed policies are deployed along with the enforced policies but do not actually grant or deny permissions. Instead Windows logs an audit event (4818) any time the result of the access check using the staged policy is different from the result of an access check using the enforced policy.

- Dynamic Access Control enables you to stage changes
- Staging is implemented with Proposed Permissions
- Permissions are compared to current permissions
- Every attempt to access resource is logged in Security log of file server

## Lesson 3

# Implementing and Configuring Dynamic Access Control

To implement and configure Dynamic Access Control you must perform several steps and configure several objects. In this lesson, you will learn about implementing and configuring Dynamic Access Control.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe Prerequisites for Implementing Dynamic Access Control.
- Enable Support in AD DS for Dynamic Access Control.
- Implement claims and resource property objects.
- Implement Central Access Policy.
- Implement File Access Auditing.
- Implement Access Denied Assistance.
- Implement File Classifications.
- Implement Dynamic Access Control.

### Prerequisites for Implementing Dynamic Access Control

Because Dynamic Access Control is a new technology in Windows Server 2012, you must ensure that certain prerequisites are fulfilled before implementation.

To implement claims-based authorization for resource access, you must implement the following:

- Windows Server 2012 installed on the file server that hosts the resources being protected with Dynamic Access Control. The file server hosting the share must be a Windows Server 2012 file server to read claims and device authorization data from a Kerberos ticket, translate those SIDs and claims from the ticket into an authentication token, and compare the authorization data in the token against conditional expressions in the security descriptor.
- At least one Windows Server 2012 domain controller accessible by the Windows client computer in the user's domain. The new authorization and auditing mechanism requires extensions to AD DS. These new extensions build the Windows claim dictionary, which is where Windows stores claims for an Active Directory forest. Claims authorization also relies on the Kerberos Key Distribution Center (KDC). The Windows Server 2012 KDC contains Kerberos enhancements required to transport claims within a Kerberos ticket and Compound Identity. Windows Server 2012 KDC also includes an enhancement to support Kerberos armoring. Kerberos armoring is an implementation of Flexible Authentication Secure Tunneling (FAST). It provides a protected channel between the **LSA, Netlogon** KDC.
- Windows Server 2012 domain controllers in each domain when using claims across a forest trust.

Dynamic Access Control is a technology specific to Windows Server 2012

To deploy Dynamic Access Control, you must have these technologies:

- Domain controller running on Windows Server 2012
- File server running Windows Server 2012
- Windows 8 desktop (for device claims)

- Windows 8 client (required when using device claims). Older desktop operating systems do not support device claims.

Although Windows Server 2012 domain controller is required, there is no requirement for having a Windows Server 2012 domain and forest functional level, unless you want to use claims across forest trust. This means that you can also have domain controllers on Windows Server 2008 and Windows Server 2008 R2 with forest functional level on Windows Server 2008.



**Note:** Implementing Dynamic Access Control in a multiple forest scenario has additional setup requirements.

## Enabling Support in AD DS for Dynamic Access Control

After fulfilling software requirements for enabling Dynamic Access Control support, you must enable claim support for the Windows Server 2012 KDC. Kerberos support for Dynamic Access Control provides a mechanism for including user claim and device authorization information in a Windows authentication token. Access checks on resources, such as files and folders, use this authorization information to verify identity.

You should first use Group Policy to enable AD DS for Dynamic Access Control. Because this setting is specific to domain controllers, you can create a new Group Policy object (GPO) and link it to Domain Controllers Organizational Unit (OU), or by editing Default Domain Controllers GPO that is already linked to that OU.

Whichever method you choose you should open Group Policy Object Editor and navigate to Computer Configuration\Policies\Administrative Templates\System\KDC. In this node, open a setting called **Support Dynamic Access Control and Kerberos armoring**.

You can configure this policy setting by choosing one of the four listed options:

- Do not support Dynamic Access Control and Kerberos armoring
- Support Dynamic Access Control and Kerberos armoring
- Always provide claims and FAST RFC behavior
- Also fail unarmored authentication requests

Claims and Kerberos armoring support are disabled by default, which is the same as if this policy setting is not configured, or configured as **Do not support Dynamic Access Control and Kerberos armoring**.

The policy setting **Support Dynamic Access Control and Kerberos armoring** configures Dynamic Access Control and Kerberos armoring in a mix-mode environment, when there is a mixture of Windows Server 2012 domain controllers and domain controllers running earlier versions of Windows Server.

You use the remaining policy settings when all the domain controllers are Windows Server 2012 domain controllers and the domain functional level is configured to Windows Server 2012. The **Always provide claims and FAST RFC behavior policy** setting and the **Also fail unarmored authentication requests** policy setting enable Dynamic Access Control and Kerberos armoring for the domain. However, the latter policy setting requires all Kerberos Authentication Service (AS) and Ticket-Granting Service (TGS) communication to use Kerberos armoring.

- Dynamic Access Control support in AD DS is enabled by using Group Policy
- GPO that contain Dynamic Access Control Settings must be linked to the OU of the domain controller
- Dynamic Access Control setting is available in Computer Configuration\Policies\Administrative Templates\System\KDC node in GPO Object Editor
- Settings Support Dynamic Access Control can be configured as:
  - Do not support Dynamic Access Control and Kerberos armoring
  - Support Dynamic Access Control and Kerberos armoring
  - Always provide claims and FAST RFC behavior
  - Also fail unarmored authentication requests

Windows Server 2012 domain controllers read this configuration while other domain controllers ignore this setting.

## Implementing Claims and Resource Property Objects

After you enable support for Dynamic Access Control in AD DS, you next create and configure claims and resource property objects.

### Creating and Configuring Claim Types

The primary method to create and configure claims is to use the Active Directory Administrative Center (ADAC) console. You use ADAC to create attribute-based claims, which are the most common. However, you can also use Active Directory Module for Windows PowerShell® to create certificate-based claims. All claims are stored in the configuration partition of AD DS.

Because this partition is forest wide, all domains within that forest share the claim dictionary, and domain controllers from those respective domain issue claim information during user and computer authentication.

If you want to create attribute based claims in ADAC, you should navigate to the Dynamic Access Control node, and then open the Claim Types container. By default, no claim types are defined here.

In the **Actions** pane, when you click **Create Claim Type**, you see the list of attributes. These attributes (for user or computer objects) are used to source values for claims. When you create a claim, you associate the claim to the specific attribute. The value of that attribute is populated as a claim value. It is therefore important that information contained in Active Directory attributes that are used to source claim types contain accurate information, or remain blank.

When you select the attribute that you want to use to create a claim, you also must provide a name for the claim. The suggested name for the claim is always the same as selected attribute name. However, you can also provide an alternate or more meaningful name for the claim. Optionally, you can also provide suggested values for a claim. This is not mandatory, but if you do it, you can reduce the possibility for making mistakes.



**Note:** Claim types are sourced from AD DS attributes. That is why you must configure attributes for your computer and user accounts in AD DS with the information that is correct for the respective user or computer. Windows Server 2012 domain controllers do not issue a claim for an attribute-based claim type when the attribute for the authenticating principal is empty. Depending on the configuration of the data file's Resource Property Object attributes, a null value in a claim may result in the user being denied access to DAC-protected data.

- Claims:
  - Created for users and computers
  - Have attributes as source
  - Created in AD AC or Windows PowerShell
- Resource Property Objects:
  - Created for resources
  - Have properties as a source
  - Create in AD AC or Windows PowerShell
- Both Claims and RPOs are used in conditional expressions

### Creating and Configuring Resource Properties

Although evaluating resource properties is the very core of Dynamic Access Control, you should implement it after user and device claims have been defined. Keep in mind that if a claim does not match the specified resource property value, then access to the data is denied. To reverse the order of implementation, then, would risk inadvertently blocking users from data that they otherwise should be able to access. When you use claims to control access to files and folders, you must also provide additional information on these resources. You do this by configuring Resource Property objects. You

manage Resource Property objects in the Resource Properties container in the Dynamic Access Control node in ADAC. You can create your own resource properties or you can use one of preconfigured properties, such as Country, Department, Folder Usage, etc. All predefined Resource Property objects are disabled by default. If you want to use any of them, you should first enable it. If you want to create your own Resource Property object, you can specify the property type and allowed or suggested values.

When you create Resource Property objects you can select properties to include on the files and folders. Windows uses the values in these properties with the values from user and device claims when evaluating file authorization and auditing.

After you have configured user and device claims and resource properties, you must then protect the file and folders using conditional expressions that evaluate user and device claims against values within resource properties, or constant values. You can do this in two ways. If you want to focus on specific folders, you can use the advanced security settings editor to create conditional expressions directly in the security descriptor. Alternatively, to cover several (or all) file servers, you can create Central Policy rules and link those rules to Central Policy objects. You can then deploy Central Policy objects to file servers using Group Policy and configure the share to use the Central Policy object. Using Central Access Policies is the most efficient and preferred method for securing files and folders. It is discussed in the next topic. If you want to cover certain files with a common set of properties across various folders or files, you can also use file classification.

You can use claim and resource property objects together in conditional expressions. Windows Server 2012 and Windows 8 support one or more conditional expressions within a permission entry. Conditional expressions simply add another applicable layer to the permission entry. The results of all conditional expressions must evaluate to true for Windows to grant the permission entry for authorization. For example, if you define claim Department for a user (with a source attribute department), and defined resource property object called Dept, you can define conditional expression that says: User can access a folder (with applied resource property objects) only if user attribute department value is equal to value of property Dept on the folder. Note, however, that if the resource property of Dept has not been applied to the file(s) in question, or if Dept is a null value, then the user will be granted access to the data. To be clear – access is controlled not by the claim, but by the Resource Object. The claim must provide the correct value corresponding to the requirements set by the Resource Object. If the Resource Object does not involve a particular attribute, then additional or extra claim attributes associated with the user or device are ignored.

## Implementing Central Access Rules and Policy

Central Access Policy enables you manage and deploy consistent authorization throughout the enterprise through Central Access Rules and Central Access Policy objects.

Central Access Policy helps act as a security net that an organization applies across its servers. You use Group Policy to deploy Central Access Policy, and you apply Central Access Policy to all file servers that will use Dynamic Access Control. Central Access Policy is not mandatory for using Dynamic Access Control. It just enables you to deploy a consistent configuration to several file servers.

Central Access enables you manage and deploy consistent authorization throughout the enterprise

Central Access Policy main component is Central Access Rule

Central Access Rule specify:

- Targeted resource
- Permissions
- Conditions



The main component of Central Access Policy is Central Access Rule. In fact, Central Access Policy objects represent a collection of Central Access Rule objects that you apply to Windows Server 2012 file servers using Group Policy. You should create a Central Access Rule before you create Central Access Policy because a Central Access Rule contains multiple criteria that Windows uses when evaluating access. A Central Access Rule can use conditional expressions to target specific files and folders. Each Central Access Rule has multiple permission entry lists that you use to manage the rule's current permission entries, or proposed permission entries, or return the rule's current permission entry list to its last known list of permission entries. Each Central Access Rule can be a member of one or more Central Access Policy objects.

## Configuring Central Access Rules

You typically create and configure Central Access Rules in Active Directory Administrative Center. However, you can also use PowerShell to do the same thing.

When you start to create a new Central Access rule, you must first provide a name and description for the rule. You can also choose to protect the rule against accidental deletion.

Next, you configure Target Resources. You use the Target Resource section to create a scope of applicability for the access rule. You create the scope by using resource properties within one or more conditional expressions. To make it simple, you can keep the default value (All resources), but usually you apply some resource filtering. You can join these conditional expressions using logical operators, such as AND and OR. Additionally, you can group conditional expressions together to combine the result of two or more joined conditional expression. The Targeted Resource box displays the currently configured conditional expression that is used to control the rule's applicability.

Finally, you configure permissions. There are two choices for permissions:

- **Use following permissions as proposed permissions**

Use this option to add the permission entries in the permission list to the list of proposed permission entries for the newly created Central Access Rule. You use the proposed permission list combined with file system auditing, to model the effective access users have to the resource without changing the permission entries in the current permissions list. Proposed permissions write a special audit event to the event log that describes the proposed effective access for the user.

- **Use following permissions as current permissions**

Use this option to add the permission entries in the permission list to the list of current permissions entries for the newly created Central Access Rule. The current permissions list represents the additional permissions Windows considers when the Central Access Rule is deployed to a file server. Central Access Rules do not replace the existing security. When making authorization decisions, Windows evaluates permission entries from Central Access Rule's current permissions list, NTFS, and share permissions lists.

## Implementing File Access Auditing

The Global Object Access Auditing feature in Windows 8 and Windows Server 2012 enables you to configure object access auditing for every file and folder in the file system on the computer. You use this policy setting to centrally manage and configure Windows to monitor every file and folder on the computer. To enable object access auditing in previous versions of Windows Server, you had to configure this option in basic audit policies (in GPOs), and also turn on auditing for a specific security principal in the System Access Control List (SACL) of the object. Sometimes this approach did not easily reconcile with company policies such as "Log all administrative write activity on servers containing Finance information," because you cannot turn on object access audit logging on the server level but only on the object level.

- Global Object Access Auditing centrally manages and configures Windows to monitor every file and folder on the server
- Can be integrated with Dynamic Access Control
- New Audit policy categories in Group Policy

The new audit category in Windows Server 2008 R2 and Windows Server 2012 enables administrators to manage object access auditing using a much wider scope.


Dynamic Access Control enables you to create targeted audit policies using expressions based on user, computer and resource claims. For example, you could create an audit policy to track all Read and Write operations on files classified as High Confidential by employees who do not have a High Security Clearance attribute populated with the appropriate value. You can author expression-based audit policies directly on a file or folder or centrally via Group Policy using Global Object Access Auditing. By using this approach you do not prevent unauthorized access, but register attempts to access the content by unauthorized people.

Global Object Access Auditing includes the File system and registry subcategory.

You configure Global Object Access Auditing when you enable Object Access auditing and Global Object Access Auditing. Enabling Object Auditing turns on auditing for the computer that applies the policy setting. However, enabling auditing alone does not always generate auditing events. The resource, in this instance files and folders, must contain audit entries.

We recommend configuring Global Object Access Auditing for the enterprise by using the security policy of a domain-based GPO. The two security policy settings required to enable Global Object Access Auditing are located at these locations:

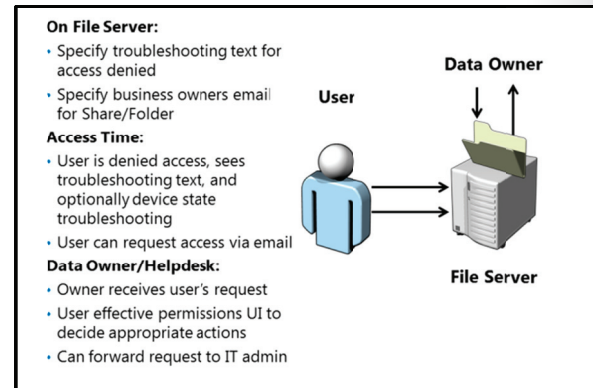
- Computer Configuration\Windows Settings\Security Settings\Advanced Audit Policy\Audit Policies\Object Access\Audit File System
- Computer Configuration\Windows Settings\Security Settings\Advanced Audit Policy\Audit Policy\Global Object Access Auditing\File System

 **Note:** If both a file or folder SACL and a Global Object Access Auditing policy (or a single registry setting SACL and a Global Object Access Auditing policy) are configured on a computer, the effective SACL is derived from combining the file or folder SACL and the Global Object Access Auditing policy. This means that an audit event is generated if an activity matches either the file or folder SACL or the Global Object Access Auditing policy.



## Implementing Access Denied Assistance

One of the most common errors that users receive when they try to access a file or folder on a remote file server is an access denied error. Usually, this error occurs when a user tries to access resource without having proper permission or because of incorrectly configured permissions or resource access control list (ACL). If you are using Dynamic Access Control, things can be even more complicated. Users, who might have permissions, but for example a relevant attribute in their account is misspelled, will not be granted access.



When users receive this kind of error, they usually try to contact the administrator to obtain access. However, administrators usually do not approve access to resources, so users are then redirected to someone else for approval.

In Windows Server 2012 there is a new technology to help both users and administrators in such situations. This technology is called Access Denied Assistance. It helps users respond to access denied issues without involving IT staff by providing information about the problem and directing users to the proper person.

### Access-denied Remediation

The Access Denied Assistance technology in Windows Server 2012 provides three ways for troubleshooting issues with access denied errors:

- **Self-remediation.** Windows Server 2012 provides a way to create customized access-denied messages that are authored by the server administrator. By using the information in these messages, users can try to self-remediate access-denied cases. For example, the user may be directed to first map to a computer using a particular drive letter. The message can also include URLs to direct the users to self-remediation websites that are provided by the organization. For example, the URL might direct the user to change their password to an application or download a refreshed copy of client-side software.
- **Remediation by the data owner.** In Windows Server 2012, administrators can define owners for shared folders. This enables users to send an email to the data owners to request access. . For example, if the user was accidentally left off a security group membership, the data owner may be able to add the user to the group. If the data owner does not know how to help the user get access, he or she can forward this information to the appropriate IT administrator. This is helpful because the number of user support requests escalated to the support desk should be limited to special, difficult-to-resolve cases.
- **Remediation by Help Desk and file server administrators.** If the user cannot self-remediate the issue or the data owner cannot help, Windows Server 2012 provides a user interface where administrators can view the effective permission for users for a file or folder so that it is easier to troubleshoot access issues. An example of when an administrator should be involved are cases where attributes – either claims and/or resource objects – have been incorrectly defined or contain incorrect information, or when the data itself seems to be corrupted.

You enable Access Denied Assistance by using group policy. You open Group Policy Object editor and navigate to Computer Configuration\Policies\Administrative Templates\System\Access-Denied Assistance. In this node, you can enable Access Denied Assistance, and also, you can provide customized messages for users. Alternatively, you can also use File Server Resource Manager console to enable access-denied assistance. However, if this feature is enabled in Group Policy, the appropriate settings in File Server Resource Manager console are disabled for configuration.

## Implementing File Classifications

To effectively implement Dynamic Access Control technology, you must have well-defined claims and resource properties. Although claims are defined by attributes for user or a device, resource properties are most often manually created and defined. File Classifications enable administrators to define automatic procedures for defining a desired property on the file, based on condition specified in classification rule. For example, you can set the property Confidentiality to High on all documents whose content contains the word "secret." You can then use this property in

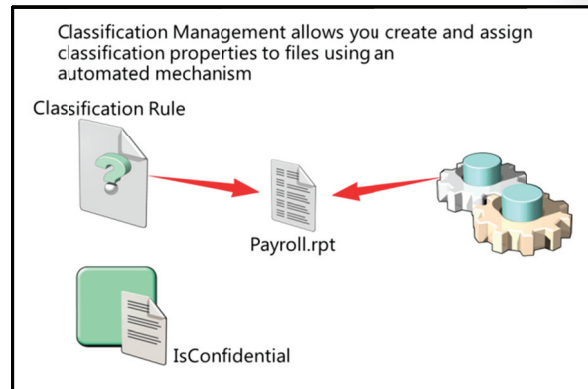
Dynamic Access Control to specify, for example, that only employees with attribute employee type set to Manager can access those documents that are classified with high confidentiality.

In Windows Server 2008 R2 and Windows Server 2012, Classification Management and File Management tasks enable administrators to manage groups of files based on various file and folder attributes. With Classification Management and File Management tasks, you can automate file and folder maintenance tasks, such as cleaning up stale data or protecting sensitive information.

Classification Management is designed to ease the burden and management of data that is spread out in the organization. Files can be classified in a variety of ways. In most scenarios, classification is performed manually. The File Classification infrastructure in Windows Server 2008 R2 enables organizations to convert these manual processes into automated policies. Administrators can specify file management policies based on a file's classification and apply corporate requirements for managing data based on business value.

You can use file classification to perform the following actions:

1. Define classification properties and values, which can be assigned to files by running classification rules.
2. Create, update, and run classification rules. Each rule assigns a single predefined property and value to files within a specified directory based on installed classification plug-ins.
3. When running a classification rule, reevaluate files that are already classified. You can choose to overwrite existing classification values or add the value to properties that support multiple values. You can also use this to de-classify files that are not in classification criteria anymore.



## Demonstration: Implementing Central Access Rules and Policies

### Demonstration Steps

1. In the Active Directory Administrative Center, create claims for department and employee type attributes.
2. Enable Resource Type for department.
3. Create Central Access rule to enable members of IT group to access resources if user department attribute matches resource department.
4. Create a Central Access Policy.

## Lab: Implementing Dynamic Access Control

### Scenario

The Research team at A. Datum performs some highly confidential work that provides much value to the business. Managers and Research departments at A. Datum frequently store files that contain business-critical information on the company file servers. The security department wants to ensure that these confidential files are only accessible to suitably authorized personnel and that all access to these files be audited.

As one of the senior network administrators at A. Datum, you are responsible for addressing these security requirements by implementing Dynamic Access Control on the file servers. You plan to work closely with the business groups and the security department in identifying which files must be secured, and who should have access to these files. Then you plan to implement Dynamic Access Control based on the company requirements.

### Objectives

- Plan Dynamic Access Control Deployment and prepare AD DS for Dynamic Access Control.
- Configure user and device claims.
- Configure resource properties and file classifications.
- Configure central access rules and policies.
- Configure and validate access remediation.

### Lab Setup

Estimated time: **90 minutes**

Virtual machines	20417A-LON-DC1 20417A-LON-SVR1 20417A-LON-CL1 20417A-LON-CL2
User Name	Adatum\Administrator
Password	Pa\$\$w0rd

For this lab, you will use the available virtual machine environment. Before you begin the lab, you must complete the following steps:

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, click **20417A-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on using the following credentials:
  - a. User name: **Adatum\Administrator**
  - b. Password: **Pa\$\$w0rd**
5. Repeat steps 2 and 3 for **20417A-LON-SVR1**, **20417A-LON-CL1** and **20417A-LON-CL2**.
6. Log on to LON-SVR1 as **Adatum\Administrator** with the password of **Pa\$\$w0rd**. Do not log on to LON-CL1 or LON-CL2 until instructed to do so.

## Exercise 1: Planning the Dynamic Access Control Implementation and Preparing AD DS for Dynamic Access Control

### Scenario

A. Datum must ensure that documents used by the Research department and managers are secured. Most of the files used by these departments are stored in shared folders dedicated to these departments, but sometimes confidential documents appear in other shared folders. Folders that belong to Research department should be accessed and modified only by members of Research department. Also, documents that are classified as highly confidential should only be accessed by Managers. The security department is also concerned that users in the Managers department are accessing the files using their home computers, which may not be highly secure. You must create a plan for securing the documents regardless of where they are located and ensure that the documents can only be accessed from authorized computers. Authorized computers for Managers are members of the security group ManagersWks.

The support department reports that a high number of calls are generated by users who cannot access resources. You must implement a technology that helps users to better understand error messages as well as enable them to automatically request access.

First, you will plan for Dynamic Access Control deployment. Then you must prepare your AD DS to support Dynamic Access Control.

The main tasks for this exercise are as follows:

1. Plan the Dynamic Access Control Deployment Based on the Security and Business Requirements.
2. Prepare AD DS to support Dynamic Access Control.

### ► Task 1: Plan the Dynamic Access Control Deployment Based on the Security and Business Requirements

- Describe how you will design Dynamic Access Control to fulfill requirements for access control, described in the scenario.

### ► Task 2: Prepare AD DS to support Dynamic Access Control

1. On the LON-DC1, from Server Manager open Active Directory Users and Computers.
2. Make new organizational unit named **Test**.
3. Move LON-CL1, LON-CL2 and LON-SVR1 computer objects into Test OU.
4. On LON-DC1, from Server Manager, open the **Group Policy Management** console.
5. Remove the Block Inheritance setting applied to the Managers OU. (This setting has been applied and used in a later module of the course.)
6. Edit the Default Domain Controllers Policy GPO.
7. In the Group Policy Management Editor, under **Computer Configuration**, expand **Policies**, expand **Administrative Templates**, expand **System**, and then click **KDC**.
8. Enable the KDC support for claims, compound authentication and Kerberos armoring policy setting.
9. Select **Supported** in **Options** section.
10. On LON-DC1, refresh Group Policy.
11. Open Active Directory Users and Computers and create a security group called **ManagersWKS** in **Users** container.

12. Add LON-CL1 to ManagersWKS group.
13. Verify that user **Aidan Delaney** is a member of **Managers** department and **Allie Bellew** is the member of the **Research** department.

**Results:** After completing this exercise you will have design for Dynamic Access Control and you will have prepared AD DS for Dynamic Access Control implementation.

## Exercise 2: Configuring User and Device Claims

### Scenario

The first step in implementing Dynamic Access Control is to configure the claims for the users and devices that access the files. In this exercise, you will review the default claims and create new claims based on the department and computer description attributes. For users, you will create a claim for department attribute. For computers, you will create claim for description attribute.

The main tasks for this exercise are as follows:

1. Review the Default Claim Types.
2. Configure Claims for Users.
3. Configure Claims for Devices.

#### ► Task 1: Review the Default Claim Types

1. On LON-DC1, in Server Manager, open the Active Directory Administrative Center.
2. Click the **Dynamic Access Control** node in Active Directory Administrative Center.
3. Open the **Claim Types** container and verify that there is no default claims defined.
4. Open the **Resource Properties** container and note that all properties are disabled by default.
5. Open **Resource Property Lists** container and then open the properties of the **Global Resource Property List**.
6. In the **Resource Properties** section review available resource properties.
7. Click **Cancel**.

#### ► Task 2: Configure Claims for Users

1. In the **Active Directory Administrative Center**, in the navigation pane click **Dynamic Access Control**.
2. Open the **Claim Types** container, and create a new claim type for users and computers using the following settings:
  - Source Attribute: **Department**
  - Display name: **Company Department**

### ► Task 3: Configure Claims for Devices

1. In the **Active Directory Administrative Center**, in the Tasks pane click **New** and select **Claim Type**.
2. Create a new claim type for computers using the following settings:
  - Source Attribute: description
  - Display name: description

**Results:** After completing this exercise you will have configured user and device claims.

## Exercise 3: Configuring Resource Properties and File Classifications

### Scenario

The second step in implementing Dynamic Access Control is to configure the resource property lists and resource property definitions. After you do this, you should make a new classification rule that classify all files that contain the word secret in the body. These files should be assigned a value of High for attribute Confidentiality. You should also assign department property to the folder that belongs to Research department.

The main tasks for this exercise are as follows:

1. Configure Resource Property Definitions.
2. Classify files.
3. Assign properties to folder.

### ► Task 1: Configure Resource Property Definitions

1. In the **Active Directory Administrative Center**, click **Dynamic Access Control** and then open the Resource Properties container.
2. Enable the **Department** and **Confidentiality** Resource Properties.
3. Open Properties for **Department** property.
4. Add **Research** as suggested value.
5. Open the **Global Resource Property List** and make sure that **Department** and **Confidentiality** are included in the list.
6. Click **Cancel**.
7. Close the Active Directory Administrative Center.

### ► Task 2: Classify files

1. On LON-SVR1, in Server Manager, add the **File Server Resource Manager**.
2. Open File Server Resource Manager.
3. Refresh **Classification Properties**. Verify that **Confidentiality** and **Department** properties are in the list.
4. Create a Classification rule with following values:
  - Name: **Set Confidentiality**
  - Scope: **C:\Docs**
  - Classification method: **Content Classifier**

- Property: **Confidentiality**
  - Value: **High**
  - Classification Parameters: String "secret"
  - Select **Re-evaluate existing property values**, and then click **Overwrite the existing value**.
5. Run the classification rule.
  6. Open Windows Explorer and open Properties for files Doc1.txt, Doc2.txt and Doc3.txt in **C:\Docs** folder.
  7. Verify values for Confidentiality. Doc1.txt and Doc2.txt should have confidentiality set to High.

### ► Task 3: Assign properties to folder

1. On LON-SVR1 open Windows Explorer.
2. Browse to **C:\Research** and open its properties.
3. On the **Classification** tab, set the **Department** value to **Research**.

**Results:** After this exercise, you will have configured resource properties and file classifications.

## Exercise 4: Configuring Central Access Rules and Policies

### Scenario

Now that you have configured claims, resource properties, and file classifications, you want to create and configure central access rules and policies.

The main tasks for this exercise are as follows:

1. Configure Central Access Policy Rules.
2. Create Central Access Policy.
3. Publish Central Access Policy with Group Policy.
4. Apply Central Access Policy to resources.
5. Configure access denied remediation settings.

### ► Task 1: Configure Central Access Policy Rules

1. On LON-DC1, in Server Manager, click **Tools** and then click **Active Directory Administrative Center**.
2. Click **Dynamic Access Control** and then open the Central Access Rules container.
3. Create a new Central Access Rule with following values :
  - Name: **Department Match**
  - Target Resource: use condition **Resource-Department-Equals-Value-Research**
  - Permissions: Remove Administrators, and then add Authenticated Users, Modify, with condition User-Company Department-Equals-Resource-Department



4. Create another Central Access Rule with following values :
  - o Name: **Access Confidential Docs**
  - o Target Resource: use condition **Resource-Confidentiality-Equals-Value-High**
  - o Permissions:  
Set first condition to be: User-Group-Member of each-Value-Managers  
Set second condition to be: Device-Group-Member of each-Value-ManagersWKS

► **Task 2: Create Central Access Policy**

1. On LON-DC1 in Active Directory Administrative Center, create a new **Central Access Policy** with following values:
  - o Name: **Protect confidential docs**
  - o Rules included: **Access Confidential Docs**
2. Create another Central Access Policy with following values:
  - o Name: **Department Match**
  - o Rules included: **Department Match**
3. Close the Active Directory Administrative Center.

► **Task 3: Publish Central Access Policy with Group Policy**

1. On LON-DC1, from the Server Manager, open the **Group Policy Management** console.
2. Create new GPO named **DAC Policy** and link it to organizational unit **Test**.
3. Edit the DAC Policy and browse to **Computer Configuration/Policies/Windows Settings /Security Settings/File System**, and then right-click **Central Access Policy**.
4. Click Manage Central Access Policies.
5. Click both **Department Match** and **Protect confidential docs**, and then click **Add**. Click **OK**.
6. Close the Group Policy Management Editor and the Group Policy Management console.

► **Task 4: Apply Central Access Policy to resources**

1. On LON-SVR1, start Windows PowerShell.
2. Refresh Group Policy on LON-SVR1.
3. Open Windows Explorer, and browse to the **C:\Docs** folder.
4. Apply the **Protect confidential docs** Central Policy to the C:\Docs folder.
5. Browse to the **C:\Research** folder.
6. Apply the **Department Match** Central Policy to the C:\Research folder.

► **Task 5: Configure access denied remediation settings**

1. On LON-DC1, open the Group Policy Management console.
2. Edit the **DAC Policy**.
3. Under Computer Configuration node, expand **Policies**, expand **Administrative Templates**, expand **System**, and then click **Access-Denied Assistance**.
4. In the right pane double-click **Customize Message for Access Denied errors**.
5. In the Customize Message for Access Denied errors window click **Enabled**.

6. In the Display the following message to users who are denied access text box type: **You are denied access because of permission policy. Please request access.**
7. Select check box **Enable users to request assistance.** Click **OK.**
8. Double-click **Enable access-denied assistance on client** for all file types and enable it.
9. Click **OK** and close the Group Policy Management Editor and the Group Policy Management console.
10. Switch to LON-SVR1, and refresh Group Policy.

**Results:** After completing this exercise you will have configured central access rules and policies.

## Exercise 5: Validating and Remediating Access Control

### Scenario

To ensure that the Dynamic Access Control settings are configured correctly, you plan to test various scenarios for users to access the files. You plan to try both approved users and devices and unapproved users and devices. You also plan to validate the access remediation configuration.

The main tasks for this exercise are as follows:

1. Verify Dynamic Access Control functionality.
2. Configure staging for Dynamic Access Policy.
3. Configure staging permissions.
4. Verify staging.
5. Use effective permissions to test Dynamic Access Control.
6. To prepare for next module.

### ► Task 1: Verify Dynamic Access Control functionality

1. Log on to LON-CL1 as **Adatum\April** with password **Pa\$\$w0rd.**
2. Click the **Desktop** tile and then open Windows Explorer.
3. Browse to **\\LON-SVR1\Docs**. Verify that you can only open Doc3.
4. Try to access **\\LON-SVR1\Research**. You should be unable to access it.
5. Log off of LON-CL1.
6. Log on to LON-CL1 **Adatum\Allie** with the password of **Pa\$\$w0rd.**
7. Open Windows Explorer and try to access **\\LON-SVR1\Research**.



**Note:** You should be able to access it as well as open files in it.

8. Log off of LON-CL1.
9. Log on to LON-CL1 as **Adatum\Aidan** with the password of **Pa\$\$w0rd.**
10. Open Windows Explorer and try to access **\\LON-SVR1\Docs**.



**Note:** You should be able to open all files in this folder.

11. Log off of LON-CL1.
12. Log on to LON-CL2, as **Adatum\Aidan** with the password of **Pa\$\$w0rd**.
13. Open Windows Explorer and try to access **\\LON-SVR1\Docs**.



**Note:** You should be unable to see Doc1 and Doc2 since LON-CL2 is not permitted to view secret documents.

#### ► Task 2: Configure staging for Dynamic Access Policy

1. On LON-DC1, open Group Policy Management.
2. Edit the **DAC Policy** GPO.
3. In the Group Policy Management Editor, browse to Computer Configuration/Policies /Windows Settings/Security Settings/Advanced Audit Policy Configuration/Audit Policies.
4. Select **Object Access**.
5. Double-click **Audit Central Access Policy Staging**. Select all three check boxes, and then click **OK**.
6. Double-click **Audit File System**. Select all three check boxes then click **OK**.
7. Close the Group Policy Management Editor and the Group Policy Management console.

#### ► Task 3: Configure staging permissions

1. On LON-DC1, open Server Manager, and then open Active Directory Administrative Center.
2. Open the Properties for the **Department Match** Central Access Rule
3. In the **Proposed permissions** section, configure a condition for Authenticated users as follows: **User-Company Department-Equals-Value-Marketing**.
4. Switch to LON-SVR1 and refresh Group Policy.

#### ► Task 4: Verify staging

1. Log on to LON-CL1 as **Adatum\Adam** with the password of **Pa\$\$w0rd**.
2. Open Windows Explorer and attempt to access **\\LON-SVR1\Research**. You will be unsuccessful. Click **Close**.
3. Switch to LON-SVR1.
4. From Server Manager, open Event Viewer and select the Security log. Look for events with Event ID 4818.

#### ► Task 5: Use effective permissions to test Dynamic Access Control

1. On LON-SVR1, open properties for **C:\Research**.
2. Open **Advanced** options for Security.
3. Click the **Effective access** tab.
4. Click **select a user**.
5. In Select User, Computer, Service Account, or Group window type **April**, and then click **Check Names**, and then click **OK**.
6. Click View effective access.
7. Review results. April should not have access to this folder.

8. Click **Include a user claim**.
9. Select **Company Department** from the drop-down list.
10. Type **Research** in **Value** text box.
11. Click **View Effective access**. April should have access now.

► **Task 6: To prepare for next module**

- When you are finished the lab, revert the virtual machines to their initial state.

**Results:** After this exercises you will have validated Dynamic Access Control functionality.

## Module Review and Takeaways

### Best Practices

- Use Central Access Policies instead of configuring conditional expressions on resources.
- Enable Access Denied Assistance settings.
- Always stage changes to Central Access Rules and Policies before implementation.
- Use file classifications to assign properties to files.

### Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Claims are not populated with appropriate values	
Conditional expression does not enable access	

### Review Questions

What is a claim?

What is the purpose of Central Access Policy?

What is Access Denied Assistance?

### Tools

Active Directory Administrative Center

**MCT USE ONLY. STUDENT USE PROHIBITED**

# Module 11

## Implementing Active Directory Domain Services

### Contents:

Module Overview	11-1
<b>Lesson 1:</b> Deploying AD DS Domain Controllers	11-2
<b>Lesson 2:</b> Configuring AD DS Domain Controllers	11-11
<b>Lesson 3:</b> Implementing Service Accounts	11-16
<b>Lesson 4:</b> Implementing Group Policy in AD DS	11-19
<b>Lesson 5:</b> Maintaining AD DS	11-28
<b>Lab:</b> Implementing AD DS	11-35
Module Review and Takeaways	11-40

## Module Overview

Active Directory® Domain Services (AD DS) is the central location for configuration information, authentication requests, and information about all the objects that are stored in an Active Directory forest. Using AD DS, you can efficiently manage users, computers, groups, printers, and other directory-enabled objects from one secure, central location. Windows PowerShell® has become the single engine for configuration and maintenance from both graphical and command-line interfaces. This module discusses deployment and configuration of domain controllers, service accounts in AD DS, Group Policy, and maintenance of AD DS.

### Objectives

After completing this module you will be able to:

- Deploy domain controllers.
- Configure domain controllers.
- Implement service accounts.
- Implement Group Policy.
- Maintain AD DS.

## Lesson 1

# Deploying AD DS Domain Controllers

To establish the Active Directory forest and the first domain in the forest, you must create at least one domain controller. In this lesson, you will learn about the new features of AD DS in Windows Server® 2012 and the various methods for deploying domain controllers.

### Lesson Objectives

After completing this lesson you will be able to:

- Describe what's new in AD DS in Windows Server 2012.
- Deploy domain controllers.
- Deploy domain controllers on a Server Core installation of Windows Server 2012.
- Deploy domain controllers using the Install From Media feature.
- Clone virtual domain controllers.
- Upgrade to AD DS in Windows Server 2012.
- Troubleshoot domain controller deployment.

### What's New in AD DS in Windows Server 2012?

Windows Server 2012 has several new features for AD DS. Windows PowerShell command-line interface is the underlying component behind installations and configurations. It enables full scripting and automation and new graphical user interfaces for previous command-line-only activities.

Some new features are described in the following table.

Windows Server 2012 includes these new features:

- New deployment methods
- Simplified administration
- Virtualized domain controllers
- Active Directory module for Windows PowerShell
- Windows PowerShell History Viewer
- Active Directory Federated Services
- Active Directory Based Activation

Feature	Improvement
Deployment	<ul style="list-style-type: none"> <li>• Server Manager now enables installation of the AD DS role on remote as well as local computers. The Active Directory Domain Services Configuration Wizard replaces Active Directory Installation Wizard (also called DCPromo).</li> <li>• Deployment now uses Windows PowerShell in the background.</li> <li>• When you install Active Directory on the member server, Windows Server 2012 performs prerequisite checks that validate domain and forest readiness.</li> </ul>
Simplified administration	<p>Improvements to configure and monitor AD DS through the Server Manager console include:</p> <ul style="list-style-type: none"> <li>• A graphical user interface for the Active Directory Recycle Bin.</li> <li>• A graphical user interface to implement fine-grained passwords.</li> </ul>



Feature	Improvement
	<ul style="list-style-type: none"> <li>Group Policy health monitoring.</li> <li>AD DS-specific performance monitoring and best practice analysis.</li> <li>Active Directory management tools, which you can open from the Server Manager console.</li> </ul>
Support for Virtualized Domain Controllers	<ul style="list-style-type: none"> <li>Improvements in the virtual environment include:</li> <li>Cloning domain controllers is now a supported option to enable automated deployment and rollback protection</li> <li>Restoration of domain controller snapshots does not disrupt the AD DS environment.</li> </ul>
Active Directory Module for Windows PowerShell	The Active Directory module has new cmdlets for replication topology management, Dynamic Access Control, and other operations. It is no longer necessary to use Active Directory Installation Wizard (also called DCPromo) to create a domain controller. When you use Windows PowerShell to install AD DS, Active Directory Installation Wizard functionality is now included in the cmdlet.
Windows PowerShell History Viewer	When administrators use the Active Directory Administrative Center, they can now view the underlying Windows PowerShell commands that are executed. This helps reduce the time required to learn the Windows PowerShell commands.
Active Directory Federated Services (AD FS)	AD FS is now included as a server role with Windows Server 2012. This version provides a less complex trust setup and management process, an ability to extend the claims attribute store and a broader scope for defining claims. AD FS services are frequently required for hybrid cloud deployments.
Active Directory Based Activation (AD BA)	Key Management Servers (KMS) are no longer required to activate computers running Windows Server 2012 and Windows® 8. Activating the initial customer-specific volume license key (CSVLK) requires a one-time contact with Microsoft activation over the Internet.

## Deploying AD DS Domain Controllers

With Windows Server 2008, you could deploy a domain controller by installing the AD DS role to add the binary files and then using Active Directory Installation Wizard to install AD DS. In Windows Server 2012 you deploy a domain controller by using Server Manager to add the AD DS role. You use a separate wizard to configure AD DS within Server Manager.

You can add the AD DS role binaries using these four methods:

- The graphical Server Manager.
- The Server Manager module.

- All configuration of domain controllers can be done through a wizard in Server Manager
- AD DS binaries can be installed using Windows PowerShell
- Dism.exe is more complex to use
- The Active Directory Installation Wizard is only supported in Unattended mode



- Dism.exe.
- Active Directory Installation Wizard (also called DCPromo)

### Using Server Manager

You can use the graphical wizard in Server Manager to install the binary files and perform all the required configuration of a domain controller. The deployment wizard uses a single expanding dialog box and can do the following:

- Install AD DS remotely.
- Install DNS by default.
- Configure the domain controller as a global catalog by default.
- Display advanced mode settings.
- Prepare schema extension and domain preparation automatically in the background.



**Note:** These new features are not backward compatible with Windows Server 2008 R2 or earlier versions of Windows Server. For more information, refer to “Understand and Troubleshoot AD DS Simplified Administration in Windows Server 8 Beta.docx” from <http://www.microsoft.com/en-us/download/details.aspx?id=29019>.

### Using Windows PowerShell

You can add AD DS binaries using the Active Directory module for local or remote installations.

### Using DISM

The Deployment Image Servicing and Management (DISM) tool is part of the Windows Automated Administration Kit (WAIK). It is more complex than, and not as flexible as, Windows PowerShell. DISM is usually associated with creating deployment images for Windows Deployment Services.

### Using Active Directory Installation Wizard

Active Directory Installation Wizard (also called DCPromo) no longer has a GUI and is only supported with the Unattend option. It is no longer recommended.



**Note:** System requirements to install Windows Server 2012 are unchanged from Windows Server 2008 R2.

## Deploying AD DS Domain Controllers on Server Core

Server Core is a version of Windows Server 2012 that has no graphical interface. Server Core provides a minimal environment for running server roles. It reduces disk space usage and maintenance, and presents a smaller attack surface.

You can now install AD DS on Server Core by using Windows PowerShell for a local or remote installation. Or you can use the GUI in Server Manager on a remote system to perform the installation.

You can install AD DS:

- Locally using Windows PowerShell cmdlets
- Remotely using either Windows PowerShell cmdlets or Server Manager

## Installing the AD DS Role Locally

To Install the AD DS Role locally:

1. Install the AD DS binary files. At the local Windows PowerShell command prompt, type the cmdlet **Install-Windowsfeature -name AD-Domain-Services**, and then press Enter.
2. Configure AD DS. At the Windows PowerShell command prompt, type the cmdlet **Install-ADDSDomainController -domainname "Adatum.com"**, with other arguments as required, and then press Enter.

## Windows PowerShell Remote Installation

You can run Windows PowerShell cmdlets against remote servers. Start by installing the AD DS binary files. Then use the `invoke-command` cmdlet. For example:

**invoke-command {install-addsdomaincontroller -domainname Adatum.com -credential (get-credential) -computername NYC-DC3}**



**Note:** Guidance for using Windows PowerShell to establish a Window Server 2012 AD DS environment can be found here: [http://technet.microsoft.com/en-us/library/hh472162#BKMK\\_PSFforest](http://technet.microsoft.com/en-us/library/hh472162#BKMK_PSFforest).

## Server Manager Remote Installation

To use Server Manager to install AD DS Role remotely, perform these high-level steps:

1. Add the Server Core computer as another computer to manage.
2. Create a server group containing the Server Core computer.
3. Use the Add Roles and Features Wizard to install AD DS.
4. Complete the configuration by running the Active Directory Domain Services Configuration Wizard.

## Deploying AD DS Domain Controllers by using Install From Media (IFM)

Another method for installing AD DS is to install from an installation media created by using the `Ntdsutil.exe` utility. Installation media is created from an existing domain controller in the form of a backup. The advantage of installing from media is that it reduces the directory replication traffic required to synchronize the new domain controller. By default, a new domain controller replicates all the data for all Directory partitions that it hosts from other domain controllers. When you use IFM the new domain controller has most of the AD DS data. It only replicates updates that have occurred since the backup media was created.



Use `Ntdsutil.exe` to create the installation media

`Ntdsutil.exe` can create the following types of installation media:

- Full (or writable) domain controller
- Full (or writable) domain controller with SYSVOL data
- Read-only domain controller with SYSVOL data
- Read-only domain controller
- Create full no defrag
- Create sysvol full no defrag

## Creating the IFM media

Windows Server 2012 has two new options that enable you to create IFM media without first performing an online defrag of the exported NTDS.DIT database file. The Ntdsutil.exe can now create six types of installation media as described in the following table.

Type of installation media	Parameter	Description
Full (or writable) domain controller	Create Full <i>PathToMediaFolder</i>	Creates installation media for a writable domain controller instance in the folder that is identified in the path.
Read-only domain controller (RODC)	Create RODC <i>PathToMediaFolder</i>	Creates installation media for an RODC in the folder that is identified in the path.
Full (or writable) domain controller with SYSVOL	Create Sysvol Full <i>PathToMediaFolder</i>	Creates installation media for a writable domain controller with SYSVOL in the folder that is identified in the path.  <b>Note:</b> Does not work on Windows Server 2012
RODC with SYSVOL	Create Sysvol RODC <i>PathToMediaFolder</i>	Creates installation media for an RODC with SYSVOL in the folder that is identified in the path.  <b>Note:</b> Does not work on Windows Server 2012
Create Full NoDefrag	Create Full NoDefrag %s	Create installation media without defragmenting for a full Active Directory domain controller or an Active Directory Lightweight Directory Services (AD LDS) instance into folder %s.
Create Sysvol Full NoDefrag	Create Sysvol Full NoDefrag %s	Create installation media with SYSVOL without defragmenting for a full Active Directory domain controller or an AD/LDS instance into folder %s.

## Steps to Create IFM Media

To create IFM media, perform the following steps on an existing domain controller that is running the same operating system as the destination computer:

1. Enter the ntdsutil context. At the Windows command prompt type **NTDSUTIL**, and then press Enter.
2. At the NTDSUTIL: prompt type **Activate instance NTDS**, and then press Enter.
3. Type **IFM**.
4. At the IFM: prompt, type the command for the type of installation media you want to create. For example, to create media for a writable domain controller with SYSVOL to a folder named Media, type **Create Sysvol Full C:\Media**.

To use IFM to create additional domain controllers in the domain, you can refer to a shared folder or removable media where you store the installation media on the **Install from Media** page in the Active Directory Domain Services Installation Wizard or by using the **/ReplicationSourcePath** parameter during an unattended installation.

## Install From Media Characteristics

IFM has the following characteristics:

- Installation from media does not work across different operating system versions. You must generate media from an existing Windows Server 2012 domain controller to install AD DS on a computer running Windows Server 2012.
- When the Active Directory Recycle Bin is enabled, any installation media that was created before the Active Directory Recycle Bin was enabled is no longer valid. Create new installation media while Active Directory Recycle Bin is enabled.
- To create the IFM you must have permissions to make a backup on a domain controller.

## Deploying AD DS Read-Only Domain Controllers

The read-only domain controller (RODC) was introduced with Windows Server 2008. An RODC hosts read-only partitions of the AD DS database. This means that no AD DS change requests are made directly to the database copy stored by RODC. Instead, AD DS modifications are forwarded to RODCs through replication with a writable domain controller. All RODC AD DS replication uses a one-way, in-coming only connection from a domain controller that has a writable AD DS database copy.

RODCs provide:

- Unidirectional replication
- Credential caching
- Administrative role separation
- Read-only DNS
- RODC filtered attribute set



RODCs are primarily designed for branch office deployments where you cannot guarantee the physical security of the AD DS computers. By deploying an RODC in a branch office, you can give users a local domain controller to facilitate efficient AD DS log on and Group Policy application, even if the WAN link to the main office (where read/write domain controllers are located) is not available. A locally based RODC configured to cache passwords of local users ensures faster logons compared to logging on across a slow network connection to authenticate with a remote domain controller.

## Characteristics of RODC

RODCs have the following characteristics:

- Server Core installations support RODCs.
- An RODC cannot hold an operations master role.
- An RODC cannot be a site bridgehead server.
- RODCs only support incoming replication.
- Caching of credentials of users and computers can be explicitly enabled or denied. This can be configured in the Active Directory Configuration Wizard. By default, no user credentials are cached.
- Users can be delegated administrative rights to a specific RODC without being granted rights to AD DS. This can be configured in the Active Directory Configuration Wizard.
- RODCs support read-only Domain Name System (DNS).
- RODC can use the IFM feature for deployment.

## Preparing to Install RODC

Several prerequisites must be in place before you install and RODC. They are:

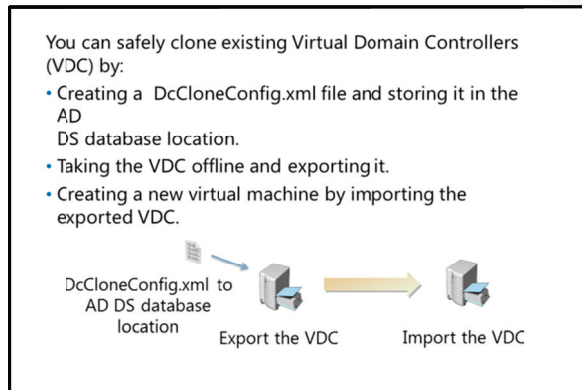
- Forest functional level must be at least 2003. The Windows Server 2012 Active Directory Configuration Wizard does not let you continue if the domain is not able to support an RODC.
- There must be a writable domain controller running Windows 2008 or later versions in the same domain.
- The domain must be prepared with the **Adprep.exe /rodcprep** command. Windows Server 2012 performs this step automatically when you install a writable domain controller.

## Installing the RODC

You can install an RODC through the Active Directory Configuration Wizard. On the **Additional Domain Controller Options** page, select the check box for **Read-only domain controller (RODC)**.

## Cloning Virtual AD DS Domain Controllers

Windows Server 2012 introduces virtualized domain controller cloning. Cloning a virtualized domain controller presents challenges. For example, two domain controllers cannot coexist in the same forest with the same name, invocation ID, and security identifier. In versions of Windows earlier than Windows Server 2012, you created virtualized domain controllers by deploying a Sysprepped base server image and manually promoting it to be a domain controller. Windows Server 2012 provides specific virtualization capabilities to AD DS Virtualized Domain Controllers (VDCs) to resolve those issues.



Windows Server 2012 VDCs have two new capabilities:

- Domain controllers can be safely cloned to deploy additional capacity and save configuration time.
- Accidental restoration of domain controller snapshots does not disrupt the AD DS environment.

## Safe Cloning

A cloned domain controller automatically syspreps (based on settings in DefaultDCCloneAllowList.xml) and promotes with the existing local AD DS data as installation media.

## Safe Backup and Restore

Rolling back to a previous snapshot of a VDC is problematic because Active Directory uses multi-master replication that relies on transactions being assigned numeric values called Update Sequence Numbers (USNs). The VDC tries to assign USNs to prior transactions that have already been assigned to valid transactions. This causes inconsistencies in the Active Directory database. Windows Server 2012 implements a process that is known as USN rollback protection. With this in place the VDC does replicate and must be forcibly demoted or manually restored non-authoritatively.

Windows Server 2012 now detects rollbacks and non-authoritatively synchronizes the delta of changes between a domain controller and its partners for AD DS and SYSVOL. You can now use snapshots without risk of permanently disabling domain controllers and requiring manually forced demotion, metadata cleanup, and re-promotion.

## Creating a VDC Clone

To create a VDC clone in Windows Server 2012, perform the following high level steps:

1. Create a DcCloneConfig.xml file that contains the unique server configuration.
2. Copy this file into the location of the AD DS database (C:\Windows\NTDS by default).
3. Take the VDC offline and export or copy it.
4. Create a new virtual machine by importing the exported one. This virtual machine is automatically promoted as a unique domain controller.



**Note:** There is no graphical interface to create the cloning xml files. However, there is a Windows PowerShell script in development for out of band release, and the XML schema is included.

## Upgrading to Windows Server 2012 AD DS

You can upgrade an existing domain controller to Windows Server 2012. You can only upgrade a domain controller created in Windows Server 2008 x64 or Windows Server 2008 R2. You cannot perform an in-place upgrade on Windows Server 2003.

To perform an in-place upgrade of a computer that has the AD DS role installed, you must first use Adprep.exe /forestprep and Adprep.exe /domainprep to prepare the forest and domain. An in-place operating system upgrade does not perform automatic schema and domain preparation. Adprep.exe is included on the installation media in the \Support\Adprep folder. There are no additional configuration steps after that point and you can continue to running the OS upgrade.

- Only domain controllers running Windows Server 2008 x64 or Windows Server 2008 R2 can be upgraded
- You cannot perform an in-place upgrade on a Windows Server 2003 domain controller
- Forestprep and Domainprep must both be run manually prior to upgrading



**Note:** We recommend a clean installation.

## Troubleshooting AD DS Domain Controller Deployments

If you encounter errors when you create a domain controller, you can use troubleshooting tools and methodologies to resolve the problem. There are also logs and utilities available.

Use a troubleshooting methodology

- Check the logs
- Use diagnostic tools
- Check for simple issues first





## Logging Options

The built-in logs are the most important tool for troubleshooting issues with domain controller promotion and demotion. There are many logs created during the installation and promotion of a domain controller, as shown in the following table.

Phase	Log
Server Manager or AD DS Deployment Windows PowerShell operations	<ul style="list-style-type: none"> <li>• %systemroot%\debug\dcpromoui.log</li> <li>• %systemroot%\debug\dcpromoui*.log</li> </ul>
Installation/Promotion of the domain controller	<ul style="list-style-type: none"> <li>• %systemroot%\debug\dcpromo.log</li> <li>• %systemroot%\debug\dcpromo*.log</li> <li>• Event viewer\Windows logs\System</li> <li>• Event viewer\Windows logs\Application</li> <li>• Event viewer\Applications and services logs\Directory Service</li> <li>• Event viewer\Applications and services logs\File Replication Service</li> <li>• Event viewer\Applications and services logs\DFS Replication</li> </ul>

## Tools and Commands for Troubleshooting Domain Controller Configuration

If the logs do not provide enough information, you can use the following tools for troubleshooting:

- **Dcdiag.exe.** Runs multiple tests to assess the overall health of AD DS.
- **Repadmin.exe** – Assists administrators in diagnosing replication problems between Windows domain controllers.
- **AutoRuns.exe** –Shows you what programs are configured to run during system bootup or logon, and shows you the entries in the order Windows processes them.
- **Task Manager** –Provides detailed information about how to run applications, processes, and services and provides performance and networking statistics.
- **MSInfo32.exe** –Displays a comprehensive view of your hardware, system components, and software environment.
- **Network Monitor** – Enables capturing and protocol analysis of network traffic.

## Methodology for Troubleshooting

Many errors are easy to correct. Check these items first:

- Is this a syntax error? Check the naming, credentials, and syntax of Windows PowerShell.
- Did the prerequisite check fail? Resolve the issue and try again.
- Did the error occur during the promotion phase? Examine the logs. Use Dcdiag and Repadmin to validate Active Directory health.
- Check for third-party software that may be preventing the promotion and remove it.



## Lesson 2

# Configuring AD DS Domain Controllers

After you install AD DS and create new domain controllers, you must address several Active Directory configuration issues. You can address some of these issues, such as creating a global catalog, during or after the promotion. You address others after the promotion.

### Lesson Objectives

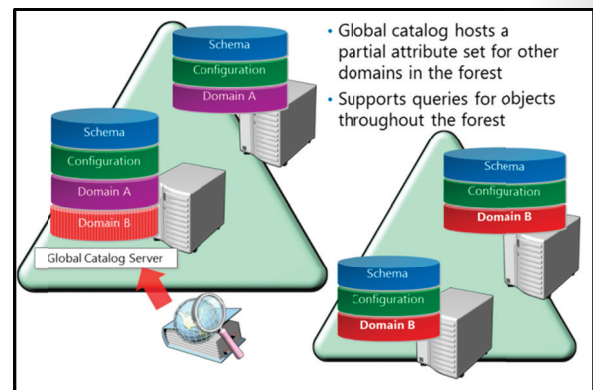
After completing this lesson you will be able to:

- Configure the global catalog.
- Configure universal group membership caching.
- Configure operations masters.

Manage domain and forest functional levels.

### Configuring the Global Catalog

The global catalog is a special partition of Active Directory that stores information about all Active Directory objects. It does not contain all attributes of all objects, but instead contains a subset of attributes that are useful for searching. The global catalog mainly occurs in a multi-domain environment. It enables searches across domain boundaries to find objects in Active Directory. The global catalog acts as an index of Active Directory. Certain applications rely on the global catalog, such as Exchange Server.



### Global Catalog Characteristics

Global catalogs are unique to Active Directory and have the following characteristics:

- The global catalog can only exist on a domain controller.
- At least one global catalog must exist in every forest.
- It is possible and frequently desirable to have multiple global catalogs. For example, have a global catalog in each AD DS site so that user authentication occurs in a timely, efficient manner.
- Global catalogs can be created during the promotion process or at any time after.
- Global catalogs can affect replication traffic.
- Global catalogs listen on ports 3268/3269 by default.

### Creating a Global Catalog

The first domain controller in the forest is a global catalog because at least one global catalog is required per forest. You can remove the domain controller's designation as a global catalog later after you have created other global catalogs.

For each additional domain controller, you can create a global catalog by ensuring that you select the check box in the Active Directory Configuration Wizard during the promotion. By default, all domain controllers are assumed to be global catalogs.

You can also add or remove the global catalog from a domain controller by using Active Directory Sites and Services MMC and editing the properties of the NTDS Settings node of the domain controller.

Alternatively, you can use the Active Directory module of Windows PowerShell to enable a global catalog.

## Configuring Universal Group Membership Caching

Universal groups include users and groups from multiple domains in a forest. The membership of universal groups is replicated in the global catalog. When a user logs on, the user's universal group membership is obtained from a global catalog server. If a global catalog is not available then universal group membership is not available. Configuring universal group membership caching addresses this problem.



**Note:** This problem does not arise when every domain controller is a global catalog.

- Universal group membership replicated in the global catalog:
  - Normal logon: User's token built with universal groups from global catalog
  - Global catalog not available at logon: domain controller denies authentication
- If every domain controller is a global catalog, this is never a problem
- If connectivity to a global catalog is not reliable:
  - Domain controllers can cache universal group membership for a user when user logs on
  - Global catalog later not available: User authenticated with cached universal groups
- In sites with unreliable connectivity to global catalog, enable universal group membership caching
- Right-click **NTDS Settings** for site, and select **Properties**:
  - Enables Universal Group Membership Caching for all domain controllers on the site

You can alleviate denial of authentication by enabling Universal Group Membership Caching on the local AD DS site. With this enabled, by default all domain controllers in that site obtain universal group membership information from a global catalog for a user when the user first logs on to the site. The domain controller caches that information indefinitely, as long as it can update universal group membership information every eight hours. If the local domain controller cannot contact a global catalog, then the cached group membership information is considered invalid after seven days. This value is called the 'staleness interval' and is set in the registry. If a network outage of less than seven days prevents the local domain controller from contacting the global catalog, the user is still authenticated successfully by using the cached group information.

### Enabling Universal Group Membership Caching

You can also enable Universal Group Membership Caching on a domain controller by using Active Directory Sites and Services MMC, and editing the properties of the NTDS Settings node of the domain controller.

You can also use the Active Directory module for Windows PowerShell to enable Universal Group Membership Caching.

## Configuring Operations Masters

In any replicated database, such as AD DS, some tasks must be performed by only one AD DS replica holder because they are impractical to perform in a multi-master manner. For example, only one domain controller can be in charge of synchronizing the time across the domain. In an Active Directory domain, operations masters, also known as flexible single master operations, or FSMO, are domain controllers that additionally provide a specific function. There are five specific operations master roles that must be filled. Any domain controller that meets the prerequisites can perform these roles.

- Forest-wide:
  - Domain naming: Adds/removes domains to/from the forest
  - Schema: Makes changes to the schema
- Domain-wide:
  - RID: Provides “pools” of RIDs to domain controllers, which use them for SIDs
  - Infrastructure: Tracks changes to objects in other domains that are members of groups in this domain
- PDC: Plays several very important roles:
  - Emulates a Primary Domain Controller (PDC): compatibility
  - Special password update handling
  - Default target for Group Policy updates
  - Master time source for domain
  - Domain master browser




**Note:** A RODC cannot host any operation master roles because, by design, it cannot directly modify the copy of AD DS it holds.

Two of the operations master roles only exist one time for the whole forest. These two roles exist only in the Forest Root Domain and are shown in the following table.

Role	Description
Domain Naming Operations Master	You use the domain naming role when you add or remove domains in the forest. When you add or remove a domain, the domain naming master must be available, or the operation fails.
Schema Operations Master	The domain controller holding the schema master role is responsible for making any changes to the forest's schema. All other domain controllers hold read-only replicas of the schema. If you want to modify the schema or install an application that modifies the schema, try to do it directly on the domain controller holding the schema master role. Otherwise, the changes that you request must be sent to the schema master to be written into the schema. If the Schema Master is inaccessible, all attempts to modify the schema will fail.

These roles can be transferred to other domain controllers if required. If a domain controller that is currently holding a role should stop functioning, the role can be forcibly seized by another domain controller.

The other three roles exist in every domain in the forest. They are shown in the following table.

Role	Description
Relative Identifier (RID) Operations Master	<p>The SID of a security principal must be unique. Any read/write domain controller in a domain can create accounts, and therefore, issue SIDs. Active Directory domain controllers generate SIDs by incorporating a unique RID into the domain SID. The RID master for the domain allocates pools of unique RIDs to each domain controller in its domain. In the past it was possible for a domain to reach the limit of the RID issuance (maximum possible of <math>2^{30}</math> or 1,073,741,823). New safeguards were put into place for Windows Server 2012 RID Masters, which include issuing warnings in Event logs when overall RIDs allocated are approaching 10% of usage. You can also increment the number of RIDs allocated to <math>2^{31}</math> (grand total of 2,147,483,648 SIDs).</p> <p> <b>Note:</b> This is the only one of the five FSMO roles that was improved in Windows Server 2012. All other roles retain same functionality as earlier versions.</p>
Infrastructure Operations Master	<p>In a multi-domain environment, it is common for a local object to reference security principals in other domains. For example, a group can include members from another domain. If the security principal in the other domain is moved or renamed, the infrastructure master in the same domain as the local group updates each remote group member's attribute accordingly.</p>
PDC Emulator Operations Master	<p>Emulates a Primary Domain Controller (PDC) and is probably the most important FSMO role for day-to-day functionality.</p> <p><b>Password handling.</b> When passwords are changed, the PDC emulator is updated immediately.</p> <p><b>Focus of Group Policy.</b> When Group Policy objects (GPOs) are being created or edited the action is being performed, by default, on the PDC emulator.</p> <p><b>Time source for the domain.</b> The PDC emulator provides the time source for all computers joined to AD DS to synchronize to.</p> <p><b>Domain Master Browser.</b> When you open the Network window and see the list of computers, you are seeing a list that is created by the browser service.</p>

These roles can be transferred to any domain controller in the domain. They do not all have to run on the same domain controller. For example, one domain controller might hold the PDC Emulator role while another holds the RID Master role. If a domain controller that is currently holding a role should stop functioning, the role can be forcibly seized by another domain controller.

## Managing Domain and Forest Functional Levels

By raising the functional levels, you can enable functionality offered by new versions of Windows. New features are not backward-compatible with older version of Windows Server. Similarly, until all domain controllers are running Windows Server 2008, or 2008 R2 or Windows Server 2012 you cannot implement its improvements to AD DS.

There are two major requirements for raising the functional level:

- All domain controllers must run the correct version of Windows Server.
- You must raise functional levels manually.

- Domain functional levels
- Forest functional levels
- New functionality requires that domain controllers run:
  - Windows 2000
  - Windows Server 2003
  - Windows Server 2008
  - Windows Server 2008 R2
  - Windows Server 2012
- Active Directory Domains and Trusts
- Cannot raise functional level while domain controllers are running previous Windows versions
- Cannot add domain controllers running previous Windows versions after raising functional level



**Note:** The operating system version of the domain controller determines the functional levels. Member servers can be running any version of Windows Server except for Windows NT 4.0. If you raise the functional level to Windows Server 2008, Windows NT 4.0 can no longer be a domain member.

Raising the functional level of either the domain or the forest is a one-way operation. You can never lower a functional level. Therefore, after you have raised the domain functional level to Windows Server 2008, for example, you cannot at a later date add a domain controller running at Windows Server 2003 to the same domain.

A forest can have domains that run at different functional levels, but after the forest functional level is raised, you cannot add a domain controller running a lower version of Windows to any domain in the forest.

Windows Server 2012 forest functional level and domain functional level do not implement new features from Windows 2008 R2 functional level.

## Lesson 3

# Implementing Service Accounts

One common issue that most organizations face is how to securely manage accounts that are used for network services. Many applications use services that require an account for service startup and authentication. As with typical user accounts, you must also effectively manage service accounts to ensure security and reliability.

### Lesson Objectives

After completing this lesson you will be able to:

- Describe managed service accounts.
- Describe group managed service accounts.
- Configure managed service accounts.
- Manage service principal names.

### What Are Managed Service Accounts

Applications are frequently configured to execute non-interactively on servers that use the security authentication context of the Local Service, Network Service, or Local System accounts. Because these accounts are typically shared by many applications and processes, you cannot isolate their credentials. That is to say, you cannot customize the security settings of these accounts without also affecting all applications and processes that are mapped to them. A Managed Service Account provides an application with its own unique service account. In Windows Server 2012, administrators no longer have to manually administer the credentials for this account.

Used to automate password and SPN management for service accounts used by services and applications

- Requires a Windows Server 2008 R2 server installed with:
  - Microsoft .NET Framework 3.5.x
  - Active Directory module for Windows PowerShell
- Recommended to run with AD DS configured at the Windows Server 2008 R2 functional level
- Can be used in a Windows Server 2003 or Windows Server 2008 AD DS environment:
  - With Windows Server 2008 R2 schema updates
  - With Active Directory Management Gateway Service

Managed service accounts in Windows Server 2012 offer the following benefits:

- Automatic password management. A managed service account automatically maintains its own password including password changes. This can better isolate services from other services on the computer.
- Simplified Service Principal Name (SPN) management. SPN management can be automatically managed if the AD DS domain is configured at the Windows Server 2008 R2 domain functional level. For example, if the samAccountName property of the computer is changed, or if the DNS host name property is modified, the managed service account SPN automatically changes from the old name to the new name for all managed service accounts on the computer.

### Requirements for Using Managed Service Accounts

To use a managed service account, the server that runs the service or application must run Windows Server 2008 R2 or later versions. You must also ensure that the .NET Framework 3.5.x, and the Active Directory Module for Windows PowerShell are both installed on the server.



**Note:** In versions of Windows earlier than Windows Server 2012, Managed service accounts could not be shared between multiple computers. Each Managed Service Account had to be unique to the computer where the application was run. This type of service account is known as a Standalone Managed Service Account. New in Windows Server 2012 is the ability to create Managed Service Accounts that can be shared with more than one computer (for example, for a clustered set of servers). These types of Managed Service Accounts are called Group Managed Service accounts. They are discussed in the next lesson.

## Managing Service Principle Names

Service Principle Names (SPNs) represent the accounts in whose security context a service executes. SPNs support mutual authentication between a client application and a service. SPNs are built either from information that a client computer knows about a service or from a trusted third-party, such as Active Directory. SPNs are associated with accounts and an account can have a different SPN for each service it is used to authenticate and execute.

- Service Principle Names (SPNs) represent the user accounts under which services run.
- SPNs support mutual authentication between applications and services.
- An account can have a different SPN for each service it authenticates and executes.
- The basic syntax of a SPN is:  
`< service type >/< instance name >:< port number >/< service name >`

The basic syntax of a SPN is as follows.

```
< service type >/< instance name >:< port number >/< service name >
```

The elements of the syntax have the meanings described in the following table.

Element	Description
Service type	The type of service, such as www for World Wide Web service.
Instance name	The name of the instance of the service. Either the host name or IP address of the server that is running the service.
Port number	Port number that is used by the host for the service if it differs from the default.
Service name	This may be the DNS name of the host, or of a replicated service, or of a domain; or it can be the distinguished name of a service connection point object or of a remote procedure call (RPC) service object.

If service name and instance name are the same, as they are for most host-based services, then you can abbreviate a service principal name to two components, as follows.

```
< service type >/< instance name>
```

## Service Names in Active Directory

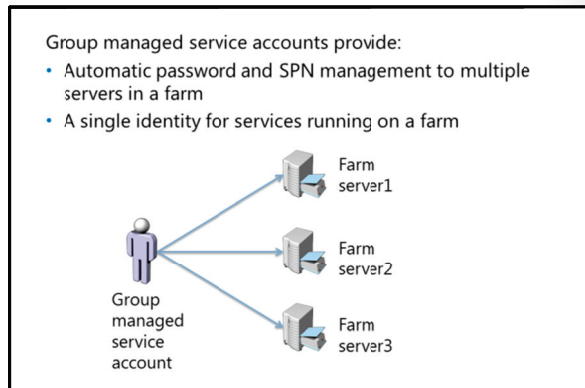
The syntax for service names in Active Directory includes the distinguished name of the instance of the service. The syntax is as follows.


```
< service type >/< host name >:< port number >/< distinguished name >
```




## What Are Group Managed Service Accounts?

As discussed in the previous lesson, Standalone Managed Service Accounts are managed domain-based accounts (that now include automatic password management and simplified SPN management for the service account) for single servers. Group Managed Service Accounts provide the same functionality but for multiple servers. When you connect to a service hosted on a server farm, such as the Network Load Balance (NLB) service, all computers that are running an instance of that service must use the same security principal. When a Group Managed Service Account is used as the service principal, the Windows Server 2012 AD DS manages the password for the account instead of relying on the administrator to manage the password.



 **Note:** Group Managed Service Accounts can only be configured and administered on computers that are running Windows Server 2012.

The group Managed Service Account has features to deal correctly with hosts that are kept offline for an extended time period. This means that you can deploy a server farm that uses a single Group Managed Security Account identity to which existing client computers can authenticate without knowing the instance of the service to which they are connecting.

 **Note:** For Windows Server 2012, the Windows PowerShell cmdlets default to managing the group Managed Service Accounts instead of the original standalone Managed Service Accounts.

## Demonstration: Configuring Group Managed Service Accounts

In this demonstration you will see how to create a group managed service account and associate the account with a server.

### Demonstration Steps

1. Log on to LON-DC1 as Administrator.
2. Create the KDS root key using the New-KdsRootKey cmdlet. Make the effective time minus 10 hours so the key is effective immediately.
3. Create the new service account named Webservice for the host LON-DC1.
4. Associate the Webservice managed account with Lon-DC1.
5. Verify the group managed service account was created by using the Get-ADServiceAccount cmdlet.



## Lesson 4

# Implementing Group Policy in AD DS

Group Policy has become the major tool for controlling the computing environment in an organization. This lesson points out the new features for Windows Server 2012 and describes some management techniques for controlling users and computers.

### Lesson Objectives

After completing this lesson you will be able to:

- Describe the new features in Group Policy.
- Manage Group Policy objects (GPOs).
- Configure Group Policy processing.
- Describe Group Policy client-side extensions.
- Troubleshoot Group Policy.
- Describe best practices for Group Policy implementation.

### What's New in Group Policy in Windows Server 2012?

Group Policy was introduced in Windows 2000. Each successive Windows version has introduced new tools or management features, such as the Group Policy Management Console (GPMC). Group Policy in Windows Server 2012 includes the following new features.

#### Graphical User Interface for Managing Fine-Grained Password Policy

New in Windows Server 2012 is the ability to manage this GPO object set from the console of the Active Directory Administrative Center. Managing domain user account password policy by group membership was an option since the initial release of Windows Server 2008. When it is enabled, any password policy associated with the user's group membership takes precedence over the default of the domain account policy. However, in earlier versions of Windows Server there was no single interface for implementing and managing type of GPO. The new GUI simplifies using this feature.

New features for Windows Server 2012 include:

- Group Policy Infrastructure Status
- Remote Policy Refresh
- New RSOP Logging Data

#### Group Policy Infrastructure Status

The Group Policy Infrastructure Status tool is a new tab in the GPMC. It displays the status of Active Directory and SYSVOL replication as it relates to Group Policy. This feature enables you to detect the current status by comparing the replication status of all domain controllers.

#### Remote Policy Refresh

You can now use GPMC to target an organizational unit (OU) and force Group Policy refresh on all its computers and their currently logged-on users. Right-click any organizational unit in the GPMC, and then click **Group Policy Update**. The update occurs within 10 minutes (randomized on each targeted computer) to prevent overwhelming a domain controller.

Also, a new Windows PowerShell cmdlet, named **Invoke-GpUpdate**, functions in the same manner as the command line GpUpdate utility.

### New RSOP Logging Data

When you use the Group Policy Results wizard or GpResult /H command line tool to generate an HTML Resultant Set of Policy (RSOP) report, you now see an updated Summary section that provides information such as network speed and whether a policy is functioning correctly or not.



**Note:** Remote RSOP logging and Group Policy refresh require you to open firewall ports on the targeted computers. This means enabling incoming communication for RPC, WMI/DCOM, event logs, and scheduled tasks.

## Managing GPOs

You must manage group policies as any other object in Active Directory. Group Policy must be created, edited, applied to containers, and backed up. The GPMC is the main tool for managing Group Policy.

### Creating, Editing, and Linking Policies

Group Policy management has the following characteristics:

- Create GPOs in the Group Policy Objects folder in the GPMC. You must have administrative rights in the domain or membership in the Group Policy Creator Owners group to create GPOs.
- Edit GPOs by using the Group Policy Management Editor. You can use policies to configure and apply thousands of settings.
- You can link GPOs to containers by using the GPMC. You can link a single GPO to multiple containers.

The GPMC is the main Group Policy management tool. It is used to:

- Create GPOs
- Edit GPOs through the GPO Editor
- Link GPOs
- Back up GPOs
- Restore GPOs
- Copy GPOs
- Import GPOs

### Backing Up and Restoring GPOs

You should back up Group Policies regularly. The first time that you back up a GPO, you must specify the location of the backup folder.

To back up GPOs in the GPMC, use the following procedures:

- To back up individual GPOs, right-click the GPO, and then click **Back Up**.
- To back up all GPOs, right-click the GPO folder, and then click **Back Up All**.

To restore an existing GPO to an earlier version of the GPO:

1. Open the Group Policy Objects folder.
2. Right-click the GPO that you want to restore.
3. Click **Restore from Backup**.

To restore a deleted GPO:

1. Right-click the **Group Policy Objects** folder.
2. Click **Manage Backups**.

3. Click the policy that you want to restore from the backup folder.
4. Click **Restore**.

### Copy or Import GPOs

By using the import and copy operations in the GPMC, you can transfer GPOs across domains and across forests. This is useful if you maintain separate test and production environments and want to replicate the content from one environment to the other. The GPMC enables you to modify certain settings as part of the import or copy operation. Specifically, you can modify references to security principals, such as users, groups, and computers, and to Universal Naming Convention (UNC) paths that exist in the GPO. You can modify security principals and UNC paths in the destination GPO by using a migration table with the import or copy operation. For example, the test environment might use a different UNC path for folder redirection than the production environment. You can use a migration table to map the test environment UNC path of the production UNC path.

A copy operation uses an existing GPO as its source and creates a new GPO as the destination. The administrator can choose to preserve the existing permissions or use the default GPO permissions. To copy an existing GPO:

1. Right-click the GPO.
2. Click **Copy**.
3. Paste the GPO into the Group Policy Object folder.

The import operation transfers settings into an existing GPO in Active Directory using a backed up GPO as the source. Importing does not modify the permissions or links associated with the destination GPO. Importing does not merge with any existing settings in the destination GPO, but will overwrite all settings. To import a GPO:

1. Right-click the GPO you are importing settings into.
2. Click **Import Settings**.
3. Follow the steps in the Import Settings Wizard.

### Configuring Group Policy Processing

When you link a Group Policy to a container, the settings affect all users, groups, or computers in that container and all child containers under that parent. For example, a GPO linked to the domain container inherits down to all child containers in the domain. Because you can link GPOs directly to the site, domain, or OU containers, there is the potential for settings in different GPOs to conflict. For example, a setting in a GPO at the domain level might be enabled while the same setting in a GPO linked to an OU may be disabled. This conflict is resolved through precedence. GPO settings are applied in the following order:

1. Local policies
2. Site linked GPOs
3. Domain linked GPOs

- GPOs are applied in an order known as precedence. When multiple policies apply to the same container the precedence can be set.
- GPO settings inherit down and merge to provide the cumulative effect of all settings. Inherited GPOs can be viewed on the **Inheritance** tab.
- Inheritance can be blocked. Inheritance cannot be blocked for only selected GPOs – it is all or none.
- GPOs can be enforced. Enforcement overrides blocking inheritance and conflicting settings.
- Loopback applies the user settings from the policy that applied the loopback setting. It is typically used for Remote Desktop Services and special cases.
- Security filtering. Permissions on the GPOs can control which object receive settings.
- WMI Filters. WMI can query for conditions under which the GPO settings are applied.

4. OU linked GPOs
5. Child OU linked GPOs

Policy settings inherit down and merge so that objects receive the cumulative effect of all GPOs. If you link multiple GPOs to the same container then they are applied in the order in which they were linked. However, you can set precedence to control the order of application to that container. If there is a conflict in GPO settings, the last GPO applied has precedence and is the effective one. In other words, the user or computer receives all the GPO settings in the path of their container and linked directly to their container, but if there is a conflict, the latest setting is the one in effect.

Group Policy provides mechanisms to modify the way GPO settings are processed. You can block inheritance and enforce policies.

### Blocking Inheritance

You can configure a domain or OU to prevent the inheritance of policy settings. This option blocks all inherited Group Policy settings from GPOs linked to parents in the Group Policy hierarchy. You cannot use it to block only selected inherited policies. It does not block GPOs that are linked directly to the container. You should use the **Block Inheritance** option sparingly. When you block inheritance, you make it more difficult to evaluate Group Policy precedence and inheritance.

### Enforcing a GPO Link

You can set a GPO link to be Enforced. When you set a GPO link to Enforced, that GPO takes the highest level of precedence. Policy settings in that GPO then prevail over any conflicting policy settings in other GPOs. In addition, a link that is enforced applies to child containers even when those containers are set to Block Inheritance. The Enforced option causes the policy to apply to all objects within its scope. The Enforced setting causes policies to override any conflicting policies and applies regardless of any other settings.

### Loopback Processing

By default a user receives the settings from GPOs inherited by, and linked to, the OU where their user account resides. There are situations, however, in which you might want to configure a user differently, depending on the computer that is being used. For example, you might want to lock down and standardize user desktops when users log on to computers in closely managed environments, such as conference rooms, reception areas, laboratories, classrooms, and kiosks. You might also want to apply specific settings for virtual desktop infrastructure (VDI) scenarios. This includes remote virtual machines and Remote Desktop Services (RDS), known as Terminal Services in earlier versions.

The loopback setting a user's typical GPO settings to be disregarded and applies the user settings associated with the GPO instead.

The loopback setting is located in the Computer Configuration\Policies\Administrative Templates\System\Group Policy folder in the GPO.



**Note:** There is an option in the loopback setting to merge the loopback user settings with their typical settings. But the default is to replace their typical settings with the settings in the loopback GPO.

## Security Filtering

Each GPO has a Discretionary Access Control List (DACL) that defines permissions to the GPO. You must apply two permissions, Allow Read and Allow Apply Group Policy, to a user or computer. By default, Authenticated Users have the Allow Apply Group Policy permission on each new GPO. This means that by default, all users and computers are affected by the GPOs settings. Therefore, by adjusting the permissions on the GPO you can control who receives them. There are two approaches to do this.

- To apply the GPO to only some users, groups or computers:
  1. Remove the Authenticated Users group from the DACL.
  2. Add the users, groups or computers you want to receive the policies.
  3. Grant them Read and Apply Group Policy permissions.
- To prevent some users, groups or computers from receiving the GPO settings:
  4. Add them to the DACL.
  5. Deny them the Apply Group Policy permission.

You access the DACL from the **Delegation, Advanced** tab of the GPO.

## WMI Filtering

You can also use Windows Management Instrumentation (WMI) to control the scope of GPO application, depending on attributes of the destination computer. You can use WMI queries to check for hardware or software conditions that must exist for settings to be applied. For example, a WMI query may check for an operating system version, make or model, or the RAM in the system to determine whether GPO settings should be applied. WMI filters can query for hundreds of different parameters.

## Group Policy Client Side Extensions

The Group Policy Client service determines which GPOs to apply to the client. This service downloads any GPOs that are not already cached. Then, a series of processes called client-side extensions interpret the settings in a GPO and make appropriate changes to the local computer or to the currently logged-on user. There are client-side extensions for each major category of policy setting. For example, there is a security client-side extension that applies security changes, a client-side extension that executes startup and logon scripts, a client-side extension that installs software, and a client-side extension that makes changes to registry keys and values. Each new version of Windows has added client-side extensions to extend the functional reach of Group Policy. There are several dozen client-side extensions now in Windows.

- How GPOs and their settings are applied
- Group Policy Client retrieves ordered list of GPOs
- GPOs are downloaded, and then cached
- Components called CSEs process the settings to apply the changes:
  - One for each major category of policy settings: Security, registry, script, software installation, mapped drive preferences, and so on.
  - Most CSEs apply settings only if the GPO as a whole has changed:
    - Improves performance
    - Security CSE applies changes every 16 hours
  - GPO application is client computer driven (pull)



**Note:** For client computers running Windows XP to accept Group Policy Preferences the client-side extensions for Windows XP preferences must be downloaded and installed on each client computer.

Group Policy is applied at the client computer side at startup for computer settings and when users log on for user settings. Group Policy is also refreshed on the client computer at regular, configurable intervals. The default interval is 90 minutes. The Group Policy client pulls the GPOs from the domain, triggering the client-side extensions to apply settings locally. Group Policy is not a push technology.



**Note:** You can manually refresh Group Policy from the GPMC in Windows Server 2012 or by using the GpUpdate command prompt utility on the client workstation.

Policies remain in force on the client even if the client is not connected to the corporate LAN. For example, mobile laptop users continue to have the GPO settings enforced because those settings are cached on the client. But mobile laptop users receive no changes to policy settings until they reconnect to the LAN.



**Note:** If client computers use cached credentials to speed up the logon process, then the user does not see the effect of several GPO settings until after two logons.

Policies are not re-applied on the client systems unless a change in a policy setting is detected. An important exception to the default policy processing settings is settings managed by the security client-side extension. Security settings are reapplied every 16 hours even if a GPO has not changed.



**Note:** You can configure client-side extensions to reapply policy settings at background refresh even if the GPO has not changed. To do this, define the settings in the Computer Configuration\Policies\Administrative Templates\System\ Group Policy node. To configure a client-side extension:

1. Open its policy processing policy setting, such as Registry Policy Processing for the Registry client-side extension.
2. Click **Enabled**.
3. Select the **Process even if the Group Policy objects have not changed** check box.

### Group Policies over Slow Links

If a slow network connection is detected then certain client-side extensions do not process GPO settings. For example, installing software is not practical across a slow network. By default, a slow connection is defined as 500 KBPS. However, you can configure this value in Group Policy. Also, you can configure each client-side extension in Group Policy to process even if a slow connection is detected.

These settings are always applied, even across a slow connection:

- Security settings
- Administrative Templates
- IPsec
- Encrypting File System (EFS)

These settings are not applied across a slow connection:

- Quotas
- Internet Explorer Maintenance
- Folder Redirection
- Scripts

- Wireless Network settings
- Software installations



**Note:** Older clients, such as Windows XP, use Ping to determine network speed. If you block Internet Control Message Protocol (ICMP) traffic, the connection always appears as a slow connection. Clients that are running Windows Vista or later versions use Network Location Awareness to determine connection speed.

## Troubleshooting Group Policy

There may be times when you must troubleshoot Group Policy. There are two main issues that can occur with Group Policy processing:

- Policies are not being applied to the client computer.
- Policies are applied, but the results are inconsistent or incorrect.

These two issues might arise for the following reasons:

- AD DS replication issues may prevent all domain controllers from receiving policies or policy updates.
- GPOs may be linked incorrectly to containers.
- Slow network conditions may exist.
- Policy filtering may be set.
- Inheritance or enforcement settings may be applied.
- The loopback setting may be turned on.
- Local computer policies may affect the results.

Start to troubleshoot by determining the scope of the issue. For example, is the issue widespread, or only affecting a single client? If the issue affects a single client, you should check for physical issues, such as incorrect configurations. These issues are usually easy to diagnose.

Check Event Viewer entries, Windows logs, and application and service logs. These can provide valuable information about the cause of issues. Log entries frequently direct you to the area in which to begin an investigation.

Most Group Policy issues are caused by:

- Inheritance
- Filtering
- Replication

Group Policy issues can be caused by Group Policy-specific issues, or they can be caused by unrelated issues like network connectivity or authentication problems.

Key Group Policy troubleshooting areas:

- Inheritance
- Security group or WMI filtering
- Replication
- Policy refresh



## Troubleshooting Inheritance

If none of the users or computers in an OU or child OUs receive policies that were linked to higher levels, it may be because of inheritance blocking. The GPMC displays a blue exclamation mark when inheritance is blocked. RSOP lists the GPOs that are being applied, and the GPOs that are being blocked. You can generate Group Policy results at the destination computer or from the GPMC through the Group Policy Results Wizard.

## Troubleshooting Filtering

GPO filtering may result from:

- Security filtering
- WMI filtering

Symptoms of filtering issues may appear as inconsistent application of policies in an OU. If some users, groups, or computers have filtering applied, they do not receive policies that other users in the same OU receive.



**Note:** If a WMI filter is deleted, the links to the WMI filter are not automatically deleted. If there is a link to a non-existent WMI filter, the GPO with that link is not processed until the link is removed or the filter is restored.

## Troubleshooting Replication

Group Policy information takes time to propagate or replicate from one domain controller to another.

Replication issues are most noticeable in remote sites with slow connections and long replication latency. You can use the new Status tab in the GPMC on Windows Server 2012 to determine the replication health of the GPO. If replication is an issue, you must determine whether the problem is with the File Replication Service (FRS) or with AD DS replication. There are two simple tests that you can use to determine the issue:

- For SYSVOL replication, put a small test file into the SYSVOL directory. See whether it replicates to other domain controllers.
- For AD DS replication, create a test object, such as an OU. See whether it replicates to other domain controllers.

## Troubleshooting Policy Refresh

Some users rarely restart or even log off their systems. Several Group Policy settings cannot be refreshed during a typical refresh cycle. Some settings require a logoff or a restart to be applied. In fact, because of cached credentials, many settings require two logons for the user to see the effect of the setting. If some users do not receive the policy settings, ensure that they restart or log off and on two times to rule out the effect of cached credentials.



## Best Practices for Implementing Group Policy

Group Policy is a very powerful tool, but you must apply it correctly. Implementing a Group Policy solution involves planning, designing, deploying, and maintaining the solution. There are some best practices that you should follow.

### Plan Your Deployment

Define the scope of application of Group Policy. Define what types of settings are global to all users and computers and design or modify the OU structure to accommodate Group Policy application. You should design the OU structure with Group Policy in mind and enhance the inherited nature of Group Policy settings by grouping objects in a hierarchy that enables that flow of Group Policy settings.

- Plan the Group Policy deployment
- Create standard desktop configurations
- Do not use the Default GPOs for other purposes
- Use inheritance modifications sparingly
- Employ Loopback processing for special case scenarios
- Implement a change request process

### Create Standard Desktop Configurations

One of the goals of controlling the computing environment is to provide consistency. Standard desktop configurations for various user types or departments can make system repair or replacement a simpler task if many of the configuration settings are delivered by using Group Policy.

### Do Not Use the Default Domain Policy or Default Domain Controllers Policy for Other Purposes

These two default policies provide basic settings for the domain, such as password policies, and for domain controllers, such as auditing settings. If you want to apply other configuration settings to the domain or to domain controllers, create new policies. Use the default policies for password, auditing and security settings only.

### Use Inheritance Modifications and Filtering Sparingly

Heavy use of blocking and enforcing of policies make troubleshooting more difficult. Also try to avoid security and WMI filtering unless it is required.

### Use Loopback Processing for Special Case Scenarios

Loopback can solve issues with desktop standardization for scenarios where the system users log on to special purpose systems, such as Remote Desktop Services or kiosk computers.

### Implement a Change Request Process

Limit changes to Group Policy settings to a small group of administrators. All changes should be approved and documented. Consider using the Advanced Group Policy Management (AGPM) tool available with the Microsoft Desktop Optimization Pack (MDOP).

## Lesson 5

# Maintaining AD DS

Maintaining the health of the AD DS is an important aspect of an administrator's job. In this lesson, you will learn how to use Windows Server Backup to effectively backup and restore AD DS and domain controllers. You will also learn how to optimize and protect your directory service so that if a domain controller does fail, you can restore it as quickly as possible.

### Lesson Objectives

After completing this module, you will be able to:

- Describe options for backing up AD DS.
- Describe options for restoring AD DS.
- Describe the Active Directory Recycle Bin.
- Describe AD DS snapshots.
- Optimize the AD DS database.

### Options for AD DS Backup

Windows Server Backup was introduced in Windows Server 2008. It enables you to back up and restore a server, its roles, and its data. Windows Server Backup is installed as a feature in Server Manager.



**Note:** The Windows Server Backup MMC appears on the Tools list in Server Manager even though the feature is not actually installed until you manually add the feature.

- Windows Server Backup snap-in
- Wbadmin.exe
- Backups can be manual or automated
- Back up to CD/DVD/HDD
- You must back up all critical volumes for AD DS
  - System volume
  - Boot volume
  - Volumes hosting SYSVOL, AD DS database (NTDS.dit), logs

Windows Server Backup provides a snap-in administrative tool and the WBAAdmin command line tool (Wbadmin.exe). Both the snap-in and the command line enable you to perform manual or automatic backups to an internal or external disk volume, a remote share, or optical media. Backing up to tape is no longer supported by Windows Server Backup.

In earlier versions of Windows, backing up Active Directory involved creating a backup of the SystemState. In Windows Server 2012, the SystemState still exists, but it is physically larger in size. Because of interdependencies between server roles, physical configuration, and Active Directory, the SystemState is now a subset of a Full Server backup and, in some configurations, might be just as large as a full server backup. To back up a domain controller, you must back up all critical volumes fully.

Windows Server Backup enables you to perform one of the following types of backups:

- Full server
- Selected volumes
- System State
- Individual files or folders

When you use Windows Server Backup to back up the critical volumes on a domain controller, the backup includes all data that resides on the volumes that host the:

- Boot files, which consist of the Bootmgr file and the Boot Configuration Data (BCD) store.
- Windows operating system and the registry.
- SYSVOL tree.
- Active Directory database (Ntds.dit).
- Active Directory database log files.

To perform a backup, you must first install the Windows Server Backup feature. You can then use the Windows Server Backup console to create backup jobs. The Actions pane in the Windows Server Backup console enables you to start a wizard to perform a scheduled backup or a one-time backup job. The wizard prompts for a backup type, backup selection, backup destination and schedule (if performing a scheduled job).

## Options for AD DS Restore

When a domain controller or its directory is corrupted, damaged, or failed, you can restore the system by using several options.

The first option is called typical restore or nonauthoritative restore. In a normal restore operation, you restore a backup of Active Directory as of a known good date. Effectively, you roll the domain controller back in time. When AD DS restarts on the domain controller, the domain controller contacts its replication partners and requests all subsequent updates. The domain controller "catches up" with the rest of the domain by using standard replication mechanisms. Normal restore is useful when the directory on a domain controller was damaged or corrupted, but the problem has not spread to other domain controllers. This is not a method that works if you are trying to restore a deleted object and the deletion has replicated to the other domain controllers.

- Nonauthoritative (normal) restore
  - Restore domain controller to previously known good state of Active Directory
  - Domain controller is updated by using standard replication from up-to-date partners
- Authoritative restore
  - Restore domain controller to previously known good state of Active Directory
  - "Mark" objects that you want to be authoritative:
  - Windows sets the version numbers very high
  - Domain controller is updated from its up-to-date partners
  - Domain controller sends authoritative updates to its partners
- Full Server Restore
  - Typically performed in Windows Recovery Environment
- Alternate Location Restore

If the typical restore does not work, you can perform an authoritative restore. In an authoritative restore, you restore the known good version of Active Directory just as you do in a typical restore. However, before restarting the domain controller, you mark the objects that you want to recover (the deleted objects) as authoritative so that they replicate from the restored domain controller to its replication partners. Behind the scenes, when you mark objects as authoritative, Windows increments the version number of all object attributes to be so high that the version is guaranteed to be higher than the version number of the deleted object on all other domain controllers. When you restart the restored domain controller, it replicates from its replication partners all changes that are made to the directory. It also notifies its partners that it has changes, and the version numbers of the changes ensure that partners take the changes and replicate them throughout the directory service.

The third option for restoring the directory service is to restore the whole domain controller. You do this by booting to the Windows Recovery Environment and restoring a full server backup of the domain controller. By default, this is a typical restore. If you must also mark objects as authoritative, you must restart the server in the Directory Services Restore Mode and set those objects as authoritative before starting the domain controller into typical operation.

Finally, you can restore a backup of the SystemState to an alternative location. This enables you to examine files and, potentially, to mount the NTDS.dit file as described in the previous lesson. You should not copy the files from an alternative restore location over the production versions of those files. Do not do a piecemeal restore of Active Directory. This option is also used if you want to use the Install From Media option for creating a new domain controller.

## How does the Active Directory Recycle Bin Work?

The Active Directory Recycle Bin was introduced in Windows 2008 R2. You could only access this feature by using Windows PowerShell cmdlets and the Ldp.exe LDAP utility.

In Windows Server 2012 you can now access the Active Directory Recycle Bin from the Active Directory Administrative Center. This simplifies the recovery of Active Directory objects that were erroneously deleted. It lets administrators enable the Recycle Bin and locate or restore deleted objects in the domain. It is no longer required to use Windows PowerShell or Ldp.exe to enable the recycle bin or restore objects in domain partitions.

- The Active Directory Recycle Bin:
  - Cannot be disabled once it is enabled
  - Now has a user interface to simplify restoration of objects
  - Is enabled and accessed through the Active Directory Administration Center
  - Cannot restore sub-trees of object in a single operation
  - Requires the forest level be at least Windows Server 2008 R2
  - Requires Enterprise Admins
  - Increases the size of the Active Directory database
  - Objects are preserved in the recycle bin for the tombstone lifetime: 180 days by default
  - Deleted object can be viewed in the Deleted Object folder
  - Objects can be restored by selecting them and choosing Restore

### Active Directory Recycle Bin Characteristics

The Active Directory Recycle Bin has the following characteristics:

- It must be manually enabled. As soon as it is enabled, you cannot disable it.
- The Active Directory Recycle Bin cannot restore sub-trees of objects in a single action. For example, if you delete an OU with nested OUs, users, groups, and computers, restoring the base OU does not restore the child objects. That must be done in a subsequent operation.
- Active Directory Recycle Bin requires at least Windows Server 2008 R2 Forest Functional Level.
- You must be a member of the Enterprise Admin group to use the Active Directory Recycle Bin.
- The recycle bin increases the size of the Active Directory database (NTDS.DIT) on every domain controller in the forest. Disk space that is used by the recycle bin continues to increase over time as it preserves objects and all attribute data.
- Objects are preserved in the recycle bin for an amount of time to match the tombstone lifetime of the forest. This is 180 days by default.
- After the Active Directory Recycle Bin is enabled, deleted restorable objects can be viewed in the Deleted Objects folder.

### Enabling the Active Directory Recycle Bin

To enable the Active Directory Recycle Bin:

1. From the Server Manager Tools menu access the Active Directory Administrative Center.
2. In the navigation pane select the domain that you want to manage.
3. In the Tasks (right side) pane click **Enable Recycle Bin**.
4. Acknowledge the warning dialog boxes to complete the action.

## Restoring Active Directory Objects

Because many objects are intentionally deleted in typical Active Directory operations, the Active Directory Administrative Center has advanced filtering criteria, making targeted restoration easier in large environments that have many deleted objects. The restore operation supports all the standard filter criteria options as any other search. Multiple search criteria can be combined. Common search criteria include:

- Object is user/inetorgperson/computer/group/organization unit
- Name
- When deleted
- Employee ID
- First name
- Last name
- Job title
- City

As soon as you locate the object to be restored, right-click the object, and then click **Restore**.

- To restore the object to its original location, in the Tasks pane, click **Restore**.
- To restore an object to a different location, click **Restore To...**

You can restore multiple objects as long as they are all restored to the same location.

## Demonstration: Restoring AD DS Objects Using the Active Directory Recycle Bin

In this demonstration you will see how to:

- Enable the Active Directory Recycle Bin
- Use the recycle bin to restore a deleted object

### Demonstration Steps

1. Enable the Active Directory Recycle Bin.
2. Delete a current user.
3. Restore the user.

## What are AD DS Snapshots?

A snapshot captures the exact state of the directory service at the time of the snapshot. Unlike a backup, you cannot use a snapshot to restore data. However, you can use tools to explore the contents of the snapshot to examine the state of the directory service at the time the snapshot was made.

### Creating a Snapshot

You use the NTDSUtil to create and mount snapshots for viewing. To create a snapshot:

1. Open an elevated command prompt.
2. Type **ntdsutil**, and then press Enter.
3. Type **activate instance ntds**, and then press Enter.
4. Type **snapshot**, and then press Enter.
5. Type **create**, and then press Enter.
6. The command returns a message indicating that the snapshot set was generated successfully. The GUID that is displayed is important for commands in later tasks. Note the GUID or, alternatively, copy it to the Clipboard.
7. Type **quit** and then press Enter.

- Create a snapshot of Active Directory:
  - NTDSUtil
- Mount the snapshot to a unique port:
  - NTDSUtil
- Expose the snapshot:
  - Right-click the root node of Active Directory Users and Computers and select **Connect to Domain Controller**
  - Enter serverFQDN:port
- View (read-only) snapshot:
  - Cannot directly restore data from the snapshot
- Recover data:
  - Manually reenter data, or
  - Restore a backup from the same date as the snapshot

### Using the Database Mounting Tool to Mount a Snapshot

The Active Directory database mounting tool (Dsamain.exe) can improve recovery processes for the organization. It enables you to compare data as it exists in snapshots or backups that are taken at different times so that you can better decide which data to restore after data loss. This eliminates the need to restore multiple backups to compare Active Directory data.

To view the contents of a snapshot, you must mount the snapshot as a new instance of AD DS. This is also accomplished with NTDSUtil. To mount a snapshot:

1. Open an elevated command prompt.
2. Type **ntdsutil**, and then press Enter.
3. Type **activate instance ntds**, and then press Enter.
4. Type **snapshot**, and then press Enter.
5. Type **list all**, and then press Enter.  
The command returns a list of all snapshots.
6. Type **mount {GUID}**, where GUID is the GUID returned by the create snapshot command, and then press Enter.
7. Type **quit**, and then press Enter.
8. Type **quit**, and then press Enter.
9. Type **dsamain -dbpath c:\\$snap\_datetime\_volume%\windows\ntds\ntds.dit -ldapport 50000**, and then press Enter (you can use any available port number).
10. Do not close the Command Prompt window and leave the command that you just ran, Dsamain.exe, running while you continue to the next step.

## Viewing the Snapshot

After you have mounted the snapshot, you can use tools to connect to and explore the snapshot, including Active Directory Users and Computers.

To connect to a snapshot with Active Directory Users and Computers:

1. Open **Active Directory Users and Computers**.
2. Right-click the **root node**, and then click **Change Domain Controller**.
3. Click **<Type a Directory Server name[:port] here>** and enter the name of the domain controller and the port number that was used in the previous step. For example, **LON-DC1:50000** and then press Enter.
4. Click **OK**.

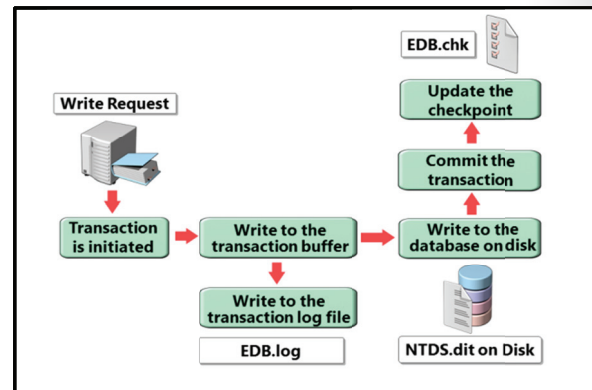
To unmount the snapshot:

1. Switch to the command prompt in which the snapshot is mounted.
2. Press Ctrl+C to stop **DSAMain.exe**.
3. Type **ntdsutil**, and then press Enter.
4. Type **activate instance ntds**, and then press Enter.
5. Type **snapshot**, and then press Enter.
6. Type **unmount GUID**, where GUID is the GUID of the snapshot, and then press Enter.
7. Type **quit**, and then press Enter.
8. Type **quit**, and then press Enter.

## AD DS Database Maintenance

The Active Directory database is stored as a file named NTDS.dit. When you install and configure AD DS, you can specify the location of the file. The default location is %systemroot%\NTDS. In the NTDS folder, there are other files that support the Active Directory database. They are:

- **EDB.log file.** The Edb.log file is the transaction log for Active Directory. When you must make a change to the directory, it is first written to the log file. The change is committed to the directory as a transaction. If the transaction fails, it can be rolled back.
- **EDB.chk.** The EDB.chk file functions like a bookmark into the log files, marking the location before which transactions are successfully committed to the database, and after which transactions remain to be committed.
- **Edbres0001.jrs and Edbres0002.jrs.** These two files are empty files of 10MB each. If the disk the database resides on should run out of space, these files provide the domain controller with the space to write pending transactions before safely shutdown AD DS services and dismounting the database.



The Active Directory database is self-maintaining. Every 12 hours, by default, each domain controller runs a process that is known as garbage collection. Garbage collection does two things. First, it removes deleted objects that have outlived their tombstone lifetime, which is 180 days by default. Second, the garbage collection process performs online defragmentation. Online defragmentation reorganizes the sectors rows of the database so that the blank rows are contiguous, very much like disk fragmentation reorganizes sectors of a disk so that free space is contiguous. However, this process does not reduce the file size of the database. It optimizes the internal order of the database. In most organizations, this will be sufficient.

To reduce the physical size of the NTDS.dit, perform offline defragmentation. To perform an offline defragmentation you must stop the AD DS. Then use the NTDSUtil to compact the database to a different location. Then replace the original NTDS.dit with the compacted version.



**Note:** Do not delete the original NTDS.dit, you only have to rename it.



## Lab: Implementing AD DS

### Scenario

A. Datum is an engineering and manufacturing company. The organization is based in London, England, but is quickly expanding the London location as well as internationally. As the company has expanded, some business requirements are changing as well. To address some business requirements, A. Datum had decided to deploy Windows Server 2012.

As the company expands, they must also expand their Active Directory infrastructure. You are assigned to implement new domain controllers and also to consider implementation of RODCs, where appropriate. Also, there are reports that Group Policies are not being applied on some computers, so you must troubleshoot. The company also wants to centralize management of all accounts that are being used for services, and to stop usage of local accounts for that purpose. Also, you must evaluate available techniques for AD DS maintenance.

### Objectives

- Deploy an RODC
- Implement Group Policy
- Configure and validate service accounts
- Maintain AD DS

### Lab Setup

Estimated time: **60 minutes**

Virtual machines	20417A-LON-DC1 20417A-LON-SVR3 20417A-LON-CL1
User Name	Adatum\Administrator
Password	Pa\$\$w0rd

For this lab, you will use the available virtual machine environment. Before you begin the lab, you must complete the following steps:

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, click **20417A-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on using the following credentials:
  - a. User name: **Adatum\Administrator**
  - b. Password: **Pa\$\$w0rd**
5. Repeat steps 2 and 3 for **20417A-LON-SVR3**, and **20417A-LON-CL1**. Do not log on to LON-SVR3 or LON-CL1 until instructed to do so.

## Exercise 1: Deploying a Read-Only Domain Controller

### Scenario

As company business expands, you must add domain controllers to new locations. Some locations do not have required physical security for server rooms so you decide to implement read-only domain controllers for these locations. Those servers are already in place at the branch location performing local file and print duties. You plan to install the RODC role remotely by using Server Manager from head office. You also plan to configure the RODC to cache passwords locally for members of the Managers group and assign administrative access to the server to the IT group.

The main tasks for this exercise are as follows:

1. Add LON-SVR3 as a Server to Manage.
2. Create a New Server Group.
3. Install the RODC Role Remotely.
4. Configure the Password Replication Policy and Administrative Access.

#### ► Task 1: Add LON-SVR3 as a Server to Manage

1. Log on to LON-DC1 as **Administrator** with a password of **Pa\$\$w0rd**.
2. Use the Server Manager Dashboard to add LON-SVR3 as a server to manage.

#### ► Task 2: Create a New Server Group

1. Use the Server Manager Dashboard to create a server group named **DCs**.
2. Add both **LON-SVR3** and **LON-DC1** to the group.

#### ► Task 3: Install the RODC Role Remotely

1. Use the Server Manager Dashboard to Add the **Active Directory Domain Services** role to LON-SVR3.
2. Open the notifications and complete the Post-deployment Configuration to promote LON-SVR3 to be a **Read only domain controller (RODC)** in the existing domain.
3. Set the Directory Services Restore Mode (DSRM) password to be **Pa\$\$w0rd**.
4. Accept the defaults for all other settings.

#### ► Task 4: Configure the Password Replication Policy and Administrative Access

1. Use Active Directory Users and Computers to configure the password caching options of **LON-SVR3** in such a way that passwords are cached on the RODC for members of the **Managers** group.
2. Configure the **IT** group to have administrative access to **LON-SVR3**.

**Results:** After completing this exercise, you will have added LON-SVR3 as a server to manage, created a server group, deployed an RODC remotely, and configured the password replication policy and administrative assignments for the RODC.

## Exercise 2: Troubleshooting Group Policy

### Scenario

Support technicians report that some Group Policy settings are not being applied as they should. Company Policy requires that:

- All domain users should not have access to change their desktop background.
- All domain users except the IT group should be unable to access Registry Editor.

Currently, there are some problems in the way the GPOs that deliver those settings are being applied.

You have to investigate, troubleshoot and resolve this problem.

The main tasks for this exercise are as follows:

1. Troubleshoot Group Policy Issues.
2. Correct Issues with Group Policy Application.
3. Verify Policies Are Being Applied.

#### ► Task 1: Troubleshoot Group Policy Issues

Determine the issue by logging on to LON-CL1 as an IT group user and a Manager group user. Check whether the policies are being applied correctly.

1. Log on as **Brad** with the password of **Pa\$\$w0rd**. Attempt to change the desktop background and attempt to start the Registry Editor.
2. Use GPREsult to determine the RSOP and then log off of LON-CL1.
3. Log on as **Bill** with the password of **Pa\$\$w0rd**. Attempt to change the desktop background and attempt to start the Registry Editor.
4. Use GPREsult to determine the RSOP.
5. Analyze the RSOP results to determine the problem.
6. Log off of LON-CL1.

#### ► Task 2: Correct Issues with Group Policy Application

1. Log on to LON-DC1 as **Administrator** with a password of **Pa\$\$w0rd**.
2. Use the Group Policy Management console to investigate and correct the issues.
3. Check the current status of the Managers OU.
4. Remove the block inheritance setting from the Managers OU to resolve the issue.
5. Think of a way to ensure that the Prohibit Registry Tools GPO will not be applied to IT group users.
6. Use Security Filtering to deny access to the policy to the IT security group.
7. Close the Group Policy Management console.

#### ► Task 3: Verify Policies Are Being Applied

1. Log on to LON-CL1 as **Bill** with a password of **Pa\$\$w0rd** and run the GPREsults utility.
2. Log off of LON-CL1.

3. Log on to LON-CL1 as **Brad** with a password of **Pa\$\$w0rd** and run the GPResult utility.
4. Log off of LON-CL1.

**Results:** After completing this exercise, you will be able to troubleshoot Group Policy issues, correct issues to apply Group Policy, and verify policies are being applied.

## Exercise 3: Implementing Service Accounts in AD DS

### Scenario

To this point, there was no consistent policy about accounts that were used for services. On some servers, local accounts were used, while others were using domain accounts. Also, password management for these accounts was not consistent. Some of them were having non-expiring passwords, while others were updated with new passwords manually. You decide to implement Managed Service Accounts to replace all these techniques. You will create the account and assign the account to the Web service DefaultAppPool.

The main tasks for this exercise are as follows:

1. Create and Associate a Managed Service Account.
2. Configure the Web Server Application Pool to Use the Group Managed Service Account.

#### ► Task 1: Create and Associate a Managed Service Account

1. Log on to LON-DC1 as **Administrator** with a password of **Pa\$\$w0rd**.
2. Create the KDS root key using the New-KdsRootKey cmdlet. Make the effective time minus 10 hours so the key will be effective immediately.
3. Create the new service account named Webservice for the host LON-DC1.
4. Associate the Webservice managed account with Lon-DC1.
5. Verify the group managed service account was created by using the **Get-ADServiceAccount** cmdlet.
6. Install the Webservice service account.

#### ► Task 2: Configure the Web Server Application Pool to Use the Group Managed Service Account

1. On LON-DC1, configure the **DefaultAppPool** to use the **Webservice\$** account as the identity.
2. Stop and start the application pool.

**Results:** After completing this exercise, you will have created and associated a managed service account, installed a managed service account on a web server, and verified password change for a managed service account.

## Exercise 4: Maintaining AD DS

### Scenario

As a part of maintenance plan, you are assigned with task to evaluate possibilities to quickly restore accidentally deleted objects. You decided to enable and test Active Directory snapshots and the AD DS Recycle Bin.

The main tasks for this exercise are as follows:

1. Create and View Active Directory Snapshots.
2. Enable the Active Directory Recycle Bin.
3. Delete a test user.
4. Restore the Deleted User.
5. To Prepare for the Next Module.

#### ► Task 1: Create and View Active Directory Snapshots

1. Switch to LON-DC1.
2. Start a command prompt using elevated credentials.
3. Run the following commands:
  - Ntdsutil
  - Snapshot
  - Activate instance ntds
  - Create
4. Mount the snapshot as a new instance of AD DS by running the Mount {GUID} command.
5. Close ntdsutil.
6. Use the **dsamain** command to expose the snapshot to LDAP port 50000.
7. Use Active Directory Users and Computers to delete Allie Bellew from the Research OU.
8. Use Active Directory Users and Computers to connect LON-DC1 to the snapshot instance at port 50000.

#### ► Task 2: Enable the Active Directory Recycle Bin

- Use the Active Directory Administration Center to enable the Recycle Bin.

#### ► Task 3: Delete a test user

- Delete Aidan Delaney from the Managers OU.

#### ► Task 4: Restore the Deleted User

- Restore the deleted user from the Deleted Object folder.

#### ► To prepare for the next module

- When you are finished the lab, revert the virtual machines to their initial state.

**Results:** After completing this exercise, you will have created and viewed Active Directory snapshots, enabled the Active Directory Recycle Bin, deleted a user as a test, and used the Active Directory Administrative Center to restore a deleted user account.

## Module Review and Takeaways

### Best Practices

- When cloning VDCs, delete snapshots before copying or exporting VDCs.
- When cloning VDCs, we recommend copying disks manually if there is only one drive. We recommend Export for VMs with more than one drive or other complex customizations such as multiple NICs.
- At least one global catalog should exist in every site.
- AD DS should be at the minimum Windows Server 2008 R2 level to provide fully automatic password and SPN management for managed service accounts.
- GPOs should be backed up after any changes are made.
- Do not use volumes that contain backups of GPOs or AD DS data for other uses.

### Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Domain controller promotion fails	
Group Policy is not being applied correctly	
You have to restore a version of AD DS and do not know which backup to restore from	

### Review Question

You have a mixture of client computers running Windows XP and Windows 8. After you configure several settings in the Administrative Templates and Preferences of a GPO, Windows XP users report that some settings are being applied while others are not.

### Real-world Issues and Scenarios

You have a large company with multiple branch offices. Some branch offices have fast, redundant connections while others have slow, unreliable connections.

When you have branch offices across WAN links, what solutions are available to facilitate client logons in the branch offices?

What if security is a concern?

What can you do to help prevent network interruptions from preventing users from logging on?

**Tools**

Tool	Use	Location
Server Manager	A central location for all aspects of server management	Open by default on logon or can be accessed from the task bar
Active Directory Users and Computers Active Directory Sites and Services Active Directory Domains and Trusts	Control all aspects of Active Directory management	Can be accessed from the <b>Tools</b> drop-down menu in Server Manager
GPMC	Control all aspects of Group Policy management	Can be accessed from the <b>Tools</b> drop-down menu in Server Manager
Active Directory Best Practices Analyzer	Can detect best practices violations and provide help implement best practices.	Server Manager Dashboard
Active Directory Recycle Bin	Restore object that were deleted in error from AD DS.	Can be accessed from the Active Directory Administration Center

**MCT USE ONLY. STUDENT USE PROHIBITED**



# Module 12

## Implementing Active Directory Federation Services

### Contents:

Module Overview	12-1
<b>Lesson 1:</b> Overview of Active Directory Federation Services	12-2
<b>Lesson 2:</b> Deploying Active Directory Federation Services	12-11
<b>Lesson 3:</b> Implementing AD FS for a Single Organization	12-17
<b>Lesson 4:</b> Deploying AD FS in a Business to Business Federation Scenario	12-23
<b>Lab:</b> Implementing AD FS	12-28
Module Review and Takeaways	12-36

## Module Overview

Active Directory® Federation Services (AD FS) in Windows Server® 2012 provides flexibility for organizations that want to enable their users to log on to applications that may be located on a local network, at a partner company, or in an online service. AD FS enables an organization to manage its own user accounts, and users only have to remember one set of credentials. However, those credentials can be used to provide access to a variety of applications, located in a variety of locations.

This module provides an overview of AD FS, and details how to configure AD FS in both a single organization scenario and in a partner organization scenario.

### Objectives

- Describe the identity-federation business scenarios, and how you can use AD FS to address the scenarios.
- Configure the AD FS prerequisites, and deploy the AD FS services.
- Implement AD FS to enable SSO in a single organization.
- Implement AD FS to enable SSO between federated partners.

## Lesson 1

# Overview of Active Directory Federation Services

AD FS is the Microsoft® implementation of an identity federation framework that enables organizations to establish federation trusts and share resources across organizational boundaries. AD FS is compliant with common web-services standards to enable interoperability with other identity federation implementations.

AD FS is designed to address a variety of business scenarios, where the typical authentication mechanisms used in a single organization do not work. This lesson provides an overview of the concepts and standards that are implemented in AD FS, and also provides an overview of the business scenarios that you can address with AD FS.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe identity federation.
- Describe claims-based authentication.
- Describe web services.
- Describe AD FS.
- Explain how AD FS enables SSO within a single organization.
- Explain how AD FS enables SSO between business partners.
- Explain how AD FS enables SSO between on-premises and cloud-based services.

### What Is Identity Federation?

Identity federation enables the distribution of identification, authentication, and authorization across organizational and platform boundaries. You can implement identity federation within a single organization to enable access to diverse web applications, or between two organizations that have a relationship of trust between them.

To establish an identity federation partnership, both partners agree to create a federated trust relationship. This federated trust is based on an ongoing business relationship, and enables the organizations to implement business processes identified in the business relationship.

#### Identity Federation:

- Enables distributed identification, authentication, and authorization across organizational and platform boundaries.
- Requires a federated trust relationship between two organizations or entities.
- Enables organizations to retain control over who can access resources.
- Enables organizations to retain control of their user and group accounts.



**Note:** A federated trust is not the same as a forest trust that organizations can configure between forests in Active Directory® Domain Services (AD DS). In a federated trust, the AD FS servers in two organizations never have to communicate directly with each other.

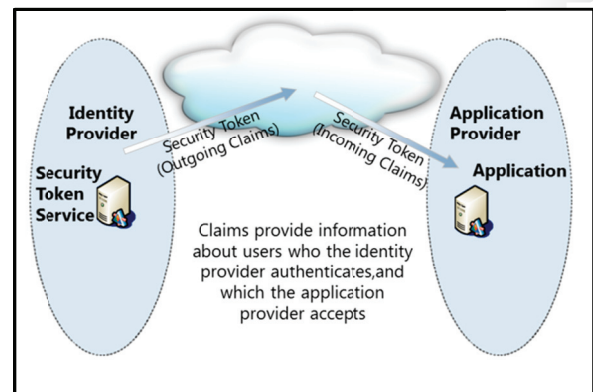
As a part of the federated trust, each partner defines what resources are accessible to the other organization, and how to enable access to the resources. For example, to update a sales forecast, a sales representative may need to collect information from a supplier's database that is hosted on the supplier's

network. The domain administrator for the sales representative is responsible for ensuring that the appropriate sales representatives are members of the group that requires access to the supplier's database. The administrator of the organization in which the database is located is responsible for ensuring that the partner's employees only have access to the data that they require.

In an identity federation solution, user identities and their associated credentials are stored, owned, and managed by the organization in which the user is located. As part of the identity federation trust, each organization also defines how the user identities are shared securely to restrict access to resources. Each partner must define the services that it makes available to trusted partners and customers, and also define which other organizations and users it trusts, what types of credentials and requests it accepts, and its privacy policies, to ensure that private information is not accessible across the trust.

## What is Claims-Based Identity?

Claims-based authentication addresses issues with extending typical authentication and authorization mechanisms outside the boundaries associated with that mechanism. For example, in most organizations, users are authenticated by an AD DS domain controller when they log on to the network. If the user provides the right credentials to the domain controller, the user is granted a security token. Applications that are running on servers in the same AD DS environment trust the security tokens that the AD DS domain controllers provide. This is because the servers can communicate with the same domain controllers where the users authenticated.



The problem with this authentication is that it does not extend easily outside the boundaries of the AD DS forest. Although it is possible to implement Kerberos or NTLM-based trusts between two AD DS forests, servers on both sides of the trust must communicate with domain controllers in the other forest to make authentication and authorization decisions. The problem becomes even more complicated when users have to access resources hosted in cloud-based systems, such as Microsoft Azure™ or Microsoft Office 365.

Claims-based authentication provides a mechanism for separating user authentication and authorization from individual applications. With claims-based authentication, users can authenticate to a directory service in their organization, and be granted a claim based on that authentication. The claim then can be presented to an application that is running in a different organization. The application is designed to enable user access to the information or features based on the claims presented.

The claim used in claims based authentication is just a statement about a user that is defined in one organization or technology and trusted in another organization or technology. The claim could include a variety of information. For example, the claim could define the user's e-mail address, user principal name (UPN), and information about all of the groups to which the user belongs. This information is collected from the authentication mechanism when the user authenticates successfully.

The organization that manages the application defines what types of claims the application will accept. For example, the application may require the user's email address to verify the user identity, and also use the group membership presented inside the claim to determine what level of access the user should have within the application.

## Web Services Overview

For claims-based authentication to work, organizations have to agree on the format for exchanging claims. Rather than have each business define this format, a set of specifications have been developed that any organization can use if it wants to implement a federated identity solution. This set of specifications is identified broadly as web services.

Web services are the set of specifications that an enterprise uses for building connected applications and services, whose functionality and interfaces are exposed to potential users through web-technology standards. These standards can include Extensible Markup Language (XML), Simple Object Access Protocol (SOAP), Web Services Description Language (WSDL), and HTTP. The goal for creating web applications by using web services is to simplify interoperability for applications across multiple development platforms, technologies, and networks.

To enhance interoperability, a set of industry standards defines web services, which are based on the following standards:

- Most web services use XML to transmit data through HTTP. XML enables developers to create their own customized tags, enabling the definition, transmission, validation, and interpretation of data between applications and organizations.
- Web services expose useful functionality to web users through a standard web protocol. In most cases, the protocol used is SOAP. SOAP is the communications protocol for XML web services. SOAP is a specification that defines the XML format for messages. Essentially, it describes what a valid XML document looks like.
- Web services provide a way to describe their interfaces in enough detail to enable a user to build a client application to communicate with the service. This description is usually provided in an XML document called a WSDL document. In other words, a WSDL file is an XML document that describes a set of SOAP messages and how the messages are exchanged.
- Web services are registered so that potential users can find them easily. This is done with Universal Discovery Description and Integration (UDDI). A UDDI directory entry is an XML file that describes a business and the services it offers.

Web services use a set of open specifications to develop applications that can interoperate across boundaries

**Web services:**

- Are developed using industry standards such as XML, SOAP, WSDL, and UDDI
- Define the security specifications used by Identity Federation systems
- Define the SAML standard for exchanging claims between federation partners

### WS-\* Security Specifications

There are many components included in web-services specifications (also known as “WS-\* specifications”). However, the most relevant specifications for an AD FS environment are the WS-Security specifications. The specifications that are part of the Web Service Security specifications include the following:

- **WS-Security.** WS-Security describes enhancements to SOAP messaging to provide quality of protection through message integrity, message confidentiality, and single message authentication. WS-Security also provides a general-purpose, but extensible, mechanism for associating security tokens with messages and how to encode binary security tokens—specifically X.509 certificates and Kerberos tickets—in SOAP messages.
- **WS-Trust.** WS-Trust defines extensions that build on WS-Security to request and issue security tokens and manage trust relationships.
- **WS-Federation.** WS-Federation defines mechanisms that WS Security can use to enable identity, attribute, authentication, and authorization federation across different trust realms.

- **WS-Federation Passive Requestor Profile.** This WS-Security extension describes how passive clients, such as web browsers, can be authenticated and authorized, and how the clients can submit claims in a federation scenario. Passive requestors of this profile are limited to the HTTP or HTTPS protocol.
- **WS-Federation Active Requestor Profile.** This WS-Security extension describes how active clients, such as SOAP-based mobile device applications, can be authenticated and authorized, and how the clients can submit claims in a federation scenario.

### Security Assertion Markup Language

The Security Assertion Markup Language (SAML) is an XML-based standard for exchanging claims between an identity provider and a service or application provider. SAML assumes that a user has been authenticated by an identity provider, and that the identity provider has populated the appropriate claim information in the security token. When the user is authenticated, the Identity Provider passes a SAML assertion to the service provider. On the basis of this assertion, the service provider can make authorization and personalization decisions within an application. The communication between federated servers is based around an XML document storing the X.509 certificate for token-signing, and the SAML 1.1 token.

### What Is AD FS?

AD FS is the Microsoft implementation of an identity-federation solution that can use claims based authentication. AD FS provides the mechanisms to implement both the identify-provider and service-provider components in an identity-federation deployment.


AD FS provides the following features:

- **Enterprise claims provider for claims-based applications:** You can configure an AD FS server as a claims provider, which means that the server can issue claims about authenticated users. This enables an organization to provide its users with access to claims-aware applications in another organization by using SSO.
- **Federation Service for identity federation across domains:** This service offers federated web SSO across domains. This enhances security and reduces overhead for IT administrators.

AD FS is the Microsoft identity federation solution that can use claims-based authentication

AD FS includes the following features:

- Web SSO
- Web services interoperability
- Support for passive and smart clients
- Extensible architecture
- Enhanced security

 **Note:** The Windows Server 2012 version of AD FS is built on AD FS version 2.0, which was the second generation of AD FS that Microsoft released. The first version, AD FS 1.0, required AD FS web agents to be installed on all web servers that were using AD FS, and provided both claims aware and NT token-based authentication. AD FS 1.0 did not support active clients or SAML.

### AD FS Features

The following are some of the key AD FS features:

- **Web SSO.** Many organizations have deployed AD DS. After authenticating to AD DS through authentication that integrates with Windows users can access all other resources that they have permission to access within the AD DS forest boundaries. AD FS extends this capability to Internet-facing applications, enabling customers, partners, and suppliers to have a similar, streamlined user experience when they access an organization's web-based applications.

- **Web Services interoperability.** AD FS is compatible with the web services specifications. AD FS employs the federation specification of WS-\*, called WS-Federation. WS-Federation makes it possible for environments that do not use the Windows identity model to federate with Windows environments.
- **Passive and smart client support.** Because AD FS is based on the WS-\* architecture, it supports federated communications between any WS-enabled endpoints, including communications between servers and passive clients, such as browsers. AD FS on Windows Server 2012 also enables access for SOAP-based smart clients, such as servers, mobile phones, personal digital assistants (PDAs), and desktop applications. AD FS implements the WS-Federation Passive Requestor Profile and WS-Federation Active Requestor Profile standards for client support.
- **Extensible architecture.** AD FS provides an extensible architecture that supports various security token types, including SAML and Kerberos authentication, as well as the ability to perform custom claims transformations. For example, AD FS can convert from one token type to another or add custom business logic as a variable in an access request. Organizations can use this extensibility to modify AD FS to coexist with their current security infrastructure and business policies.
- **Enhanced security.** AD FS also increases the security of federated solutions by delegating responsibility of account management to the organization closest to the user. Each individual organization in a federation continues to manage its own identities, and is capable of securely sharing and accepting identities and credentials from other members' sources.

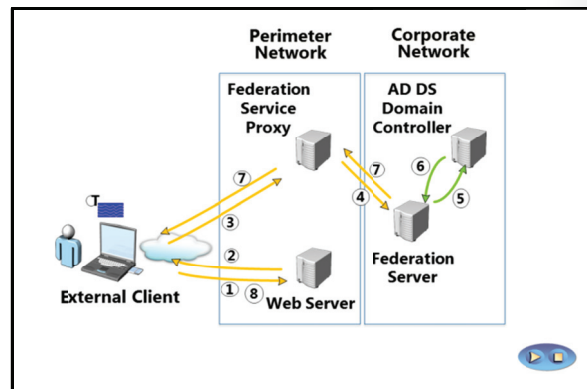



**Additional Reading:** For information on the different identity federation products that can interoperate with AD FS, and for step by step guides on how to configure the products, see the AD FS 2.0 Step-by-Step and How To Guides, located at <http://technet.microsoft.com/en-us/library/adfs2-step-by-step-guides%28v=ws.10%29.aspx>.

## How AD FS Enables SSO in a Single Organization


For many organizations, configuring access to applications and services may not require an AD FS deployment. If all users are members of the same AD DS forest, and if all applications are running on servers that are members of the same forest, you typically can use AD DS authentication to provide application access. However, there are several scenarios in which you can use AD FS, and enable SSO, to optimize the user experience, including:

- The applications may not be running on Windows servers or on any servers that support AD DS authentication. The applications may require SAML or web services for authentication and authorization.
- Large organizations frequently have multiple domains and forests that may be the results of mergers and acquisitions. Users in multiple forests might require access to the same applications.
- Users from outside the office might require access to applications that are running on internal servers. The external users may be logging on to the applications from computers that are not part of the internal domain.



 **Note:** Implementing AD FS does not necessarily mean that users are not prompted for authentication when they access applications. Depending on the scenario, users may be prompted for their credentials. However, the key point is that users always authenticate by using their internal credentials. They never have to remember alternate credentials for the application.

Organizations can use AD FS to enable SSO in these scenarios. Because all users and the application are in the same organization, the organization only has to deploy a single federation server. This server can operate as the claims provider so that it authenticates user requests and issues the claims. The same server also is the relying provider, or the consumer of the claims to provide authorization for application access.

 **Note:** The slide and the following description use the terms Federation Server and Federation Service Proxy to describe AD FS server roles. The Federation Server is responsible for issuing claims, and in this scenario, also is responsible for consuming the claims. The Federation Service Proxy is a proxy component that we recommend is used in a deployment where users outside the network need to access the AD FS environment. The next lesson covers these components in more detail.

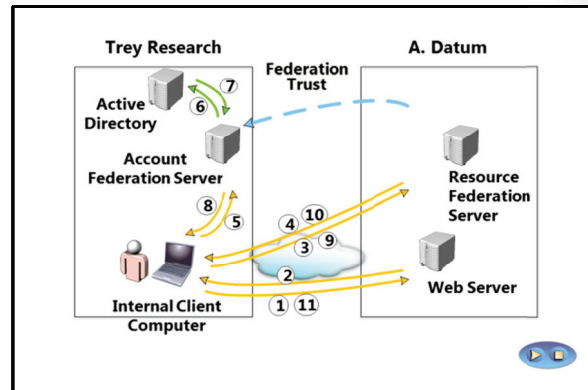
The following steps describe the communication flow in this scenario:

1. The client computer, which is located outside the network, must access a web-based application on the web server. The client computer sends an HTTPS request to the web server.
2. The web server receives the request, and identifies that the client computer does not have a claim. The web server redirects the client computer to the Federation Service proxy.
3. The client computer sends an HTTPS request to the Federation Service proxy. Depending on the scenario, the Federation Service proxy may prompt the user for authentication or use Windows Integrated authentication to collect the user credentials.
4. The Federation Service proxy passes the request and the credentials to Federation Server.
5. The Federation Server uses AD DS to authenticate the user.
6. If authentication is successful, the federation server collects AD DS information about the user, which is used to generate the user's claims.
7. If the authentication is successful, the authentication information and other information is collected in a security token and passed back to the client computer, through the Federation Service proxy.
8. The client presents the token to the web server. The web resource receives the request, validates the signed tokens, and uses the claims in the user's token to provide access to the application.



## How AD FS Enables SSO in a Business-to Business-Federation

One of the most common scenarios for deploying AD FS is to provide SSO in a business-to-business (B2B) federation. In the scenario, the organization that requires access to another organization's application or service can manage their own user accounts, and define their own authentication mechanisms. The other organization can define what applications and services are exposed to users outside the organization and what claims it accepts to provide application access. To enable application or service sharing in this scenario, the organizations just have to establish a federation trust, and then define the rules for exchange claims between the two organizations.



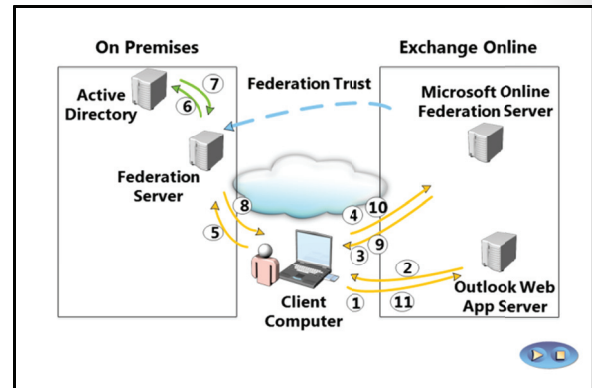
The slide above shows the flow of traffic in a federated B2B scenario using a claims-aware web application. In this scenario, users at Trey Research have to access a web-based application at A. Datum. The AD FS authentication process follows these steps:

1. A user at Trey Research, using a web browser, establishes an HTTPS connection to the web server at A. Datum.
2. The web application receives the request, and then verifies that the user does not have a valid token stored in a web browser cookie. Because the user is not authenticated, the web application redirects the client to the federation server at A. Datum, by using an HTTP 302 redirect message.
3. The client computer sends an HTTPS request to the A. Datum's federation server. The federation server determines the user's home realm. In this case, the home realm is Trey Research.
4. The client computer is redirected again to the federation server in the user's home realm, Trey Research.
5. The client computer sends an HTTPS request to the Trey Research federation server.
6. If the client computer is logged on to the domain already, the federation server can take the user's Kerberos ticket, and then request authentication from AD DS on the user's behalf, by using Windows Integrated Authentication.
7. The AD DS domain controller authenticates the user, and sends the success message back to the federation server, along with other information about the user that the federation server can use to generate the user's claims.
8. The federation server creates the claim for the user based on the rules defined for the federation partner. The claims data is placed in a digitally signed security token, and then sent to the client computer. The client computer then posts it back to the A. Datum's federation server.
9. A. Datum's federation server validates that the security token came from a trusted federation partner.
10. A. Datum's federation server creates and signs a new token, which it sends to the client computer. The client computer then sends the token back to the original URL requested.
11. The application on the web server receives the request, and validates the signed tokens. The web server issues the client a session cookie that indicates that it has authenticated successfully. The federation server issues a file-based persistent cookie (good for 30 days by default) to eliminate the home-realm discovery step during the cookie lifetime. The application then provides access to the application, based on the claims that the user provides.



## How AD FS Enables SSO with Online Services

As organizations move services and applications to cloud-based services, it is increasingly important that these organizations have some way to simplify the authentication and authorization experience for their users as they consume the cloud-based services. Cloud-based services add another level of complexity to the IT environment, as those services are located outside the direct administrative control of the IT administrators, and the services may be running on many different platforms.



You can use AD FS to provide an SSO experience to users across the various cloud-based platforms available. For example, once users are authenticated with AD DS credentials, they then could access Microsoft Online Services, such as hosted Microsoft Exchange Online or Microsoft SharePoint® Online, by using those domain credentials. AD FS also provides single sign-on to non-Microsoft cloud providers. Because AD FS is based on open standards, AD FS can interoperate with any compliant claims-based system.

The process for accessing a cloud-based application is quite similar to the B2B scenario. One example of a cloud-based service that uses AD FS for authentication is a hybrid Exchange Online deployment. In this type of deployment, an organization has deployed some or all of their mailboxes in an Office 365 Exchange Online environment. However, the organization manages all of their user accounts in their on-premises AD DS environment. The deployment uses the Microsoft Online Services Directory Synchronization tool to synchronize user-account information from the on-premises deployment to the Exchange Online deployment.

When users try to log on to their Exchange Online mailbox, the user must be authenticated by using their internal AD DS credentials. If the user tries to logon directly to the Exchange Online environment, they are redirected back to the internal AD FS deployment to authenticate before the user is given access.

The following steps describe how a user tries to access their online mailbox by using a web browser:

1. The user opens a web browser, and then sends an HTTPS request to the Exchange Online Outlook Web App server.
2. The Outlook Web App server receives the request, and then verifies that the user is part of a hybrid Exchange Server deployment. If this is the case, the server redirects the client computer to the Microsoft Online federation server.
3. The client computer sends an HTTPS request to the Microsoft Online federation server.
4. The client computer is redirected again to the on-premises federation server.
5. The client computer sends an HTTPS request to the on-premises federation server.
6. If the client computer is logged on to the domain already, the federation server can take the user's Kerberos ticket, and then request authentication from AD DS on the user's behalf, by using Windows Integrated Authentication. If the user is logging on from outside the network, or from a computer that is not a member of the internal domain, the user is prompted for credentials.
7. The AD DS domain controller authenticates the user, and sends the success message back to the federation server, along with other information about the user that can be used to generate the user's claims.

8. The federation server creates the claim for the user, based on the rules that are defined during the AD FS server setup. The claims data is placed in a digitally signed security token, and then sent to the client computer. The client computer then posts it back to the Microsoft Online federation server.
9. The Microsoft Online federation server validates that the security token came from a trusted federation partner. This trust is configured when you configure the hybrid Exchange environment.
10. The Microsoft Online federation server creates and signs a new token, which it sends to the client computer. The client computer then sends the token back to the Outlook Web App server.
11. The Outlook Web App server receives the request and validates the signed tokens. The server issues the client a session cookie, which indicates that it has successfully authenticated. The user then is granted access to their Exchange server mailbox.

## Lesson 2

# Deploying Active Directory Federation Services

Now that you have an understanding of how AD FS works, the next step is deploying the service. Before deploying AD FS, you must understand the components that you deploy, and the prerequisites that you must meet, especially with regards to certificates. This lesson provides an overview of deploying the AD FS server role in Windows Server 2012.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe the components that an AD FS deployment can include.
- List the prerequisites for an AD FS deployment.
- Describe the Public Key Infrastructure (PKI) and certificate requirements for an AD FS deployment.
- Describe the AD FS federation server roles.
- Install the AD FS server role.

### AD FS Components

AD FS is installed as a server role in Windows Server 2012. However, there are many different components that you can install and configure in an AD FS deployment. The following table lists the AD FS components.

- Federation Server
- Federation Server Proxy
- Claims
- Claim Rules
- Attribute Store
- Claims Providers
- Relying Parties
- Claims Provider Trust
- Relying Party Trust
- Certificates
- Endpoints

Component	What does it do?
Federation Server	The federation server issues, manages, and validates requests that involve identity claims. All implementations of AD FS require at least one Federation Service.
Federation Server Proxy	The Federation Server proxy is an optional component that typically is deployed in a perimeter network. The Federation Server proxy does not add any functionality to the AD FS deployment, but is deployed just to provide a layer of security for connections from the Internet to the Federation Server.
Claims	A claim is a statement that one object makes about another object, such as a user. The claim could include the user's name, job title, or any other factor that might be used in an authentication scenario.
Claim Rules	Claim rules determine how federation servers process claims. For example, a claim rule may state that an email address is accepted as a valid claim, or that a group name from one organization is translated into an application-specific role in the other organization. The rules usually are processed in real time, as claims are made.

Component	What does it do?
Attribute Store	An attribute store is used by AD FS to look up claim values. AD DS is a common attribute store, and is available by default if AD FS is installed on a domain-joined server.
Claims Providers	A claims provider enables one side of the AD FS authentication and authorization process. The claims provider manages the user authentication, and then issues the claims that the user presents to a relying party.
Relying Parties	The relying party enables the second side of the AD FS authentication and authorization process. The relying party is a web service that consumes claims from the claims provider. The relying party server must have the Windows Identity Foundation (WIF) installed or use AD FS 1.0's claims-aware agent.
Claims Provider Trust	This is configuration data that defines rules under which a client may request claims from a claims provider and subsequently submit them to a relying party. The trust consists of various identifiers, such as names, groups and various rules.
Relying Party Trust	This is the AD FS configuration data that is used to provide claims about a user or client to a relying party. It consists of various identifiers, such as names, groups, and various rules.
Certificates	AD FS uses digital certificates when communicating over SSL or as part of the token-issuing process, the token-receiving process, and the metadata-publishing process.
Endpoints	Endpoints are mechanisms that enable access to the AD FS technologies, including token issuance and metadata publishing. AD FS comes with built-in endpoints that are responsible for a specific functionality.



**Note:** Many of these components are described in more detail in the remainder of this module.

## AD FS Prerequisites

Before deploying AD FS, you must ensure that your internal network meets some basic prerequisites. The configuration of the following network services is critical for a successful AD FS deployment:

- Network connectivity: TCP/IP connectivity must exist between:
  - The client computer
  - A domain controller
  - Federation Service server
  - Federation Service Proxy server (when applicable)
  - An application server that is integrated with AD FS
  - Web server running the AD FS Web Agent (AD FS v1 only)

Infrastructure critical to a successful AD FS deployment include:

- TCP/IP network connectivity
- AD DS
- Attribute stores
- DNS
- Compatible operating systems



- **AD DS:** AD DS is a critical piece of AD FS. Domain controllers should be running Windows Server 2003 Service Pack 1 (SP1) at a minimum. In both AD FS v1 and AD FS, federation servers must be joined to an AD DS domain. The Federation Service proxy does not have to be domain-joined. In fact, we recommend that this component be installed on a workgroup-joined computer as a security best practice. Although you can install AD FS on a domain controller, we do not recommend this due to security implications.
- **Attribute stores.** AD FS uses an attribute store to build claim information. The attribute store contains information about users – this information is extracted from the store by the AD FS server after the user has been authenticated. AD FS supports the following attribute stores:
  - Active Directory Application Mode (ADAM) in Windows Server 2003
  - Active Directory Lightweight Directory Services (AD LDS) in Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012
  - Microsoft SQL Server 2005 (all editions)
  - Microsoft SQL Server 2008 (all editions)
  - A custom attribute store



**Note:** AD DS can be used both as the authentication provider and as an attribute store. AD FS also can use AD LDS as an attribute store. In AD FS v1, you can use AD LDS as an authentication store, but in the current version of AD FS, you only can use AD LDS as an attribute store.

- **Domain Name System (DNS):** Name resolution allows clients to find federation servers. The client computers must resolve the DNS names for all federation servers that they connect to, as well as the web applications that the client computer is trying to use. If the client computer is external to the network, the client computer must resolve the DNS name for the federation service proxy, not the internal federation server. The Federation Service proxy must resolve the name of the internal federation server. If internal users have to access the internal federation server directly, and external users have to connect through the Federation Server proxy, you require a split DNS.
- **Operating-system prerequisites:** You can only deploy the Windows Server 2012 version of AD FS as a server role on a Windows Server 2012 server. AD FS 2.0, which is almost identical to the Windows Server 2012 version, can be installed on Windows Server 2008 Service Pack 2 (SP2) or Windows Server 2008 R2.

## PKI and Certificate Requirements

AD FS is designed to enable computers to communicate securely, even though they may be located in different locations. In this scenario, most of the communications between computers passes through the Internet. To provide security for the network traffic, all communications are protected by using SSL. This factor means that it is important to choose and assign SSL certificates correctly to the AD FS servers. To provide SSL security, AD FS servers use certificates in the following three ways.

- AD FS federation services require:
  - Service Communication Certificates
  - Token-Signing Certificates
  - Token-Decrypting Certificates
- When choosing certificates, ensure that the Service Communication Certificate and the Token-Signing Certificate are trusted by all federation partners and clients

## Service Communication Certificates

This certificate is used to secure SSL communications to the websites running on the AD FS server and is bound to the default web site on the AD FS server. You can choose which certificate to use when you configure the AD FS server role on the server, and can change the assigned certificate after deployment by using the AD FS management console. This certificate also is called a server authentication certificate.

## Token-Signing Certificates

The token-signing certificate is used to sign every token issued a federation server. This certificate is critical in an AD FS deployment, because the token signature indicates which federation server issued the token. The claims provider uses this certificate to identify itself, and also by the Replying Party to verify that the token is coming from a trusted Federation partner.

The relying party also requires a token-signing certificate to sign the tokens that it prepares for other AD FS components, such as web applications and clients. These tokens must be signed by the relying party's token-signing certificate in order for the destination applications to validate them.

When you configure a Federation Server, the server assigns a self-signed certificate as the token-signing certificate. Because no other parties trust the self-signed certificate, it is important that you replace the self-signed certificate with a trusted certificate. You can configure multiple token-signing certificates on the federation server, but only the primary certificate is used to sign tokens.

## Token-Decrypting Certificates

Token-decrypting certificates encrypt the entire user token before transmitting the token across the network. To provide this functionality, the relying party federation server sends the certificate to the claims provider federation server. The certificate is sent without the private key. The claims provider server uses the public key from the certificate to encrypt the user token. When the token is returned to the relying party federation server, it uses the private key from the certificate to decrypt the token. This provides an extra layer of security when transmitting the certificates across the Internet.

When you configure a Federation Server, the server assigns a self-signed certificate as the token-decrypting certificate. Because no other parties have to trust this certificate, it is possible to continue to use this certificate without replacing it with a trusted certificate.



**Note:** Federation server proxies only require a service communication certificate. The certificate is used to enable SSL communication for all client connection. Since the federation server proxy does not issue any tokens, it does not need the other two types of certificates. Web servers that are deployed as part of an AD FS deployment also should be configured with SSL server certificates to enable secure communications with client computers.

## Choosing a Certification Authority

AD FS federation servers can use self-signed certificates, certificates from an internal, private certification authority (CA), or certificates that have been purchased from an external public CA.

The most important factor when choosing the certificates in most AD FS deployments is that the certificates be trusted by all parties involved. This means that if you are configuring an AD FS deployment that interacts with other organizations, you are almost certainly going to use a public CA, because all partners trust the certificates issued by the public CA automatically.

If you are deploying AD FS just for your organization, and all servers and client computers are under your control, you can consider using a certificate from an internal private CA. If you deploy an enterprise CA on Windows Server 2012, you can use Group Policy to ensure that all computers in the organization automatically trust the certificates that the internal CA issues. Using an internal CA can decrease the cost of the certificates significantly.

 **Note:** Deploying an internal CA using Active Directory Certificate Services is very easy, but it is critical that you plan and implement the deployment carefully.


When you install the AD FS server role, the server is configured with self-signed certificates. These certificates are not trusted by any other systems, so you must replace the server communications certificate and the token-signing certificates with a trusted certificate. It is not critical that you replace the token-decrypting certificate with a trusted certificate.

## Federation Server Roles

When you deploy the AD FS server role, and configure the server, you can choose which role the server plays in an AD FS deployment. You can configure an AD FS server in one of three roles:

- **Claims Provider.** A claims provider is a federation server that provides signed tokens containing claims to users. Claims provider federation servers are deployed in organizations where user accounts are located. When a user requests a token, the claims provider federation server verifies the user authentication by using AD DS, and then collects information from an attribute store, such as AD DS or AD LDS, to populate the user claim with the attributes required by the partner organization. The server issues tokens in the Security Assertion Markup Language (SAML) format. The claims provider federation server also protects the contents of security tokens in transit by signing and optionally encrypting them.
- **Relying Party.** A relying party is a federation server that receives security tokens from a trusted claims provider. The relying party federation servers are deployed in organizations that provide application access to claims provider organizations. The relying party accepts and validates the claim, and then issues new security tokens that the web server can use to provide appropriate access to the application.

AD FS Server Role	Description
Claims Provider federation server	<ul style="list-style-type: none"> <li>• Authenticates internal users</li> <li>• Issues signed tokens containing user claims</li> </ul>
Relying Party federation server	<ul style="list-style-type: none"> <li>• Consumes tokens from the Claims Provider</li> <li>• Issues tokens for application access</li> </ul>
Federation server proxy	<ul style="list-style-type: none"> <li>• Deployed in a perimeter network</li> <li>• Provides a layer of security for internal federation servers</li> </ul>

 **Note:** A single AD FS server can operate as both a claims provider and a relying party, even with the same partner organizations. The AD FS server functions as a claims provider when it is authenticating users and providing tokens for another organization, but also can accept tokens from the same or another organization in a relying party role.

- **Federation Server Proxy.** A federation server proxy provides an extra level of security for AD FS traffic coming from the Internet to the internal AD FS federation servers. Federation server proxies can be deployed in both the claims provider and relying party organizations. On the claims provider side, the proxy collects the authentication information from client computers and passes it to the claims provider federation server for processing. The federation server issues a security token to the proxy, which sends it to the relying party proxy. The relying party federation server proxy accepts these tokens, and then passes them on to the internal federation server. The relying party federation server issues a security token for the web application, and then sends the token to the proxy, which then forwards the token to the client. The federation server proxy does not provide any tokens or create claims. It only forwards requests from clients to internal AD FS servers.





**Note:** You cannot configure a federation server proxy as a claims provider or a Relying Provider. The claims provider and Relying Provider must be members of an AD DS domain. You must configure the federation server proxy as a member of a workgroup, and then deploy it in a perimeter network.

## Demonstration: Installing the AD FS Server Role

In this demonstration, you will see how to install and complete the initial configuration of the AD FS server role in Windows Server 2012. The instructor will install the server role, and then run the AD FS Federation Server Configuration Wizard to configure the server as a standalone federation server.

### Demonstration Steps

1. On LON-DC1, in Server Manager, add the Active Directory Federation Services server role.
2. Run the AD FS Federation Server Configuration Wizard by using the following parameters:
  - a. Create a new federation services
  - b. Create a stand-alone deployment
  - c. Use the LON-DC1.Adatum certificate
  - d. Choose a service name of LON-DC1.Adatum.com
3. Open Windows Internet Explorer®, and then connect to <https://lon-dc1.adatum.com/federationmetadata/2007-06/federationmetadata.xml>.



## Lesson 3

# Implementing AD FS for a Single Organization

The simplest deployment scenario for AD FS is within a single organization. In this scenario, a single AD FS server can operate both as the claims provider and as the Relying Provider. All users in this scenario are internal to the organization, as is the application that the users are accessing.

This lesson provides details on the components that are required to configure in a single organization deployment of AD FS. These components include configuring claims, claim rules, claims provider trusts, and relying party trusts.

### Lesson Objectives

After completing this lesson, you will be able to:

- Describe AD FS claims.
- Describe AD FS claim rules.
- Describe claims provider trusts.
- Describe relying provider trusts.
- Configure claims provider and relying provider trusts.

### What are AD FS Claims?

AD FS claims provide the link between the claims provider and Relying Provider roles in an AD FS deployment. The claims provider creates the claims and the Relying Provider consumes the claims. AD FS claims provide a standards-based and flexible way for claims provider organizations to provide very specific information about users in their organizations, and a way for Relying Providers to define exactly what information they require to provide application access.

An AD FS claim is a statement made about a particular subject (such as a user) by a trusted entity (such as a claims provider). The claim information provides the details that the application requires to enable access to claims-aware applications.

- Claims are used to provide information about users from the Claims Provider to the Relying Partner
- AD FS:
  - Provides a default set of built-in claims
  - Enables the creation of custom claims
  - Requires that each claim have a unique URI
- Claims can be:
  - Retrieved from an attribute store
  - Calculated based on retrieved values
  - Transformed into alternate values

### Claim Types

Each AD FS claim has a claim type, such as Email Address, UPN, or Last Name. Users can be issued claims based on any defined claim type. So a user might be issued a claim with a type of Last Name and a value of Weber. AD FS provides several built-in claim types, or you can create new ones based on the organization requirements.



**Note:** In AD FS 1.0, you could configure claims as identity claims, group claims or custom claims. These claim types do not apply to AD FS 2.0 or later. Essentially, all claims are now considered custom claims.

Each AD FS claim type is identified by a Uniform Resource Identifier (URI) that uniquely identifies the claim type. This information is provided as part of the AD FS server metadata. For example, if the claims provider organization and the Relying Provider organization decide to use a claim type of AccountNumber, both organizations must configure a claim type with this name. The claim type is published, and the claim type URI must be identical on both AD FS servers.

### How Claim Values are Populated

The claims issued by a claims provider contain the information that is required by the relying party to enable appropriate application access. One of the first steps in planning an AD FS deployment is to define exactly what information the applications must have about each user to provide that application access. Once this information is defined, the claims are defined on the claims provider federation server. The information required to populate the claim can be obtained in several ways:

- The claim can be retrieved from an attribute store. Frequently, the information required for the claim is already stored in an attribute store that is available to the federation server. For example, an organization might decide that the claim should include the user's UPN, email address, and group memberships. This information is already stored in AD DS, so the federation server can just retrieve this information from AD DS when creating the claim. Since AD FS can use AD DS, AD LDS, Microsoft SQL Server, a third-party Lightweight Directory Access Protocol (LDAP) directory, or a custom attribute store to populate claims, you can define almost any value within the claim.
- The claim can be calculated based on collected information – claims provider federation servers can also calculate information based on information gathered from an attribute store. For example, you may want to provide information about a person's salary within a claim. This information is likely stored in a Human Resources database, but the actual value may be considered confidential. You can define a claim that categorizes salaries within an organization, and then have the AD FS server calculate which category a specific user belongs to. In this way, the claim only includes the salary category information, not the actual user salary.
- The claim can be transformed from one value to another. In some cases, the information stored in an attribute store does not exactly match the information that the application requires when making authorization information. For example, the application may have different user roles defined that do not directly match the attributes stored in any attribute store. However, the application role may correlate to AD DS group membership. For example, users in the Sales group may correlate to one application role, while users in the Sales Management group may correlate to a different application role. To establish the correlation in AD FS, you can configure a claims transformation that takes the value that the claims provider provides, and translates the value to a claim that is useful to the relying party's application.

### What Are AD FS Claim Rules?

Claims rules define how AD FS servers send and consume claims. Claims rules define the business logic that is applied to claims that claims providers provide, and to claims that the relying parties accept. You can use claim rules to:

- Define which incoming claims are accepted from one or more claims providers.
- Define which outbound claims are provided to one or more relying parties.

- Claims rules define how claims are sent and consumed by AD FS servers
- Claims provider rules are acceptance transform rules
- Relying party rules can be:
  - Issuance transform rules
  - Issuance authorization rules
  - Delegation authorization rules
- AD FS servers provide default claims rules, templates and a syntax for creating claims rules

- Apply authorization rules to enable access to a specific relying party for one or more users or groups of users.

You can define two types of claim rules:

- Claim rules for a claims provider trust. A claims provider trust is the AD FS trust relationship configured between an AD FS server and a claims provider. You can configure claim rules to define how the claims provider processes and issues claims.
- Claim rules for a relying party trust. A relying party trust is the AD FS trust relationship configured between an AD FS server and a relying party. You can configure claim rules that define how the relying party accepts claims from the claims provider.

Claims rules on an AD FS claims provider are all considered acceptance transform rules. These rules determine what types of claims are accepted from the claims provider and then sent to a relying party trust. When configuring AD FS within a single organization, there is a default claims provider trust configured with the local AD DS domain, so this rule set defines the claims that are accepted from AD DS.

There are three types of claim rules for a relying party trust:

- Issuance Transform Rules: These rules define the claims that are sent to the relying party that has been defined in the relying party trust.
- Issuance Authorization Rules: These rules define which users are permitted or denied access to the relying party that has been defined in the relying party trust. This rule set can include rules that explicitly permit access to a relying party, and/or rules that explicitly deny access to a relying party.
- Delegation Authorization Rules: These rules define the claims that specify which users can act on behalf of other users when accessing the relying party. This rule set can include rules that explicitly permit delegates for a relying party, or rules that explicitly deny delegates to a relying party.

A single claim rule associated with a single federated trust relationship. This means that you cannot create a set of rules for one trust and then reuse those rules for other trusts that you configure on your federation server.

AD FS servers are preconfigured with a set of default rules, as well as several default templates that you can use to create the most common claims rules. You can also create custom claim rules by using the AD FS claim rule language.

## What Is a Claims Provider Trust?

You configure a claims provider trust on the relying party federation server. The claims provider trust identifies the claims provider, and also describes how the relying party consumes the claims that the claims provider issues. You must configure a claims provider trust for each claims provider.

By default, an AD FS server is configured with a claims provider trust named Active Directory. This trust defines the claim rules, which are all acceptance transform rules that define how the AD FS server accepts AD DS credentials. For example, the default claim rules on the claims provider trust include rules that pass the user names, SIDs, and group SIDs to the relying party. In a single-organization AD FS deployment, where AD DS authenticates all users, the default claims provider trust may be the only required claims provider trust.

- Claims provider trusts:
  - Are configured on the relying party federation server
  - Identify the claims provider
  - Configure the claims rules for the claims provider
- In a single organization scenario, a claims provider trust called Active Directory defines how AD DS user credentials are processed
- Additional claims provider trusts can be configured:
  - By importing the federation metadata
  - By importing a configuration file
  - By manually configuring the trust

When you expand the AD FS deployment to include other organizations, you must create additional claims provider trusts for each federated organization. You have three options when configuring a claims provider trust:

- Import data about the claims provider through the federation metadata. If the AD FS federation server or federation proxy server is accessible through the network from your AD FS federation server, you can enter the host name or URL for the partner federation server. Your AD FS connects to the partner server, and downloads the federation metadata from the server. The federation metadata includes all information required to configure the claims provider trust. As part of the federation metadata download, your federation server also downloads the SSL certificate that the partner federation server uses.
- Import data about the claims provider from a file. Use this option if the partner federation server is not directly accessible from your federation server, but where the partner organization has exported its configuration, and then provided you the information in a file. The configuration file must include the configuration information for the partner organization, as well as the SSL certificate that the partner federation server uses.
- Manually configure the claims provider trust. Use this option if you want to configure all of the settings for the claims provider trust directly. When you choose this option, you must provide the features that the claims provider supports, as well as the URL used to access the claims provider AD FS servers. Furthermore, you must add the SSL certificate that the partner organization uses.

## What Is a Relying Party Trust?

A relying party trust is defined on the claims provider federation server. The relying party trust identifies the relying party, and also defines the claims rules that define how the relying party accepts and process claims from the claims provider.

In a single-organization scenario, the relying party trust defines how the AD FS server interacts with the applications deployed within the application. When you configure the relying party trust in a single organization, you provide the URL for the internal application and configure settings such whether the application supports SAML 2.0 or whether it requires AD FS 1.0 tokens, the SSL certificate and URL used by the web server, and the application's issuance-authorization rules.

The process for configuring relying party trust is very similar to the claims provider trust. When you expand the AD FS deployment to include other organizations, you must create additional relying party trusts for each federated organization. You have three options when configuring a relying party trust:

- Import data about the relying party through the federation metadata. If the AD FS federation server or federation proxy server is accessible through the network from your AD FS federation server, you can enter the host name or URL for the partner federation server. Your AD FS connects to the partner server, and downloads the federation metadata from the server. The federation metadata includes all the information required to configure the relying party trust. As part of the federation metadata download, your federation server also downloads the SSL certificate that the partner federation server uses.

- Relying party trusts:
  - Are configured on the claims provider federation server
  - Identify the relying party
  - Configure the claims rules for the relying party
- In a single organization scenario, a relying party trust defines the connection to internal applications
- Additional relying party trusts can be configured:
  - By importing the federation metadata
  - By importing a configuration file
  - By manually configuring the trust

- Import data about the relying party from a file. Use this option if the partner federation server is not directly accessible from your federation server, but where the partner organization has exported its configuration and provided you the information in a file. The configuration file must include the configuration information for the partner organization, as well as the SSL certificate that the partner federation server uses.

Manually configure the claims provider trust. Use this option if you want to configure all of the settings for the claims provide trust directly.

## Demonstration: Configuring Claims Provider and Relying Party Trusts

In this demonstration, you will see how to configure claims provider trusts and relying party trusts. The instructor will show how to edit the default Active Directory claims provider trust, and will create a new relying party trust and show how to configure the trust.

### Demonstration Steps

1. In the AD FS 2.0 Management console, go to the claims provider **Trusts**, highlight the **Active Directory** store, and then go to **Edit Claim Rules**.
2. In the **Edit Claim Rules for Active Directory** dialog on the **Acceptance Transform Rules** tab, start the Add Transform Claim Rule Wizard, and complete the wizard with the following settings:
  - Under **Claim rule template** select **Send LDAP Attributes as Claims**.
  - Name the claim rule **Outbound LDAP Attribute Rule**.
  - Choose Active Directory as the Attribute Store.
3. In the **Mapping of LDAP attributes to outgoing claim types**, select the following values:
  - E-Mail-Addresses to E-Mail Address
  - User-Principal-Name to UPN
4. On LON-SVR1, from the Start screen, start the Windows Identity Foundation Federation Utility.
5. Complete the wizard with the following settings:
  - Point to the web.config file of the WIF sample application by pointing to **C:\inetpub\wwwroot\AdatumTestApp\web.config**.
  - Specify an **Application URI** box by typing **https://lon-svr1.adatum.com/AdatumTestApp/**.
  - Select **Use an existing STS**, and enter a path **https://lon-dc1.adatum.com/federationmetadata/2007-06/federationmetadata.xml**.
  - Disable certificate chain validation.
  - Select **No encryption**.
6. In the AD FS 2.0 Management console, in the middle pane, click **Required: Add a trusted relying party**.

7. Complete the Add relying party Wizard with the following settings:
  - Select **Import data about the relying party published online or on a local network**, and type **https://lon-svr1.adatum.com/adatumtestapp**.
  - Specify a **Display** name of **ADatum Test App**.
  - Select **Permit all users to access this relying party**.
  - Select **Permit access for all users**.
  - Select to open the **Edit Claims Rules for WIF Sample Claims App** check box when the wizard is complete.

## Lesson 4

# Deploying AD FS in a Business to Business Federation Scenario

A second common scenario for implementing AD FS is in a B2B federation scenario. In this scenario, users in one organization have to be able to access an application in another organization. AD FS in this scenario enables SSO. Users always log on to their home AD DS environment, but are granted access to the partner application based on the claims acquired from their local AD FS server.

Configuring AD FS in a B2B federation scenario is quite similar to configuring AD FS in a single organization scenario. The primary difference is that now the claims provider trusts and the relying provider trusts refer to external organizations rather than internal AD DS or application.

This lesson describes how to configure AD FS in a B2B scenario.

### Lesson Objectives

After completing this lesson, you will be able to:

- Configure the account partner in a B2B federation scenario.
- Configure the resource partner in a B2B federation scenario.
- Describe how claims transformations work.
- Describe how home-realm discovery works.
- Configure claims rules.

### Configuring an Account Partner

In a B2B AD FS scenario, the terminology used to describe the parties involved in the AD FS deployment changes slightly. In this scenario, the claims provider organization is also called the account partner organization. An account partner organization is the organization in which the user accounts are stored in an attribute store. An account partner handles the following tasks:

- Gathering credentials from users by using a web-based service, and then authenticating those credentials.
- Building up claims for users, and then packaging the claims into security tokens. The tokens can then be presented across a federation trust to gain access to federation resources located at the resource partner organization.

- An account partner is a claims provider in a B2B federation scenario

- To configure an account partner:

1. Implement the physical topology
2. Add an attribute store
3. Configure a relying party trust
4. Add a claim description
5. Prepare client computers for federation

Configuring the account partner organization to prepare for federation involves the following steps:

1. Implement the physical topology for the account partner deployment. This step could include deciding on the number of federation servers and federation server proxies to deploy, the locations where these will be deployed and configuring the required DNS records and certificates.
2. Add an attribute store. Use the AD FS management console to add the attribute store. In most cases, you use the default Active Directory attribute store, which also must be used for authentication. However, you also can add other attribute stores, if necessary, to build user claims.



3. Connect to a resource partner organization by creating a relying party trust. The easiest way to do this is to use the federation metadata URL that the resource partner organization provides. With this option, your AD FS server automatically collects the information that the relying party trust requires.
4. Add a claim description. The claim description lists the claims that your organization provides to the relying partner. This information may include user names, email addresses, group membership information, or other identifying information about a user.
5. Prepare client computers for federation. This may involve two steps:
  - o Add the account partner federation server to the trusted sites list in the browser of client computers. By adding the account partner federation server to the trusted sites list on the client computers, you enable Windows Integrated Authentication, which means that users are not prompted for authentication if they are already logged into the domain. You can use Group Policy objects (GPOs) to assign the URL to the trusted site.
  - o Configure certificate trusts. This is an optional step that is required only if one more of the servers accessed by the clients do not have trusted certificates. The client computer may have to connect to the account federation servers, resource federation servers or federation proxy servers, and the destination web servers. If any of these certificates are not from a trusted public CA, you may have to add the appropriate certificate or root certificate to the certificate store on the clients. You can do this by using GPOs.

## Configuring a Resource Partner

The resource partner organization is the relying party in a B2B federation scenario. The resource partner organization is where the resources exist and are made accessible to account partner organizations. The resource partner handles the following tasks:

- Accepts and validates security tokens that the account-partner federation server produces.
- Consumes the claims from the security tokens, and then provides new claims to its web servers after making an authorization decision.

- An resource partner is a relying party in a B2B federation scenario
- To configure an relying party:
  1. Implement the physical topology
  2. Add an attribute store
  3. Configure a claims provider trust
  4. Create claim rule sets for the claims provider trust

The web servers must have Windows Identity Framework (WIF) installed or have the AD FS 1.x Claims-Aware Web Agent role services installed to externalize the identity logic and accept claims.



**Note:** Microsoft offers WIF to provide a set of consistent development tools that enable developers to integrate claims-based authentication and authorization into their applications. WIF also includes a Software Development Kit (SDK) and sample applications. You use a WIF sample application in the lab for this module.

Configuring the resource partner organization is similar to configuring the account partner organization, and consists of the following steps:

1. Implement the physical topology for the resource partner deployment. The planning and implementation steps are the same as the account partner, with the addition of planning the web server location and configuration.



2. Add an attribute store. On the resource partner, the attribute store is used to populate the claims that are offered to the client, which presents them to the web server.
3. Connect to an account partner organization by creating a claims provider trust.
4. Create claim rule sets for the claims provider trust.

## Configuring Claims Rules for Business to Business Scenarios

In a single organization deployment of AD FS, it may be quite easy to design and implement claims rules. In many cases, you may need to just provide the user name or group name collected from the claim to the web server. In a B2B scenario, it is more likely that you have to configure more complicated claims rules to define user access between widely varying systems.

Claim rules define how account partners (claims providers) create claims, and how resource partners (relying parties) consume claims. AD FS provides several templates that you can use when configuring claim rules:

- Organization to organization scenarios may require more complex claims rules
- You can create claims rules by using the following templates:
  - Send LDAP attributes as claims
  - Send group membership as a claim
  - Pass through or filter an incoming claim
  - Transform an incoming claim
  - Permit or deny users based on an incoming claim
- You can also create custom rules by using the AD FS Claim Rule Language

- **Send LDAP Attribute as Claims** rule template. Use this template when you select specific attributes in an LDAP attribute store to populate claims. You can configure multiple LDAP attributes as individual claims in a single claim rule created from this template. For example, you can create a rule that extracts the **displayName** and **givenName** AD DS attributes from all authenticated users, and then send these values as outgoing claims to be sent to a relying party.
- **Send Group Membership as a Claim** rule template. Use this template to send a particular claim type and associated claim value based on the user's AD DS security group membership. For example, you might use this template to create a rule that sends a group claim type with a value of **SalesAdmin** if the user is a member of the Sales Manager security group within their AD DS domain. This rule only issues a single claim, based on the AD DS group that you select as a part of the template.
- **Pass Through or Filter an Incoming Claim** rule template. Use this template to set additional restrictions on which claims are submitted to relying parties. For example, you might want to use a user email address as a claim, but only forward the email address if the domain suffix on the email address is adatum.com. When using this template, you can either pass through whatever claim you extract from the attribute store, or you can configure rules that filter whether the claim passes through based on various criteria.
- **Transform an Incoming Claim** rule template. Use this template to map the value of an attribute in the claims provider attribute store to a different value in the relying party attribute store. For example, you may want to provide all members of the Marketing department at A. Datum limited access to a purchasing application at Trey Research. At Trey Research, the attribute used to define the limited access level may have an attribute of **LimitedPurchaser**. To address this scenario, you can configure a claims rule that transforms an outgoing claim where the Department value is Marketing to an incoming claim where the **ApplicationAccess** attribute is **LimitedPurchaser**. Rules created from this template must have a one-to-one relationship between the claim at the claims provider and the claim at the relying partner.

- **Permit or Deny Users based on an Incoming Claim** rule template. This template is available only when you are configuring Issuance Authorization Rules or Delegation Authorization Rules on a relying party Trust. Use this template to create rules that enable or deny access by users to a relying party, based on the type and value of an incoming claim. This claim rule template allows you to perform an authorization check on the claims provider before claims are even sent to a relying party. For example, you can use this rule template to create a rule that only permits users from the Sales group to access a relying party, authentication requests from members of other groups are not even sent to the relying party.

If none of the built-in claim rule templates provide the functionality that you are looking for, you can create more complex rules using the AD FS Claim Rule Language. By creating a custom rule, you can extract claims information from multiple attribute stores and also combine claim types into a single claim rule.

## How Home Realm Discovery Works

Some resource partner organizations hosting claims-aware applications may want to enable multiple account partners to access the applications. In this scenario, when users connect to the web application, there must be some mechanism for directing the users to the AD FS federation server in their home domain rather than to another organization's federation server. The process for directing clients to the appropriate account partner is called home realm discovery.

Home realm discovery occurs after the client connects to the relying parties web site and the client has been redirected to the relying party's federation server. At this point, the relying party's federation server must redirect the client to the Federation Server in the client's home realm, so that the user can be authenticated. If there are multiple claims providers configured on the relying party federation server, it has to know to which federation server to redirect the client.

At a high level, there are three main ways implement home realm discovery:

1. Ask users to select their home realm. With this option, when the user is redirected to the relying party's federation server, the federation server can display a web page that requests that the user identify the company they work for. Once the user selects the appropriate company, the federation server uses that information to redirect the client computer to the appropriate home federation server for authentication.
2. Modify the link for the web application to include a "Whr" string that specifies the user's home realm. The relying party's Federation Server uses this string to automatically redirect the user to the appropriate home realm. This means that the user does not have to be prompted to select the home realm, because the "Whr" string in the URL that the user clicks relays the needed information to the relying party's Federation Server. The modified link might look something like <https://www.adatum.com/OrderApp/?whr=urn:federation:TreyResearch>.

- Home realm discovery is required on the resource partner when it has configured AD FS federations with account partners
- To enable home realm discovery, you can:
  - Prompt the user for home realm information
  - Modify the URL for the web application to specify the home realm
  - Configure a SAML profile called IdPInitiated SSO to direct users to the account partner site first

3. If the remote application is SAML 2.0-compliant, users can use a SAML profile called IdPInitiated SSO. This SAML profile configures users to access their local claims provider first, which can prepare the user's token with the claims required to access the partner web application. This process changes the normal process for accessing the web application by having the users log on to the claims provider federation server first, and then prompting them to select which application they want to access so that their token can be created with the appropriate information.



**Note:** The home realm discovery process occurs the first time the user tries to access a web application. After the user successfully authenticates, a home-realm discovery cookie is issued to the client so that the user does not have to go through the process the next time. This home-realm discovery cookie expires after a month, unless the cookie cache is cleared sooner.

## Demonstration: Configuring Claims Rules

In this demonstration, you will see how to configure claims rules. You will see how to configure claims rules on a relying party trust that forwards a group name as part of the claim. You will also see how to configure a claims rule that limits access to the application only to members of a particular group.

### Demonstration Steps

1. On LON-DC1, edit the Adatum Test App relying party trust by creating a new Issuance Transform Rule that passes through or filters an incoming claim. Name the rule **Send Group Name rule**, and configure the rule to use an incoming claim type of group.
2. Delete the Issuance Authorization Rule that grants access to all users.
3. Create a new Issuance Authorization Rule that permits or denies user access based on the incoming claim. Configure the rule with the name **Permit Production Group Rule**, an **Incoming claim type** of **Group**, an **Incoming claim value** of **Production**, and select the option to **Permit access to users with this incoming claim**.
4. Create a new Issuance Authorization Rule that permits or denies user access based on the incoming claim. Configure the rule with the name **Allow A Datum Users**, an **Incoming claim type** of **UPN**, an **Incoming claim value** of **@adatum.com**, and select the option to **Permit access to users with this incoming claim**, and then click **Finish**.
5. Open the **Allow A Datum Users rule** properties, and show the claims rule language to the students.

## Lab: Implementing AD FS

### Scenario

A. Datum has set up a variety of business relationships with other companies and customers. Some of these partner companies and customers must access business applications that are running on the A. Datum network. The business groups at A. Datum want to provide a maximum level of functionality and access to these companies. The security and operations departments want to ensure that the partners and customers can only access the resources to which they require access, and that implementing the solution does not significantly increase the workload for the operations team.

A. Datum is also working on migrating some parts of their network infrastructure to online services, including Windows Azure and Office 365.

To meet these business requirements, A. Datum plans to implement AD FS. In the initial deployment, the company plans to use AD FS to implement single sign on for internal users accessing an application on a web server. A. Datum also has entered into a partnership with another company, Trey Research. Trey Research users must be able to access the same application.

As one of the senior network administrators at A. Datum, it is your responsibility to implement the AD FS solution. As a proof of concept, you plan to deploy a sample claims aware application, and then configure AD FS to enable both internal users and Trey Research users to access the same application.

### Objectives

- Configure the AD FS prerequisites.
- Install and configure AD FS.
- Configure and validate SSO for single organization.
- Configure and validate SSO for a business federation scenario.

### Lab Setup

Estimated time: **90 minutes**

Virtual Machines	20417A-LON-DC1 20417A-LON-SVR1 20417A-LON-CL1 20417A-MUN-DC1
User Name	Adatum\Administrator
Password	Pa\$\$w0rd

For this lab, you will use the available virtual machine environment. Before you begin the lab, you must complete the following steps:

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V® Manager, click **20417A-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.

4. Log on using the following credentials:
  - a. User name: **Adatum\Administrator**
  - b. Password: **Pa\$\$w0rd**
5. Repeat steps 2 to 4 for **20417A-LON-SVR1**, **20417A-LON-CL1**, and **20417A-MUN-DC1**.
  - a. Do not log on to 20417A-LON-CL1 at this point.
  - b. On 20417A-MUN-DC1, log in as **TreyResearch\Administrator** with the password **Pa\$\$w0rd**.

## Exercise 1: Configuring AD FS Prerequisites

### Scenario

To deploy AD FS at A. Datum, you must verify that all required components are configured. You plan to verify that AD CS is deployed in the organization, and then configure the certificates required for AD FS on the AD FS server and on the web servers. You also plan to configure the DNS forwarders to enable communication between Adatum.com and TreyResearch.com.

The main tasks for this exercise are as follows:

1. Configure DNS forwarders.
2. Exchange root certificates to enable certificate trusts.
3. Request and install a certificate for the web server.
4. Bind the certificate to the claims aware application on the web server and verify application access.

#### ► Task 1: Configure DNS forwarders

1. On LON-DC1, create a new conditional forwarder for the TreyResearch.com domain, by using the DNS server IP address of 172.16.10.10.
2. On MUN-DC1, create a new conditional forwarder for the Adatum.com domain, by using the DNS server IP address of 172.16.0.10.

#### ► Task 2: Exchange root certificates to enable certificate trusts

1. On LON-DC1, copy the **MUN-DC1.TreyResearch.com\_TreyResearch-MUN-DC1-CA.crt** from **\\MUN-DC1.treyresearch.com\certenroll** to the **Documents** folder.
2. Create a new Microsoft Management Console (MMC), and then add the **Group Policy Management Editor**.
3. Edit the Default Domain Policy Group Policy Object, and import the copied root certificate to the **Trusted Root Certification Authorities** folder.
4. On MUN-DC1, copy the **LON-DC1.Adatum.com\_Adatum-LON-DC1-CA.crt** from **\\LON-DC1.Adatum.com\certenroll** to the **Documents** folder.
5. Create a new MMC, and then add the **Certificates** snap-in focused on the Local Computer.
6. Import the copied root certificate to the **Trusted Root Certification Authorities** folder.

#### ► Task 3: Request and install a certificate for the web server

1. On LON-SVR1, open the **Internet Information Services (IIS) Manager**.
2. Request a new Domain Certificate for the server by using the following parameters:
  - Common name: **LON-SVR1.adatum.com**
  - Organization: **A. Datum**

- Organization unit: **IT**
  - City/locality: **London**
  - State/province: **England**
  - Country/region: **GB**
3. Request the certificate from the default CA.

► **Task 4: Bind the certificate to the claims aware application on the web server and verify application access**

1. On LON-SVR1, in Internet Information Services, create a new HTTPS site binding, and then select the newly created certificate.
2. On LON-DC1, open Internet Explorer, and then connect to **https://lon-svr1.adatum.com/adatumtestapp**.
3. Verify that you can connect to the site, but that you receive a 401 access denied error. This is expected because you have not yet configured AD FS for authentication.
4. Close Internet Explorer.

**Results:** In this exercise, you configured DNS forwarding to enable name resolution between A. Datum and Trey Research, and you exchanged root certificates between the two organizations. You also installed and configured a web certificate on the application server.

## Exercise 2: Installing and Configuring AD FS

### Scenario

To start the AD FS implementation, you plan to install AD FS on the A. Datum domain controller, and then configure the server as a standalone federation server. You also plan to configure the server to use a CA-signed token-signing certificate.

The main tasks for this exercise are as follows:

1. Install and configure AD FS 2.0.
2. Create a stand-alone Federation Server by using the AD FS Federation Server Configuration Wizard.
3. Verify that FederationMetaData.xml is present and contains valid data.

► **Task 1: Install and configure AD FS 2.0**

- On LON-DC1, in Server Manager, add the Active Directory Federation Services server role.

► **Task 2: Create a stand-alone Federation Server by using the AD FS Federation Server Configuration Wizard**

- On LON-DC1, run the AD FS Federation Server Configuration Wizard using the following parameters:
  - a. Create a new federation service.
  - b. Create a standalone deployment.
  - c. Use the LON-DC1.Adatum certificate.
  - d. Choose a service name of LON-DC1.Adatum.com

► **Task 3: Verify that FederationMetaData.xml is present and contains valid data**

1. On LON-CL1, log on as **Adatum\Brad**, using the password **Pa\$\$w0rd**.
2. Open Internet Explorer.
3. Open Internet Options, and then add **https://LON-DC1.Adatum.com** and **https://LON-SVR1.adatum.com** to the Local intranet zone.
4. Connect to **https://lon-dc1.adatum.com/federationmetadata/2007-06/federationmetadata.xml**.
5. Verify that the xml file opens successfully, and then scroll through its contents.
6. Close Internet Explorer.

**Results:** In this exercise, you installed and configured the AD FS server role, and then verified a successful installation by viewing the Federation Meta Data .xml contents.

### Exercise 3: Configure AD FS for a Single Organization

#### Scenario

The first scenario for implementing the proof-of-concept AD FS application is to ensure that internal users can use SSO to access the web application. You plan to configure the AD FS server and the web application to enable this scenario. You also want to verify that internal users can access the application.

The main tasks for this exercise are as follows:

1. Configure a Token Signing Certificate for LON-DC1.Adatum.com.
2. Configure the Active Directory Claims Provider Trust.
3. Configure the claims application to trust incoming claims by running the WIF Federation Utility.
4. Configure a relying party trust for the claims aware application.
5. Configure claim rules for the relying party trust.
6. Test the access to the claims aware application.

► **Task 1: Configure a Token Signing Certificate for LON-DC1.Adatum.com**

1. On LON-DC1, use the **set-ADFSProperties -AutoCertificateRollover \$False** command to enable modification of the assigned certificates.
2. In the AD FS Management console, add the **LON-DC1.Adatum.com** certificate as a new token signing certificate.

Verify that the certificate has a subject of **CN=LON-DC1.Adatum.com**. If no name is listed under the Subject when you add the certificate, delete the certificate, and then add the next certificate in the list.

3. Make the new certificate the primary certificate, and then remove the old certificate.

► **Task 2: Configure the Active Directory Claims Provider Trust**

1. In the AD FS 2.0 Management console, go to the claims provider **Trusts**, highlight the **Active Directory** store, and then go to **Edit Claim Rules**.
2. In the **Edit Claim Rules for Active Directory** dialog box on the **Acceptance Transform Rules** tab, launch the **Add Transform Claim Rule Wizard**, and then complete the wizard with the following settings:
  - a. Select **Send LDAP Attributes as Claims** under **Claim rule template**.
  - b. Name the claim rule **Outbound LDAP Attribute Rule**.
  - c. Choose Active Directory as the Attribute Store.
  - d. In the **Mapping of LDAP attributes to outgoing claim types**, select the following values:
    - E-Mail-Addresses to E-Mail Address
    - User-Principal-Name to UPN
    - Display-Name to Name

► **Task 3: Configure the claims application to trust incoming claims by running the WIF Federation Utility**

1. On LON-SVR1, launch the WIF Federation Utility from the Start screen.
2. Complete the wizard with the following settings:
  - Point to the web.config file of the WIF sample application by pointing to **C:\inetpub\wwwroot\AdatumTestApp\web.config**.
  - Specify an **Application URI** box by typing **https://lon-svr1.adatum.com/AdatumTestApp/**.
  - Select to **Use an existing STS**, and then enter a path **https://lon-dc1.adatum.com/federationmetadata/2007-06/federationmetadata.xml**.
  - Select **No encryption**.

► **Task 4: Configure a relying party trust for the claims aware application**

1. In the AD FS 2.0 Management console, click **Required: Add a trusted relying party**, in the middle pane.
2. Complete the Add relying party Wizard with the following settings:
  - Choose to **Import data about the relying party published online or on a local network** and type **https://lon-svr1.adatum.com/adatumtestapp**.
  - Specify a **Display** name of **ADatum Test App**.
  - Choose to **Permit all users to access this relying party**.
  - Choose to Permit access for all users.
  - Select the option to open the **Edit Claims Rules for WIF Sample Claims App** when the wizard is complete.



► **Task 5: Configure claim rules for the relying party trust**

1. In the **Edit Claim Rules for WIF Sample Claims App** properties dialog box, choose to **Add a Rule** on the **Issuance Transform Rules** tab.
2. Complete the Add Transform Claim Rule Wizard with the following settings:
  - Choose **Pass through of Filter an Incoming Claim** in the **Claim rule template** drop-down list.
  - Name the claim rule **Pass Through Windows Account Name**.
  - Select **Windows account name** in the **incoming claim type** drop-down list.
  - Create three more rules to pass through **E-Mail Address**, **UPN**, and **Name type claim**.

► **Task 6: Test the access to the claims aware application**

1. On LON-CL1, open Internet Explorer, and then connect to **https://lon-svr1.adatum.com/AdatumTestApp/**
2. Verify that you can access the application.

**Results:** After this exercise, you configured a token signing certificate and configured a claims provider trust for Adatum.com. You also configured the sample application to trust incoming claims and configured a relying party trust and associated claim rules. You also tested access to the sample WIF application in a single organization scenario.

## Exercise 4: Configure AD FS for Federated Business Partners

### Scenario

The second deployment scenario is to enable Trey Research users to access the web application. You plan to configure the integration of AD FS at Trey Research with AD FS at A. Datum, and then verify that Trey Research users can access the application. You also want to confirm that you can configure access based on user groups. You must ensure that all users at A. Datum, but only users in the Production group at Trey Research, can access the application.

The main tasks for this exercise are as follows:

1. Add a claims provider trust for the TreyResearch.com AD FS server.
2. Configure a relying party trust on MUN-DC1 for A. Datum's claim aware application.
3. Verify access to the A. Datum Test Application for Trey Research users.
4. Configure claim rules for the claim provider trust and the relying party trust to allow access only for a certain group.
5. Verify restrictions and accessibility to the claims aware application.
6. To shut down the virtual machines.

► **Task 1: Add a claims provider trust for the TreyResearch.com AD FS server**

1. On LON-DC1, in the ASDFS 2.0 Management console, go to **Trust Relationships**, go to claims provider **Trusts**, and then choose to **Add claims provider Trust**.
2. Complete the Add claims provider Trust Wizard with the following settings:
  - Choose **Import data about the claims provider published online or on a local network** and enter **https://mun-dc1.treyresearch.com** as the data source.
  - In **Display Name** enter **mun-dc1.treyresearch.com**.
  - Complete the wizard.

3. In the **Edit Claim Rules for the mun-dc1.treyresearch.com** properties dialog, use the following values:
  - **Add a Rule** to the Acceptance Transform Rules.
  - Choose **Pass Through or Filter an Incoming claim** in the **Claim rule template** list.
  - Use **Pass through Windows account name** rule as the claim rule name.
  - Choose **Windows account name** as the incoming claim type, and then choose to **Pass through all claim values**.
  - Complete the rule.
4. On LON-DC1, run the following command in Windows PowerShell:

```
Set-ADFSClaimsProviderTrust -TargetName "nyc-dc1.contoso.com" -
SigningCertificateRevocationCheck None
```

### ► Task 2: Configure a relying party trust on MUN-DC1 for A. Datum's claim aware application

1. On MUN-DC1, in the AD FS Management console, open the Add relying party Trust Wizard, and then complete it with the following settings:
  - Choose to **Import data about the relying party published online or on a local network** and type in **https://lon-dc1.adatum.com**.
  - Specify a **Display** name of **Adatum TestApp**.
  - Choose to **Permit all users to access this relying party**.
  - Select to open the **Edit Claim Rules for lon-dc1.adatum.com** when the wizard is complete check box.
2. In the **Edit Claim Rules for lon-dc1.adatum.com** properties dialog box, on the **Issuance Transform Rules** tab, click to add a rule with the following settings:
  - Choose **Pass Through or Filter an Incoming claim** in claim rule template list.
  - In the **Claim rule name** box, type **Pass through Windows account name rule**.
  - Choose **Windows account name** in **Incoming claim type**.
  - Choose to **Pass through all claim values**.
  - Complete the wizard.

### ► Task 3: Verify access to the A. Datum Test Application for Trey Research users

1. On MUN-DC1, open Internet Explorer, and then connect to **https://lon-svr1.adatum.com/adatumtestapp/**.
2. Select **mun-dc1.treyresearch.com** as the home realm, and then logon as **TreyResearch\April**, with the password **Pa\$\$w0rd**.
3. Verify that you can access the application.
4. Close Internet Explorer, and then connect to the same web site. Verify that you are not prompted for a home realm this time.

You are not prompted for a home realm again. Once users have selected a home realm and been authenticated by a realm authority, they are issued with an `_LSRealm` cookie by the relying party Federation Server. The default lifetime for the cookie is 30 days. Therefore, for us to log on multiple times, we should delete that cookie after each logon attempt to return to a clean state.

#### ► Task 4: Configure claim rules for the claim provider trust and the relying party trust to allow access only for a certain group

1. On MUN-DC1, in the AD FS Management Console, access the lon-dc1.adatum.com relying party trust.
2. Add a new Issuance Transform Rule that sends the group membership as a claim. Name the rule **Permit Production Group Rule**, configure the **User's Group** as **Production**, configure the **Outgoing claim type** as **Group**, and the **Outgoing claim value** as **Production**.
3. On LON-DC1, in the AD FS Management Console, edit the mun-dc1.treyresearch.com claims provider Rule, creating a new rule that passes through or filters an incoming claim with the rule name of **Send Production Group Rule**. Configure the rule with an incoming claim type of **Group**.
4. Edit the Adatum Test App relying party trust by creating a new Issuance Transform Rule that passes through or filters an incoming claim. Name the rule **Send TreyResearch Group Name** rule, and configure the rule to use an incoming claim type of group.
5. Delete the Issuance Authorization Rule that grants access to all users.
6. Create a new Issuance Authorization Rule that permits or denies user access based on the incoming claim. Configure the rule with the name **Permit TreyResearch Production Group Rule**, an **Incoming claim type** of **Group**, an **Incoming claim value** of **Production**, and select the option to **Permit access to users with this incoming claim**.
7. Create a new Issuance Authorization Rule that permits or denies user access based on the incoming claim. Configure the rule with the name **Temp**, an **Incoming claim type** of **UPN**, an **Incoming claim value** of **@adatum.com**, and select the option to **Permit access to users with this incoming claim**, and then click **Finish**.
8. Edit the Temp rule, and then copy the claim rule language into the clipboard.
9. Delete the Temp rule.
10. Create a new rule that sends claims using a custom rule named **ADatum User Access Rule**
11. Click in the **Custom rule** box, and then press Ctrl+V to paste the clipboard contents into the box. Edit the first URL to match the following text, and then click **Finish**:

```
c:[Type == "http://schemas.xmlsoap.org/ws/2005/05/identity/claims/upn", Value =~
"^(?i).+@adatum\.com$"]=> issue(Type =
"http://schemas.microsoft.com/authorization/claims/permit", Value =
"PermitUsersWithClaim");
```

#### ► Task 5: Verify restrictions and accessibility to the claims aware application

1. On MUN-DC1, verify that TreyResearch\April no longer has access to the A. Datum test app.
2. Clear the browsing history in Internet Explorer.
3. Verify that TreyResearch\morgan does have access to the A. Datum test app. Morgan is a member of the Production group.

#### ► To shut down the virtual machines

- When you are finished the lab, revert the virtual machines to their initial state.

**Results:** In this exercise, you configured a claims provider trust for Trey Research on Adatum.com and a relying party trust for Adatum on TreyResearch.com. You verified access to the A. Datum claim-aware application. Then you configured the application to restrict access from TreyResearch.com to specific groups, and you verified appropriate access.

## Module Review and Takeaways

### Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Certificate errors on the federation server	
Certificate errors on the client	
Client application failed to authenticate with AD FS	

**Question:** What are the benefits of deploying AD FS with a cloud-based application or service?

**Question:** Under what circumstances, would you choose to deploy a federation proxy server?  
Under what circumstances, do you *not* have to deploy a federation proxy server?

### Real-world Issues and Scenarios

1. Tailspin Toys is deploying a new claims-based web application. The web application needs to be accessible to both Tailspin Toys users and to Trey Research users. What AD FS components will you need to deploy at Tailspin Toys to enable this level of access?
2. Fabrikam is examining the requirements for AD FS. The company wants to use a federation proxy server for maximum security. Currently, Fabrikam has an internal network with internal DNS servers. Their internet-facing DNS is hosted by a hosting company. The perimeter network uses the hosting company's DNS servers for DNS resolution. What must the company do to prepare for the deployment?

## Course Evaluation

Your evaluation of this course will help Microsoft understand the quality of your learning experience.

Please work with your training provider to access the course evaluation form.

Microsoft will keep your answers to this survey private and confidential and will use your responses to improve your future learning experience. Your open and honest feedback is valuable and appreciated.

**MCT USE ONLY. STUDENT USE PROHIBITED**

## Module 1: Installing and Configuring Servers Based on Windows Server 2012

# Lab: Installing and Configuring Servers Based on Windows Server 2012

### Exercise 1: Install Windows Server 2012 Server Core

#### ► Task 1: Install Windows Server 2012

1. On the host machine, open the Hyper-V Manager console.
2. Click **20417A-LON-SVR5**. In the Actions pane click **Settings**.
3. Under **Hardware**, click **DVD Drive**.
4. Click **Image file**, and then click **Browse**.
5. Browse to **C:\Program Files\Microsoft Learning\20417\Drives**, and then click **Win2012\_RC.ISO**.
6. Click **Open** and then click **OK**.
7. In the Hyper-V Manager console, double-click **20417A-LON-SVR5**; this will open the Virtual Machine Connection window. From the **Action** menu, click **Start**.
8. On the **Windows Server 2012** page of the Windows Setup Wizard, verify the following settings, and then click **Next**:
  - Language to install: **English (United States)**
  - Time and currency format: **English (United States)**
  - Keyboard or input method: **US**
9. On the **Windows Server 2012** page of the Windows Setup Wizard, click **Install now**.
10. On the **Select the operating system you want to install** page of the Windows Setup Wizard, select **Windows Server 2012 Release Candidate Datacenter (Server Core Installation)**, and then click **Next**.
11. On the **License terms** page of the Windows Setup Wizard, review the operating system license terms. Select the **I accept the license terms** check box, and then click **Next**.
12. On the **Which type of installation do you want?** page of the Windows Setup Wizard, click **Custom: Install Windows Only (Advanced)**.
13. On the **Where do you want to install Windows?** page of the Windows Setup Wizard, verify that **Drive 0 Unallocated Space** has sufficient space for the Windows Server 2012 operating system, and then click **Next**:
  - Depending on the speed of the host computer, the installation will take approximately 20 minutes.
  - The virtual machine will restart several times during this process.
14. Click **OK**, and then in both the **Password** and **Confirm password** boxes type **Pa\$\$w0rd**, and then click **OK**.

**► Task 2: Convert a Windows Server 2012 Server Core installation to a full installation**

1. If necessary, log on to **LON-SVR5** using the **Administrator** account with the password **Pa\$\$w0rd**.
2. At the command prompt type and press Enter:

```
mkdir c:\mount
```

3. Issue the following command and press Enter to mount the Windows Server 2012 full installation image:

```
dism.exe /mount-image /ImageFile:d:\sources\install.wim /Index:4 /Mountdir:c:\mount /readonly
```

4. Start Windows PowerShell by issuing the command:

```
PowerShell.exe
```

5. Load the ServerManager module by issuing the command and pressing Enter:

```
Import-Module ServerManager
```

6. Install the Windows Server 2012 GUI components of server core by issuing the following command and pressing Enter:

```
Install-WindowsFeature -IncludeAllSubfeature User-Interfaces-Infra -  
Source:c:\mount\windows
```

7. When prompted, restart the server by issuing the following command and pressing Enter.

```
Shutdown /r /t 5
```

8. Log on to **LON-SVR5** as **Administrator** with the password of **Pa\$\$w0rd** and verify the presence of the full GUI components.

**► Task 3: Convert a Windows Server 2012 full installation to a Server Core installation**

1. If necessary, log on to **LON-SVR5** and verify that the full graphic environment is present.
2. Click **Internet Explorer**.
3. Click **Close** to close the message informing you that you cannot open Internet Explorer with the built-in Administrator account.
4. On the **Start** screen, click **Windows PowerShell**.
5. Enter the following command and press Enter:

```
Import-Module ServerManager
```

6. Enter the following command and press Enter:

```
Uninstall-WindowsFeature User-Interfaces-Infra
```

7. Enter the following command to restart LON-SVR5:

```
Shutdown /r /t 5
```

8. Log on to **LON-SVR5** as **Administrator** with the password of **Pa\$\$w0rd** and verify that it now configured to use the Server Core configuration.



## Exercise 2: Configure a Computer Running a Server Core Installation of Windows Server 2012

### ► Task 1: Configure the network

1. If necessary, log on to **LON-SVR5** using the account **Administrator** with password **Pa\$\$w0rd**.
2. At the command prompt, type **sconfig**.
3. Type **2** and press Enter to select **Computer Name**:
4. Enter the computer name **LON-SVR5** and press Enter.
5. On the **Restart** dialog box, click **Yes**.
6. Log on to **LON-SVR5** as **Administrator** with the password of **Pa\$\$w0rd**.
7. At the command prompt, type **hostname** and press Enter to verify the computer's name.
8. At the command prompt, type **sconfig** and press Enter.
9. To configure Network Settings, type **8** and press Enter.
10. Type the index number of the network adapter that you want to configure and press Enter.
11. To set the Network Adapter Address, on the **Network Adapter Settings** page, type **1** and press Enter.
12. To select static IP address configuration, type **S** and press Enter.
13. At the **Enter static IP address:** prompt, type **172.16.0.111** and press Enter.
14. At the **Enter subnet mask** prompt, type **255.255.0.0** and press Enter.
15. At the **Enter default gateway** prompt, type **172.16.0.1** and press Enter.
16. To configure the DNS server address, on the Network Adapter Settings page, type **2** and press Enter.
17. At the **Enter new preferred DNS server** prompt, type **172.16.0.10** and press Enter.
18. In the **Network Settings** dialog box, click **OK**.
19. To not configure an alternative DNS server address, press Enter.
20. To return to the main menu, type **4** and press Enter.
21. To exit sconfig, type **15** and press Enter.
22. To verify connectivity to the domain controller from LON-SVR5, type **ping lon-dc1.adatum.com** and press Enter.

### ► Task 2: Add the server to the domain

1. Ensure that you are logged on to LON-SVR5 using the account **Administrator** with password **Pa\$\$w0rd**.
2. At the command prompt, type **sconfig** and press Enter.
3. To switch to configure Domain/Workgroup, type **1** and press Enter.
4. To join a domain, type **D** and press Enter.
5. At the **Name of domain to join** prompt, type **adatum.com** and press Enter.
6. At the **Specify an authorized domain\user** prompt, type **adatum\administrator** and press Enter.
7. At the **Type the password associated with the domain user** prompt, type **Pa\$\$w0rd** and press Enter.

8. At the **Change Computer Name** prompt, click **Yes**.
9. At the **Enter new computer name** prompt, press Enter.
10. To restart the server, type **13** and press Enter.
11. In the **Restart** dialog box, click **Yes**.
12. Log on to **LON-SVR5** with the **adatum\administrator** account and a password of **Pa\$\$w0rd**.

### ► Task 3: Configure Windows Firewall

1. Ensure that you are logged on to LON-SVR5 using the account **Adatum\Administrator** with password **Pa\$\$w0rd**.
2. At the command prompt, type **sconfig.cmd** and press Enter.
3. To switch to **Configure Remote Management**, type **4** and press Enter.
4. To enable Remote Management, type **1** and press Enter.
5. On the **Configure Remote Management** dialog box, click **OK**.
6. To return to the main menu, type **4** and press Enter.
7. To return to the command prompt, type **15** and press Enter.
8. At the command prompt, type **PowerShell.exe** and then press Enter.
9. To view the enabled Firewall rules on LON-SVR5 that allow traffic, at the Windows PowerShell prompt, type the following command:

```
Get-NetFirewallRule | Where-Object {$_.Action -eq "Allow"} | Format-Table -Property DisplayName
```

10. To view all disabled Firewall rules on LON-SVR5, type the following command:

```
Get-NetFirewallRule | Where-Object {$_.Enabled -eq "False"} | Format-Table -Property Displayname
```

11. To view all NetFirewallRule related Windows PowerShell cmdlets, type the following command:

```
Get-Command -Noun NetFirewallRule
```

12. To view the status of the Remote Desktop inbound firewall rule, type the following command:

```
Get-NetFirewallRule RemoteDesktop-UserMode-In-TCP
```

13. To enable the Remote Desktop Inbound Firewall rule, type the following command:

```
Enable-NetFirewallRule RemoteDesktop-UserMode-In-TCP
```

14. To verify that the Remote Desktop Inbound Firewall rule is enabled, type the following command:

```
Get-NetFirewallRule RemoteDesktop-UserMode-In-TCP
```

15. To disable the Remote Desktop Inbound Firewall Rule, type the following command:

```
Disable-NetFirewallRule RemoteDesktop-UserMode-In-TCP
```

16. To verify that the Remote Desktop Inbound Firewall Rule is disabled, type the following command:

```
Get-NetFirewallRule RemoteDesktop-UserMode-In-TCP
```

## Exercise 3: Configure Remote Management for servers running Windows Server 2012

### ► Task 1: Validate the WinRM configuration

1. Log on to **LON-DC1** using the **Adatum\Administrator** account with the password **Pa\$\$w0rd**.
2. In the Server Manager console, click **Local Server**, and then click **Enabled** next to **Remote Management**.
3. On the **Configure Remote Management** dialog box, clear the check next to **Enable remote management of this server from other computers**, and then click **OK**.
4. Close the Server Manager console.
5. Open **Windows PowerShell** from the Taskbar.
6. At the Windows PowerShell prompt issue the command **winrm qc**. When you are prompted, type **Y** and press Enter.
7. Open the Server Manager console. Click **Local Server**. Verify that **Remote Management** is now enabled.

### ► Task 2: Configure Server Manager for multiple server management

1. Log on to **LON-DC1** using the **Adatum\Administrator** account with the password **Pa\$\$w0rd**.
2. In the Server Manager console, click **Dashboard**, and then click **Create a server group**.
3. On the **Create Server Group** dialog box, click the **Active Directory** tab, and then click **Find Now**.
4. Click **LON-DC1** and then press and hold the Ctrl key, and then click **LON-SVR5**. To add them to a server group click the Arrow.
5. Set the Server Group Name to **LONDON-GROUP**, and then click **OK**.
6. In Server Manager click **LONDON-GROUP**.
7. In the details pane, select both **LON-DC1** and **LON-SVR5**.
8. Scroll down to the **Performance** section.
9. Click **LON-DC1**. Press and hold the Ctrl key, and then click **LON-SVR5**.
10. While both servers are selected, right-click **LON-DC1**, and then click **Start Performance Counters**.
11. Scroll up and verify that in the **Manageability** column, both **LON-DC1** and **LON-SVR5** are listed as **Online**.

### ► Task 3: Deploy a feature to the Server Core server

1. On LON-DC1, in the Server Manager console, click **LONDON-GROUP**.
2. In the **Servers** list, right-click **LON-SVR5**, and then click **Add Roles and Features**.
3. On the **Before You Begin** page of the Add Roles and Features Wizard, click **Next**.
4. On the **Select installation type** page of the Add Roles and Features Wizard, select **Role-based or feature-based installation**, and then click **Next**.
5. On the **Select destination server** page of the Add Roles and Features Wizard, ensure that **LON-SVR5.Adatum.com** is selected, and then click **Next**.
6. On the **Select server roles** page of the Add Roles and Features Wizard, click **Next**.
7. On the **Select features** page of the Add Roles and Features Wizard, select **Windows Server Backup**, and then click **Next**.

8. On the **Confirm installation selections** page of the Add Roles and Features Wizard, click **Install**.
9. To dismiss the Add Roles and Features Wizard, click **Close**.
10. In Server Manager, click the Flag and verify that the installation of the Windows Server Backup feature succeeded on LON-SVR5.

► **Task 4: To prepare for next module**

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20417A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20417A-LON-SVR5**.

## Module 2: Monitoring and Maintaining Windows Server 2012

# Lab: Monitoring and Maintaining Windows 2012 Servers

### Exercise 1: Configuring Centralized Monitoring for Windows Server 2012 Servers

#### ► Task 1: Configure Server Manager to monitor multiple servers

1. Switch to **LON-SVR1**.
2. In the Server Manager console, in the navigation pane, click **All Servers**.
3. In the Server Manager console, in the navigation pane, right-click **All Servers**, and then click **Add Servers**.
4. In the **Add Servers** dialog box, click **Find Now**.
5. In the details pane of the **Add Servers** dialog box, click **LON-DC1**, click the right-arrow button, and then click **OK**.
6. In Server Manager, hold down the Ctrl key, click **LON-DC1**, and then click **LON-SVR1** to select both the machines.
7. In Server Manager, scroll down to the **Performance** section; select both **LON-DC1** and **LON-SVR1**. Right-click the selected servers, and then click **Start Performance Counters**.

#### ► Task 2: Configure a data collector set

1. On LON-SVR1, in Server Manager, click **Tools**, and then click **Performance Monitor**.
2. In the navigation pane, expand **Data Collector Sets**, and then click **User Defined**.
3. Click the **Action** menu, click **New**, and then click **Data Collector Set**.
4. In the Create new Data Collector Set Wizard, in the **Name** box, type **Windows Server Monitoring**, select **Create manually (Advanced)**, and then click **Next**.
5. On the **What type of data do you want to include?** page, ensure that the **Create data logs** option button is selected, select the **Performance Counter** check box, and then click **Finish**.
6. In the Performance Monitor, in the navigation pane, expand **Data Collector Sets**, expand **User Defined**, click **Windows Server Monitoring**, click the **Action** menu, click **New**, and then click **Data Collector**.
7. In the Create New Data Collector Wizard, in the **Name** box, type **Base Windows Server Monitoring**, select **Performance counter data collector**, click **Next**, and then click **Add**.
8. In the **Available counters** object list, expand **Processor**, and then click **% Processor Time**. Click **Add**.
9. In the **Available counters** object list, expand **Memory**, and then click **Available Mbytes**. Click **Add**.
10. In the **Available counters** object list, expand **Logical Disk**, click **% Free Space**, click **Add**, and then click **OK**.
11. In the Create New Data Collector Wizard, in the **Sample interval** box, accept the default values, and then click **Finish**.

12. In the Performance Monitor, in the navigation pane, click **Windows Server Monitoring**, click the **Action** menu, and then click **Start**.
13. Wait at least one minute, click the **Action** menu, and then click **Stop**.
14. In the navigation pane, expand **Reports**, expand **User Defined**, expand **Windows Server Monitoring**, click **LON-SVR1\_DateTime**, and then review the report.
15. Close the Performance Monitor.

► **Task 3: Configure an event subscription**

1. Switch to **LON-SVR1**.
2. Move the mouse pointer on the lower-right corner on the screen, and then in **Search** box, type **cmd** to open the **Command Prompt**.
3. At the command prompt, type **winrm quickconfig** and then press Enter.
4. In Server Manager, click **Tools**, and then click **Computer Management**.
5. In the Computer Management console, expand **Local Users and Groups**, and then click **Groups**.
6. In the details pane, double-click **Administrators**.
7. Click **Add**, and in the **Select Users, Computers, Service Accounts or Groups** dialog box, click **Object Types**.
8. In the **Object Types** dialog box, select the **Computers** check box, and then click **OK**.
9. In the **Select Users, Computers, Service Accounts or Groups** dialog box, in the **Enter the object names to select** box, type **LON-DC1**, and then click **OK**.
10. In the **Administrators Properties** dialog box, click **OK**.
11. Switch to **LON-DC1**.
12. Move the mouse pointer on the lower-right corner on the screen, and then in **Search** box, type **cmd** to open the **Command Prompt**.
13. At the command prompt, type **wecutil qc** and then press Enter.
14. When you are prompted, type **Y** and then press Enter.
15. In Server Manager, click **Tools**, and then click **Event Viewer**.
16. In the Event Viewer, in the navigation pane, click **Subscriptions**.
17. Right-click **Subscriptions**, and then click **Create Subscription**.
18. In the **Subscription Properties** dialog box, in the **Subscription** name box, type **LON-SVR1 Events**.
19. Click **Collector Initiated**, and then click **Select Computers**.
20. In the **Computers** dialog box, click **Add Domain Computers**.
21. In the **Select Computer** dialog box, in the **Enter the object name to select** box, type **LON-SVR1**, and then click **OK**.
22. In the **Computers** dialog box, click **OK**.
23. In the **Subscription Properties – LON-SVR1 Events** dialog box, click **Select Events**.
24. In the **Query Filter** dialog box, select the **Critical**, **Warning**, **Information**, **Verbose**, and **Error** check boxes.
25. In the **Logged** list, click **Last 7 days**.

26. In the **Event logs** list, select **Windows Logs**. Click inside the **Query Filter** dialog box, and then click **OK**.
27. In the **Subscription Properties – LON-SVR1 Events** dialog box, click **OK**.
28. In Event Viewer, in the navigation pane, expand **Windows Logs**.
29. Click **Forwarded Events**, and check for events from LON-SVR1.

**Results:** After completing this exercise, you will have configured Server Manager to monitor multiple servers, configured a data collector set, and configured an event subscription.

## Exercise 2: Backing up Windows Server 2012

### ► Task 1: Install the Windows Server Backup feature

1. Switch to **LON-SVR1**.
2. In Server Manager, on the Dashboard, click **Add Roles and Features**.
3. In the Add Roles and Features Wizard, click **Next**.
4. On the **Select Installation Type** page, click **Next**.
5. On the **Select Destination Server** page, select **LON-SVR1** and then click **Next**.
6. On the **Select server roles** page, click **Next**.
7. On the **Select features** page, select **Windows Server Backup**, and then click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. On the **Installation progress** page, wait until the **Installation succeeded on LON-SVR1.adatum.com** text appears, and then click **Close**.

### ► Task 2: Configure a scheduled backup

1. Switch to **LON-SVR1**.
2. On LON-SVR1, in Server Manager, click **Tools**, and then click **Windows Server Backup**.
3. Click **Local Backup**, and then in the Actions pane, click **Backup Schedule**.
4. On the **Getting Started** page of the Backup Schedule Wizard, click **Next**.
5. On the **Select Backup Configuration** page, click **Full server (recommended)**, and then click **Next**.
6. On the **Specify Backup Time** page, next to **Select time of day**, select **1:00 AM**, and then click **Next**.
7. On the **Specify Destination Type** page, click **Backup to a shared network folder**, and then click **Next**. Review the warning, and then click **OK**.
8. On the **Specify Remote Shared Folder** page, in the **Path** box, type **\\LON-DC1\Backup**, and then click **Next**.
9. In the **Register Backup Schedule** dialog box, in the **Username** box, type **Administrator**, in the **Password** box, type **Pa\$\$w0rd**, and then click **OK**. Click **Finish**, and then click **Close**.

### ► Task 3: Complete an on-demand backup

To prepare for this task, you need to create a folder on LON-SVR1 with a name **Financial Data** on drive **C:** and within **Financial Data** folder you need to create a text file with a name **Financial Report.txt**.

1. On LON-SVR1, on the Taskbar, click on **Windows Explorer**.
2. In the Windows Explorer window, in navigation pane, click on **Local Disk (C:)**.
3. In the Windows Explorer window, in the menu, click **Home**, click **New Folder**, and then in the **New Folder** icon in details pane, type **Financial Data**.
4. In the Windows Explorer window, double-click **Financial Data** folder, right click in details pane, click **New**, click **Text Document**, and in **New Text Document** icon, type **Financial Report**.

**To complete an on-demand backup, perform the following steps:**

1. On LON-SVR1, in Server Manager, click **Tools**, and then click **Windows Server Backup**.
2. In the wbadmin – [Windows Server Backup (Local)] window, in the navigation pane, click **Local Backup**, and then in the Actions pane, click **Backup Once**.
3. On the **Backup Options** page of the Backup Once Wizard, click **Different options**, and then click **Next**.
4. On the **Select Backup Configuration** page, click **Custom**, and then click **Next**.
5. On the **Select Items for Backup** page, click **Add Items**.
6. Expand **Local disk (C:)**, select the **Financial Data** check box, click **OK**, and then click **Next**.
7. On the **Specify Destination Type** page, click **Remote shared folder**, and then click **Next**.
8. On the **Specify Remote Folder** page, type **\\LON-DC1\Backup**, and then click **Next**.
9. On the **Confirmation** page, click **Backup**.
10. On the **Backup Progress** page, click **Close** after the backup is complete.

**Results:** After completing this exercise, you will have installed the Windows Server Backup feature, configured a scheduled backup, and ran an on demand backup.

## Exercise 3: Restoring files by using Windows Server Backup

### ► Task 1: Delete a file from the file server

1. On LON-SVR1, on the Taskbar, click on **Windows Explorer**, and then in navigation pane, click on **Local Disk (C:)**.
2. In Windows Explorer in details pane, right-click **Financial Data** folder, and then click **Delete**.

### ► Task 2: View the available restores by using the Vssadmin command

1. On LON-SVR1, on the Taskbar click **Windows Powershell**.
2. At the Windows Powershell prompt, run the following command:

```
vssadmin list shadows
```

The command should display the existing shadow copy from the backup performed previously.



### ► Task 3: Restore the file from backup

1. In the Windows Server Backup console, in the Actions pane, click **Recover**.
2. On the **Getting Started** page, click **A backup stored on another location**, and then click **Next**.
3. On the **Specify Location type** page, click **Remote shared folder**, and then click **Next**.
4. On the **Specify Remote Folder** page, type `\\LON-DC1\Backup`, and then click **Next**.
5. On the **Select Backup Date** page, click **Next**.
6. On the **Select Recovery Type** page, click **Next**.
7. On the **Select Items to Recover** page, expand **LON-SVR1**, click **Local Disk (C:)** drive, and on the right pane, select **Financial Data**, and then click **Next**.
8. On the **Specify Recovery Options** page, under **Another Location**, type `C:\`, and then click **Next**.
9. On the **Confirmation** page, click **Recover**.
10. On the **Recovery Progress** page, click **Close**.
11. Locate `C:\` and ensure that the **Financial Data** folder is restored to drive C.

**Results:** After completing this exercise, you will have deleted a folder to simulate data loss, viewed available resources, and then restored the folder the backup that you created.

## Exercise 4: Implementing Microsoft Online Backup and Restore

### ► Task 1: Install the Microsoft Online Backup Service component

1. On LON-SVR1, on the taskbar, click **Windows Explorer**.
2. In the Windows Explorer window, in the navigation pane, click **Allfiles (E:)**, and in the details pane double-click `msoidcli.msi`. Click **Run**.
3. On the Microsoft Software License Terms page, click **I accept the terms in the License Agreement and Privacy Statement**, and then click **Install**. Click **Finish**.
4. In **Allfiles (E:)**, in the details pane double-click `OBSInstaller.exe`. Click **Run**.
5. In the **Microsoft Online Service Pre-Release Agreement** dialog box, select **I accept the Service Agreement terms and conditions**, and then click **OK**.
6. On the **Prerequisites Check** page, click **Next**.
7. On the **Installation Settings** page, specify the settings (if not default), and then click **Next**:
  - Installation Folder: **C:\Program Files**
  - Cache Location: **C:\Program Files\Microsoft Online Backup Service Agent**
8. On the **Microsoft Update Opt-In** page, select **I don't want to use Microsoft Update**, and then click **Install**.
9. On the **Installation** page, ensure that the **Microsoft Online Backup Service Agent installation has completed successfully** message is displayed. Clear the **Check for newer updates** check box, and then click **Finish**.

10. On LON-SVR1, move the mouse pointer on the lower-left corner of the screen, click **Start**, and then click **Microsoft Online Backup Service**.
11. On LON-SVR1, move the mouse pointer on the lower-left corner of the screen, click **Start**, and then click **Microsoft Online Backup Service Shell**.

### ► Task 2: Register the server with Microsoft Online Backup

Before you start this task, you should rename **LON-SVR1** to **YOURCITYNAME-YOURNAME**, for example **NEWYORK-ALICE**. This is because this exercise will be performed online, and therefore the computer names used in this lab should be unique. If there is more than one student in the classroom with a same name, add a number at the end of the computer name, such as **NEWYORK-ALICE-1**.

To rename **LON-SVR1**, perform the following steps:

1. In the Server Manager window, on the **Welcome to Server Manager** page, click **1. Configure this local server**.
2. In the Server Manager window, on the **Local Server** page, click **LON-SVR1**.
3. In the System Properties window, click **Change**, in the **Computer Name** box, type **YOURCITYNAME-YOURNAME**, click **OK** twice, and then click **Close**.
4. In a window that displays the message that you should restart your computer, click **Restart Now**.
5. Wait until **YOURCITYNAME-YOURNAME** is restarted, and then log on as **Adatum\Administrator** with password **Pa\$\$w0rd**.

To register the server with Microsoft Online Backup, perform the following steps:

1. Start the Microsoft Online Backup Service console, and then click **Register Server**.
2. In the Register Server Wizard, on the **Account Credentials** page, in the **Username** box, type **holuser@onlinebackupservice.onmicrosoft.com**, and in the **Password** box, type **Pa\$\$w0rd**. Click **Next**.



**Note:** In real-life scenario, you would type username and password of your Microsoft Online Backup Service subscription account.

3. On the **Proxy Configuration** page, click **Next**.
4. On the **Encryption Settings** page, in the **Enter passphrase** and **Confirm passphrase** boxes, type **Pa\$\$w0rdPa\$\$w0rd**, and then click **Register**.
5. On the **Server Registration** page, ensure that the **Microsoft Online Backup Service is now available for this server** message is displayed, and then click **Close**.

### ► Task 3: Configure an online backup

1. Switch to the Microsoft Online Backup Service console, and then click **Schedule Backup**.
2. On the **Getting started** page, click **Next**.
3. On the **Select Items to back up** page, click **Add Items**.
4. In the **Select Items** dialog box, expand **C:**, select **Financial Data**, click **OK**, and then click **Next**.
5. On the **Specify Backup Time** page, select **Saturday**, click **1:00AM**, click **Add**, and then click **Next**.
6. On the **Specify Retention Setting** page, accept the default settings, and then click **Next**.
7. On the **Confirmation** page, click **Finish**.

8. On the **Modify Backup Progress** page, click **Close**.
9. In the Microsoft Online Backup Service console, click **Back Up Now**.
10. In the Back Up Now Wizard, on the **Confirmation** page, click **Back Up**.
11. On the **Backup progress** page, wait until **Backup is successfully completed** message appears, and then click **Close**.

► **Task 4: Restore files by using the online backup**

1. On the taskbar, click **Windows Explorer**, and then in the navigation pane, click **Local Disk (C:)**.
2. In the Local Disk (C:) window, right-click the **Financial Data** folder, and then click **Delete**.
3. Switch to the Microsoft Online Backup Service console, and then click **Recover Data**.
4. In the Recover Data Wizard, on the **Getting Started** page, select **This server**, and then click **Next**.
5. On the **Select Recovery Mode** page, select **Browse for files**, and then click **Next**.
6. On the **Select Volume and Date** page, in the **Select the volume** drop-down list, select **C:\**. In the calendar, click the date when you performed the backup, in the **Time** drop-down list, click the time when you performed backup, and then click **Next**.
7. On the **Select Items to Recover** page, expand **C:\**, click the **Financial Data** folder, and then click **Next**.
8. On the **Specify Recovery Options** page, select **Original location** and **Create copies so that you have both versions**, and then click **Next**.
9. On the **Confirmation** page, click **Recover**.
10. On the **Recovery Progress** page, ensure that **File(s) recovery job succeeded** status message appears, and then click **Close**.
11. Locate **C:\** and ensure that the **Financial Data** folder is restored to drive C.

► **Task 5: Unregister the server from the Microsoft Online Backup Service**

1. Switch to the Microsoft Online Backup Service console, and then click **Unregister Server**.
2. On the **Getting started** page, click **Unregister this server**, and then click **Next**.
3. On the **Account Credentials page**, provide the following credentials:
  - Username: **holuser@onlinebackupservice.onmicrosoft.com**,
  - Password: **Pa\$\$w0rd**
4. Click **Unregister**.
5. On the **Server Unregistration** page, click **Close**.

**Results:** After completing this exercise, you will have installed the Microsoft Online Backup Service agent, registered the server with Microsoft Online Backup Service, configured a scheduled backup, and performed a restore by using Microsoft Online Backup Service.

► **Task: To prepare for the next module**

When you finish the lab, revert the virtual machines to their initial state. To do this, complete the following steps.

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20417A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20417A-LON-SVR1** and **MSL-TMG1**.

## Module 3: Managing Windows Server 2012 by Using Windows PowerShell 3.0

# Lab: Managing Servers Running Windows Server 2012 by Using Windows PowerShell 3.0

### Exercise 1: Introduction to Windows PowerShell 3.0

#### ► Task 1: Use Windows PowerShell ISE to retrieve basic information about LON-DC1

1. Start the following virtual machines: **LON-DC1**, **LON-SVR1**, and **LON-SVR2**.
2. On LON-DC1, browse to the Start screen, type **Windows PowerShell ISE** and then right-click **Windows PowerShell ISE**. In the pop-up banner, click **Run as administrator**.
3. In the Console pane, type **Get-WindowsFeature** and then press Enter.
4. In the Console pane, type **Get-ChildItem E:\ModXA\Democode**, and then press Enter.
5. In the Console pane, type **dir C:\Windows**, and then press Enter.
6. In the Console pane, type **Get-E**, press the Tab key until **Get-ExecutionPolicy** is shown, and then press the Enter key.

#### ► Task 2: Use Windows PowerShell ISE to retrieve a list of stopped services on LON-DC1

1. If necessary, open Windows PowerShell ISE as an administrator.
2. In the Console pane, type **Get-Service** and then press Enter.
3. In the Console pane, type **\$Services = Get-Service** and then press Enter.
4. In the Console pane, type **Get-Help Where-Object -examples** and then press Enter. Click **No to update help**.
5. In the Console pane, type **\$Services | Where-Object {\$\_.Status -eq "Stopped"}** and then press Enter.

#### ► Task 3: Use a Remote Windows PowerShell session to install XPS Viewer on LON-SVR1

1. In Windows PowerShell ISE, click **File**, and then click **New Remote PowerShell Tab**.
2. In the New Remote PowerShell Tab window, in the **Computer** box, type **LON-SVR1** and then click **Connect**.
3. In the Console pane, type **Get-WindowsFeature** and then press Enter.
4. In the Console pane, type **Add-WindowsFeature XPS-Viewer** and then press Enter.
5. Press the Up Arrow key two times or until **Get-WindowsFeature** appears. Press Enter to execute.
6. On the **LON-SVR1 Remote PowerShell** tab, click **Close**.

**Results:** After this exercise, you will have explored the Windows PowerShell ISE interface and used cmdlets, variables, and pipelining.

## Exercise 2: Managing AD DS by Using Windows PowerShell 3.0

### ► Task 1: Import the Active Directory PowerShell module and view the available cmdlets

1. If it is necessary, open Windows PowerShell ISE as an administrator.
2. In the Console pane, type **Import-Module ActiveDirectory** and then press Enter.
3. In the Console pane, type **Get-Command -Module ActiveDirectory** and then press Enter.

### ► Task 2: View options on how to create a report of users in the Active Directory domain

1. If it is necessary, open Windows PowerShell ISE as an administrator and import the Active Directory module.
2. Run the following command:

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

3. Run the following commands:

```
Get-ADUser -Filter * | Format-List
Get-ADUser -Filter * |
Format-List -Property GivenName, Surname
Get-ADUser -Filter * -Properties * | Format-List *
```

4. Run the following commands:

```
Get-ADUser -Filter * | Format-Table
Get-ADUser -Filter * |
Format-Table -Property GivenName, Surname
Get-ADUser -Filter * -Properties * | Format-Table
```

5. Run the following commands:

```
Get-ADOrganizationalUnit -Filter * | Format-Wide
Get-ADOrganizationalUnit -Filter * |
Format-Wide -column 3
```

6. Run the following commands:

```
Get-ADUser -Filter * | Sort-Object | Format-Wide
Get-ADUser -Filter * | Sort-Object -Property ObjectGUID | Format-Wide -Property
ObjectGUID
```

7. Run the following command:

```
Get-ADUser -Filter * | Measure-Object
```

### ► Task 3: Use a script to create new users in the domain by using a CSV-based file

1. On LON-DC1, browse to the **Start** screen and then type **Notepad.exe**. Press Enter.
2. In the Notepad window, on the **File** menu, click **Open**. Locate **E:\ModXA\Democode\LabUsers.Csv**. You will need to change the file type to All Files.
3. Close Notepad.
4. In Windows PowerShell ISE, click **File** and then click **Open**. Locate **E:\ModXA\Democode\LabUsers.ps1**. Click **Open**.

5. On line 13, modify the **\$OU** variable to read:  
**\$OU = "ou=sales, dc=adatum, dc=com"**
6. Press F5 to run the **LabUsers.ps1** script.
7. In the Console pane, type the following to verify that Luka Abrus, Marcel Truempy, Andy Brauninger, and Cynthia Cary were created:

```
Get-ADUser -Filter * -SearchBase "OU=Sales,DC=Adatum,DC=com"
```

► **Task 4: Create a script to modify the address of a user based on the day of the week**

1. If it is necessary, open Windows PowerShell ISE as an administrator and import the Active Directory module.
2. In Windows PowerShell ISE, on the **File** menu, click **Open**. Locate **E:\ModXA\Democode\Using If Statements.ps1**. Click **Open**.
3. Verify that line 9 reads:  
**\$Admin = Get-ADUser -identity Administrator -Properties StreetAddress**
4. Press F5 to run the script. Run the script a second time to view the changes.

**Results:** After completing this lab, you will have explored the Active Directory Windows PowerShell module, experienced formatting output in Windows PowerShell, used a Windows PowerShell script to create users, and used Windows PowerShell conditional loops to modify Active Directory properties.

### Exercise 3: Managing Servers by Using Windows PowerShell 3.0

► **Task 1: Install and configure Windows PowerShell Web Access**

1. On LON-DC1, open Windows PowerShell ISE, in the Console pane type the following, and then press Enter.

```
Install-WindowsFeature -Name WindowsPowerShellWebAccess -ComputerName LON-DC1 -  
IncludeManagementTools -Restart
```

2. In the Console pane, type **Install-PswaWebApplication -UseTestCertificate** and then press Enter.
3. In the Console pane, type **Add-PswaAuthorizationRule -UserName Adatum\Administrator -ComputerName \* -ConfigurationName \*** and then press Enter.

► **Task 2: Verify Windows PowerShell Web Access configuration**

1. Browse to the Start screen and then click **Internet Explorer**.
2. In the Address bar, type the following URL and then press Enter:  
**https://LON-DC1/pswa**
3. Click **Continue to this website**.
4. Sign in to Windows PowerShell Web Access by using the following information:
  - User: **Administrator**
  - Password: **Pa\$\$w0rd**
  - Computer: **LON-DC1**

5. In the Windows PowerShell Web Access command shell, type **Get-EventLog System -Newest 5** and then press Enter.
6. Type the following in the Windows PowerShell Web Access command shell:

```
Invoke-Command -ScriptBlock { Get-Eventlog Security -Newest 20 } -ComputerName LON-DC1,LON-SVR2
```

**Results:** After this exercise, you will have performed one to many management of remote servers by using Windows PowerShell, installed and configured Windows PowerShell Web Access, and managed servers by using Windows PowerShell Web Access.

#### ► To prepare for the next module

When you are finished the lab, revert the virtual machines to their initial state. To do this, perform the following steps:

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20417A-LON-SVR1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20417A-LON-SVR2** and **20417A-LON-DC1**.



## Module 4: Managing Storage for Windows Server 2012

### Lab A: Managing Storage for Servers Based on Windows Server 2012

#### Exercise 1: Configuring iSCSI Storage

##### ► Task 1: Install the iSCSI Target feature

1. Log on to **LON-DC1** with username of **Adatum\Administrator** and the password of **Pa\$\$w0rd**.
2. In Server Manager, click **Add roles and features**.
3. In the Add Roles and Features Wizard, on the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, click **Next**.
5. On the **Select destination server** page, make sure that **Select server from the server pool** is selected, and then click **Next**.
6. On the **Select server roles** page, expand **File And Storage Services (Installed)**, expand **File and iSCSI Services**, select the **iSCSI Target Server** check box, and then click **Next**.
7. On the **Select features** page, click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. When installation is complete, click **Close**.

##### ► Task 2: Configure the iSCSI targets

1. On LON-DC1, in Server Manager, in the navigation pane, click **File and Storage Services**.
2. In the File and Storage Services pane, click **iSCSI**.
3. In the iSCSI VIRTUAL DISKS pane, click **TASKS**, and then in the **TASKS** drop-down list, select **New iSCSI Virtual Disk**.
4. In the New iSCSI Virtual Disk Wizard, on the **Select iSCSI virtual disk location** page, under **Storage location**, click **C:**, and then click **Next**.
5. On the **Specify iSCSI virtual disk name** page, in the **Disk name** box, type **iSCSIDisk1**, and then click **Next**.
6. On the **Specify iSCSI virtual disk size** page, in the **Size** box, type **5**, make sure **GB** is selected in the drop-down list, and then click **Next**.
7. On the **Assign iSCSI target** page, click **New iSCSI target**, and then click **Next**.
8. On the **Specify target name** page, in the **Name** box, type **lon-svr2**, and then click **Next**.
9. On the **Specify access servers** page, click **Add**.
10. In the **Select a method to identify the initiator** dialog box, click **Enter a value for the selected type**, in the **Type** drop-down list, select **IP Address**, in the **Value** box, type **172.16.0.22**, and then click **OK**.
11. On the **Specify access servers** page, click **Add**.
12. In the **Select a method to identify the initiator** dialog box, click **Enter a value for the selected type**, in the **Type** drop-down list, select **IP Address**, in the **Value** box, type **131.107.0.2**, and then click **OK**.

13. On the **Specify access servers** page, click **Next**.
14. On the **Enable Authentication** page, click **Next**.
15. On the **Confirm selections** page, click **Create**.
16. On the **View results** page, wait until the creation is completed, and then click **Close**.
17. In the iSCSI VIRTUAL DISKS pane, click **TASKS**, and then in the **TASKS** drop-down list, select **New iSCSI Virtual Disk**.
18. In the New iSCSI Virtual Disk Wizard, on the **Select iSCSI virtual disk location** page, under **Storage location**, click **C:**, and then click **Next**.
19. On the **Specify iSCSI virtual disk name** page, in the **Disk name** box, type **iSCSIDisk2**, and then click **Next**.
20. On the **Specify iSCSI virtual disk size** page, in the **Size** box, type **5**, make sure **GB** is selected in the drop-down list, and then click **Next**.
21. On the **Assign iSCSI target** page, click **lon-svr2**, and then click **Next**.
22. On the **Confirm selections** page, click **Create**.
23. On the **View results** page, wait until the creation is completed, and then click **Close**.
24. In the iSCSI VIRTUAL DISKS pane, click **TASKS** and then in the **TASKS** drop-down list, select **New iSCSI Virtual Disk**.
25. In the New iSCSI Virtual Disk Wizard, on the **Select iSCSI virtual disk location** page, under **Storage**, click **C:**, and then click **Next**.
26. On the **Specify iSCSI virtual disk name** page, in the **Disk name** box, type **iSCSIDisk3**, and then click **Next**.
27. On the **Specify iSCSI virtual disk size** page, in the **Size** box, type **5**, make sure **GB** is selected in the drop-down list, and then click **Next**.
28. On the **Assign iSCSI target** page, click **lon-svr2**, and then click **Next**.
29. On the **Confirm selections** page, click **Create**.
30. On the **View results** page, wait until the creation is completed, and then click **Close**.
31. In the iSCSI VIRTUAL DISKS pane, click **TASKS** and then in the **TASKS** drop-down list, select **New iSCSI Virtual Disk**.
32. In the New iSCSI Virtual Disk Wizard, on the **Select iSCSI virtual disk location** page, under **Storage**, click **C:**, and then click **Next**.
33. On the **Specify iSCSI virtual disk name** page, in the **Disk name** box, type **iSCSIDisk4**, and then click **Next**.
34. On the **Specify iSCSI virtual disk size** page, in the **Size** box, type **5**, make sure **GB** is selected in the drop-down list, and then click **Next**.
35. On the **Assign iSCSI target** page, click **lon-svr2**, and then click **Next**.
36. On the **Confirm selections** page, click **Create**.
37. On the **View results** page, wait until the creation is completed, and then click **Close**.
38. In the iSCSI VIRTUAL DISKS pane, click **TASKS** and then in the **TASKS** drop-down list, click **New iSCSI Virtual Disk**.

39. In the New iSCSI Virtual Disk Wizard, on the **Select iSCSI virtual disk location** page, under **Storage**, click **C:**, and then click **Next**.
40. On the **Specify iSCSI virtual disk name** page, in the **Disk name** box, type **iSCSIDisk5**, and then click **Next**.
41. On the **Specify iSCSI virtual disk size** page, in the **Size** box, type **5**, make sure **GB** is selected in the drop-down list, and then click **Next**.
42. On the **Assign iSCSI target** page, click **lon-svr2**, and then click **Next**.
43. On the **Confirm selections** page, click **Create**.
44. On the **View results** page, wait until the creation is completed, and then click **Close**.

► **Task 3: Configure MPIO**

1. Log on to **LON-SVR2** with username of **Adatum\Administrator** and the password of **Pa\$\$w0rd**.
2. In Server Manager, click **Add roles and features**.
3. In the Add Roles and Features Wizard, on the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, click **Next**.
5. On the **Select destination server** page, make sure that **Select server from the server pool** is selected, and then click **Next**.
6. On the **Select server roles** page, click **Next**.
7. On the **Select features** page, click **Multipath I/O**, and then click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. When installation is complete, click **Close**.
10. In Server Manager, on the menu bar, click **Tools** and then in the **Tools** drop-down list, select **iSCSI Initiator**.
11. In the **Microsoft iSCSI** dialog box, click **Yes**.
12. In the **iSCSI Initiator Properties** dialog box, on the **Targets** tab, in the **Target** box, type **LON-DC1**, and then click **Quick Connect**. In the **Quick Connect** box, click **Done**.
13. Click **OK** to close the **iSCSI Initiator Properties** dialog box.
14. In Server Manager, on the menu bar, click **Tools**, and then in the **Tools** drop-down list, select **MPIO**.
15. In **MPIO Properties** dialog box, click the **Discover Multi-Paths** tab.
16. Select the **Add support for iSCSI devices** check box, and then click **Add**. When you are prompted to reboot the computer, click **Yes**.
17. After the computer restarts, log on to **LON-SVR2** with username of **Adatum\Administrator** and password of **Pa\$\$w0rd**.
18. In Server Manager, on the menu bar, click **Tools**, and then in the **Tools** drop-down list, select **MPIO**.
19. In the **MPIO Properties** dialog box, on the **MPIO Devices** tab, notice that additional **Device Hardware ID MSFT2005iSCSIBusType\_0x9** is added to the list.
20. Click **OK** to close the **MPIO Properties** dialog box.

**► Task 4: Connect to and configure the iSCSI targets**

1. On LON-SVR2, in Server Manager, on the menu bar, click **Tools** and then in the **Tools** drop-down list, select **iSCSI Initiator**.
2. In the **iSCSI Initiator Properties** dialog box, on the **Targets** tab, click **Disconnect**.
3. In the **Disconnect From All Sessions** dialog box, click **Yes**.
4. In the **iSCSI Initiator Properties** dialog box, on the **Targets** tab, click **Connect**.
5. In the Connect to Target window, click **Enable multi-path**, verify that the **Add this connection to the list of Favorite Targets** check box is selected, and then click the **Advanced** button.
6. In the **Advanced Settings** dialog box, on the **General** tab, change the **Local Adapter** from **Default** to **Microsoft iSCSI Initiator**. In the **Initiator IP** drop-down list, click **172.16.0.22** and in the **Target Portal IP** drop-down list, click **172.16.0.10 / 3260**.
7. In the **Advanced Settings** dialog box, click **OK**.
8. In the Connect to Target window, click **OK**.
9. In the **iSCSI Initiator Properties** dialog box, on the **Targets** tab, click **Connect**.
10. In Connect to Target window, click **Enable multi-path**, verify that the **Add this connection to the list of Favorite Targets** check box is selected, and then click the **Advanced** button.
11. In the **Advanced Settings** dialog box, on the **General** tab, change the **Local Adapter** from **Default** to **Microsoft iSCSI Initiator**. In the **Initiator IP** drop-down list, select **131.107.0.2** and in the **Target Portal IP** drop-down list, select **131.107.0.1 / 3260**.
12. In the **Advanced Settings** dialog box, click **OK**.
13. In the Connect to Target window, click **OK**.
14. In the **iSCSI Initiator Properties** dialog box, click the **Volumes and Devices** tab.
15. In the **iSCSI Initiator Properties** dialog box, on the **Volumes and Devices** tab, click **Auto Configure**.
16. In the **iSCSI Initiator Properties** dialog box, click the **Targets** tab.
17. In the **Targets** list, select **iqn.1991-05.com.microsoft:lon-dc1-lon-svr2-target**, and then click **Devices**.
18. In the **Devices** dialog box, click the **MPIO** button.
19. Verify that in **Load balance policy**, **Round Robin** is selected. Under **This device has the following paths**, notice that two paths are listed. Select the first path and then click the **Details** button.
20. Note the IP address of the **Source** and **Target** portals, and then click **OK**.
21. Select the second path and then click the **Details** button.
22. Verify that the Source IP address is of the second network adapter, and then click **OK**.
23. Click **OK** to close the **Device Details** dialog box.
24. Click **OK** to close the **Devices** dialog box.
25. Close the iSCSI Initiator Properties dialog box.

**Results:** After completing this exercise, you will have configured and connected to iSCSI targets.

## Exercise 2: Configuring a Redundant Storage Space

### ► Task 1: Create a storage pool by using the iSCSI disks attached to the server

1. On LON-SVR2, open Server Manager by clicking the icon on the taskbar.
2. In the navigation pane, click **File and Storage Services**, and then in the Servers pane, click **Storage Pools**.
3. In the STORAGE POOLS pane, click **TASKS**, and then in the **TASKS** drop-down list, click **New Storage Pool**.
4. In the New Storage Pool Wizard window, on the **Before you begin** page, click **Next**.
5. On the **Specify a storage pool name and subsystem** page, in the **Name** box, type **StoragePool1**, and then click **Next**.
6. On the **Select physical disks for the storage pool** page, click all five physical disks, and then click **Next**.
7. On the **Confirm selections** page, click **Create**.
8. On the **View results** page, wait until the creation is completed, then click **Close**.

### ► Task 2: Create a 3-way mirrored disk

1. On LON-SVR2, in Server Manager, in the STORAGE POOLS pane, click **StoragePool1**.
2. In the VIRTUAL DISKS pane, click **TASKS**, and then from the **TASKS** drop-down list click **New Virtual Disk**.
3. In the New Virtual Disk Wizard window, on the **Before you begin** page, click **Next**.
4. On the **Select the server and storage pool** page, click **StoragePool1**, and then click **Next**.
5. On the **Specify the virtual disk name** page, in the **Name** box, type **Mirrored vDisk**, and then click **Next**.
6. On the **Select the storage layout** page, in the **Layout** list, select **Mirror**, and then click **Next**.
7. On the **Configure the resiliency settings** page, click **Three-way mirror**, and then click **Next**.
8. On the **Specify the provisioning type** page, click **Thin**, and then click **Next**.
9. On the **Specify the size of the virtual disk** page, in the **Virtual disk size** box, type **10**, and then click **Next**.
10. On the **Confirm selections** page, click **Create**.
11. On the **View results** page, wait until the creation is completed, make sure **Create a volume when this wizard closes** is checked, and then click **Close**.
12. In the New Volume Wizard window, on the **Before you begin** page, click **Next**.
13. On the **Select the server and disk** page, in the Disk pane, click the virtual disk that is called **Mirrored vDisk**, and then click **Next**.
14. On the **Specify the size of the volume** page, click **Next** to confirm the default selection.
15. On the **Assign to a drive letter or folder** page, make sure **E** is selected in the **Drive letter** drop-down list, and then click **Next**.
16. On the **Select file system settings** page, in the **File system** drop-down list, select **ReFS**, in the **Volume label** box, type **Mirrored Volume**, and then click **Next**.

17. On the **Confirm selections** page, click **Create**.
18. On the **Completion** page, wait until the creation is completed, and then click **Close**.

► **Task 3: Copy a file to the volume and verify visibility in Windows Explorer**

1. On the Start screen, type **command prompt** and then press Enter.
2. At the command prompt, type the following command and then press Enter:

```
Copy C:\windows\system32\write.exe E:\
```

3. Close the command prompt.
4. On the taskbar, open Windows Explorer and then click **Mirrored Volume (E:)**. You should now see write.exe in the file list.
5. Close Windows Explorer.

► **Task 4: Disconnect an iSCSI disk**

1. Switch to **LON-DC1**.
2. In Server Manager, in the navigation pane, click **File and Storage Services**.
3. In the File and Storage Services pane, click **iSCSI**.
4. In the iSCSI VIRTUAL DISKS pane, in the LON-DC1 list, right-click **iSCSIDisk1.vhd**, and then click **Disable iSCSI Virtual Disk**.
5. In the Disable iSCSI Virtual Disk warning message box, click **Yes**.

► **Task 5: Verify that the file is still accessible and check the health of the virtual disk**

1. Switch to **LON-SVR2**.
2. On the taskbar, open Windows Explorer, and then click **Mirrored Volume (E:)**.
3. In the file list pane, double-click **write.exe** to make sure access to the volume is still available.
4. Close the Document - WordPad window.
5. Close Windows Explorer.
6. In Server Manager, in the STORAGE POOLS pane, on the menu bar click the **Refresh "Storage Pools"** button. Wait until all panes are refreshed. Notice the warning that appears right next to Mirrored vDisk.
7. In the VIRTUAL DISK pane, right-click **Mirrored vDisk**, in the drop-down list, select **Properties**.
8. In the Mirrored vDisk Properties window, in the navigation pane, click Health. Notice that the Health Status indicates a Warning. The Operational Status should indicate Degraded.
9. Click **OK** to close the window.

► **Task 6: Add a new iSCSI virtual disk**

1. Switch to **LON-DC1**.
2. In Server Manager, in the navigation pane, click **File and Storage Services**.
3. In the File and Storage Services pane, click **iSCSI**.
4. In the iSCSI Virtual VIRTUAL DISKS pane, click **TASKS**, and then in the **TASKS** drop-down list, select **New iSCSI Virtual Disk**.

5. In the New iSCSI Virtual Disk Wizard, on the **Select iSCSI virtual disk location** page, in the Storage location pane, click **C:**, and then click **Next**.
6. On the **Specify iSCSI virtual disk name** page, type **iSCSIDisk6**, and then click **Next**.
7. On the **Specify iSCSI virtual disk size** page, in the **Size** box, type **5**, make sure **GB** is selected in the drop-down list, and then click **Next**.
8. On the **Assign iSCSI target** page, click **lon-svr2**, and then click **Next**.
9. On the **Confirm selections** page, click **Create**.
10. On the **View results** page, wait until the creation is completed, and then click **Close**.

► **Task 7: Add the new disk to the storage pool and extend the virtual disk**

1. Switch to **LON-SVR2**.
2. In Server Manager, in the STORAGE POOLS pane, on the menu bar click the **Refresh "Storage Pools"** button.
3. In the STORAGE POOLS pane, right-click **StoragePool1**, and then in the drop-down list, select **Add Physical Disk**.
4. In the Add Physical Disk window, click **PhysicalDisk1 (LON-SVR2)**, and then click **OK**.
5. In the VIRTUAL DISKS pane, right-click **Mirrored vDisk**, and then in the drop-down list, select **Extend Virtual Disk**.
6. In the Extend Virtual Disk window, in the **New size** box, type **15**, and then click **OK**.

**Results:** After completing this exercise, you will have created a storage pool and added a new disk to the storage pool and extended the disk.

► **To prepare for the next lab**

When you finish the lab, revert the virtual machines to their initial state. To do this, perform the following steps:

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20417A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20417A-LON-SVR2**.



## Lab B: Implementing BranchCache

### Exercise 1: Performing Initial Configuration Tasks for BranchCache

#### ► Task 1: Configure LON-DC1 to use BranchCache

1. Log on to **LON-DC1** with username of **Adatum\Administrator** and the password of **Pa\$\$w0rd**.
2. Open Server Manager by clicking the icon on the taskbar.
3. Click **Add roles and features**.
4. In the Add Roles and Features Wizard, on the **Before you begin** page, click **Next**.
5. On the **Select installation type** page, click **Next**.
6. On the **Select destination server** page, make sure that **Select server from the server pool** is selected, and then click **Next**.
7. On the **Select server roles** page, expand **File And Storage Services (Installed)**, expand **File and iSCSI Services**, select the **BranchCache for Network Files** check box, and then click **Next**.
8. On the **Select features** page, click **Next**.
9. On the **Confirm installation selections** page, click **Install**.
10. After the installation has succeeded, click **Close**.
11. Click to the Start screen, type **gpedit.msc** and then press Enter.
12. In the navigation pane of the Local Group Policy Editor console, under **Computer Configuration**, expand **Administrative Templates**, expand **Network**, and then click **Lanman Server**.
13. In the **Setting** list in the Lanman Server result pane, right-click **Hash Publication for BranchCache**, and then click **Edit**.
14. In the **Hash Publication for BranchCache** dialog box, click **Enabled**, in the **Hash publication actions** list, select the **Allow hash publication only for shared folders on which BranchCache is enabled** check box, and then click **OK**.

#### ► Task 2: Simulate slow link to the branch office

1. In the navigation pane of the Local Group Policy Editor console, under **Computer Configuration**, expand **Windows Settings**, right-click **Policy-based QoS**, and then click **Create new policy**.
2. On the **Create a QoS policy** page of the Policy-based QoS Wizard, in the **Policy name** box, type **Limit to 100 KBps**, click **Specify Outbound Throttle Rate** check box, type **100**, and then click **Next**.
3. On the **This QoS policy applies to** page, click **Next**.
4. On the **Specify the source and destination IP addresses** page, click **Next**.
5. On the **Specify the protocol and port numbers** page, click **Finish**.
6. Close the Local Group Policy Editor.

#### ► Task 3: Enable a file share for BranchCache

1. Open Windows Explorer by clicking the icon on the taskbar.
2. In the Computer window, browse to Local Disk (**C:**).
3. On the menu, on the **Home** tab, click **New Folder**.
4. Type **Share** and then press Enter.



5. Right-click **Share** and then click **Properties**.
6. On the **Sharing** tab of the **Share Properties** dialog box, click **Advanced Sharing**.
7. Select the **Share this folder** check box and then click **Caching**.
8. In the **Offline Settings** dialog box, select the **Enable BranchCache** check box and then click **OK**.
9. In the **Advanced Sharing** dialog box, click **OK**.
10. In the **Share Properties** dialog box, click **Close**.
11. Click to the Start screen, type **command prompt** and then press Enter.
12. At the command prompt, type the following command and then press Enter:

```
Copy C:\windows\system32\mspaint.exe c:\share
```

13. Close the command prompt.
14. Close Windows Explorer.

#### ► Task 4: Configure client firewall rules for BranchCache

1. On LON-DC1, open Server Manager by clicking the icon on the taskbar.
2. In Server Manager, on the menu bar, click **Tools** and then select **Group Policy Management** from the Tools drop-down list.
3. In Group Policy Management, expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**, right-click **Default Domain Policy**, and then click **Edit**.
4. In the navigation pane of the Group Policy Management Editor console, under **Computer Configuration** expand **Policies**, expand **Windows Settings**, expand **Security Settings**, and then expand **Windows Firewall with Advanced Security**.
5. In the navigation pane, under **Windows Firewall with Advanced Security**, expand **Windows Firewall with Advanced Security**, and then click **Inbound Rules**.
6. On the **Action** menu of the Group Policy Management Editor console, click **New Rule**.
7. On the **Rule Type** page of the New Inbound Rule Wizard, click **Predefined**, click **BranchCache – Content Retrieval (Uses HTTP)**, and then click **Next**.
8. On the **Predefined Rules** page, click **Next**.
9. On the **Action** page, click **Finish** to create the firewall inbound rule.
10. Click **Inbound Rules**, and then on the **Action** menu of the Group Policy Management Editor console, select **New Rule**.
11. On the **Rule Type** page of the New Inbound Rule Wizard, click **Predefined**, click **BranchCache – Peer Discovery (Uses WSD)**, and then click **Next**.
12. On the **Predefined Rules** page, click **Next**.
13. On the **Action** page, click **Finish**.

**Results:** At the end of this exercise, you will have deployed BranchCache, configured a slow link, and enabled BranchCache on a file share.

## Exercise 2: Configuring BranchCache Client Computers

### ► Task 1: Configure client computers to use BranchCache in the Hosted Cache mode

1. On LON-DC1, in the navigation pane of the Group Policy Management Editor console, under **Computer Configuration**, expand **Policies**, expand **Administrative Templates**, expand **Network**, and then click **BranchCache**.
2. In the **Setting** list of the BranchCache result pane, right-click **Turn on BranchCache** and then click **Edit**.
3. In the **Turn on BranchCache** dialog box, click **Enabled** and then click **OK**.
4. In the **Setting** list of the BranchCache result pane, right-click **Set BranchCache Hosted Cache mode** and then click **Edit**.
5. In the **Set BranchCache Hosted Cache mode** dialog box, click **Enabled**, in the **Type the name of the hosted Cache server**, type **LON-SVR1.adatum.com**, and then click **OK**.
6. In the **Setting** list of the **BranchCache** result pane, right-click **Configure BranchCache for network files** and then click **Edit**.
7. In the **Configure BranchCache for network files** dialog box, click **Enabled**, in the **Type the maximum round trip network latency value (milliseconds) after which caching begins** box, type **0**, and then click **OK**. This setting is required to simulate access from a branch office and is not typically required.
8. Close the Group Policy Management Editor console.
9. Close the Group Policy Management console.
10. Start **20417A-LON-CL1**. After the computer starts, log on as **Adatum\Administrator** with the password of **Pa\$\$w0rd**.
11. On the Start screen, type **command prompt** and then press Enter.
12. At the command prompt, type the following command and then press Enter:

```
gpupdate /force
```

13. At the command prompt, type the following command and then press Enter:

```
netsh branchcache show status all
```

14. Start **20417A-LON-CL2**. After the computer starts, log on as **Adatum\Administrator** with the password of **Pa\$\$w0rd**.
15. On the Start screen, type **command prompt** and then press Enter.
16. At the command prompt, type the following command and then press Enter:

```
gpupdate /force
```

17. At the command prompt, type the following command and then press Enter:

```
netsh branchcache show status all
```

**Results:** At the end of this exercise, you will have configured the client computers for BranchCache.

### Exercise 3: Configuring BranchCache on the Branch Server

#### ► Task 1: Install the BranchCache feature on LON-SVR1

1. Start **20417A-LON-SVR1**. After the computer starts, log on as **Adatum\Administrator** with the password of **Pa\$\$w0rd**.
2. Open Server Manager by clicking the icon on the taskbar.
3. Click **Add roles and features**.
4. In the Add Roles and Features Wizard, on the **Before you begin** page, click **Next**.
5. On the **Select installation type** page, click **Next**.
6. On the **Select destination server** page, make sure that **Select server from the server pool** is selected, and then click **Next**.
7. On the **Select server roles** page, expand **File And Storage Services (Installed)**, expand **File and iSCSI Services**, click **BranchCache for Network Files** check box.
8. On the **Select server roles** page, click **Next**.
9. On the **Select features** page, click **BranchCache**, and then click **Next**.
10. On the **Confirm installation selections** page, click **Install**.
11. Close Server Manager.

#### ► Task 2: Start the BranchCache host server

1. Switch to **LON-DC1**.
2. In Server Manager, on the menu bar, click **Tools** and then select **Active Directory Users and Computers** from the Tools drop-down list.
3. Right-click **Adatum.com**, point to **New**, and then click **Organizational Unit**.
4. In the New Object - Organization Unit window, type **BranchCacheHost** and then click **OK**.
5. Click the **Computers** container.
6. Click **LON-SVR1** and drag it to **BranchCacheHost**.
7. Click **Yes** to clear the warning about moving objects.
8. Close Active Directory Users and Computers.
9. In Server Manager, on the menu bar, click **Tools** and then select **Group Policy Management** from the Tools drop-down list.
10. Under **Domains**, expand **Adatum.com**, right-click **BranchCacheHost**, and then click **Block Inheritance**.
11. On LON-DC1, close all open windows.
12. Restart **LON-SVR1** and log on as **Adatum\Administrator** with the password of **Pa\$\$w0rd**.
13. Open Windows PowerShell by clicking the icon on the taskbar.
14. At the Windows PowerShell window, type the following cmdlet, and then press Enter:

```
Enable-BCHostedServer -RegisterSCP
```

15. At the Windows PowerShell window, type the following cmdlet, and then press Enter:

```
Get-BCStatus
```

16. Close the Windows PowerShell.



**Note:** BranchCache is only available on Windows 8 Enterprise edition. This edition was not available when this course was created, so the BranchCache verification steps are not included in this lab.

**Results:** At the end of this exercise, you will have enabled the BranchCache server in the branch office.

### ► To prepare for the next module

When you finish the lab, revert the virtual machines to their initial state. To do this, perform the following steps:

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20417A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20417A-LON-SVR1**, **20417A-LON-CL1**, and **20417A-LON-CL2**.

## Module 5: Implementing Network Services

# Lab: Implementing Network Services

### Exercise 1: Configure new features in DNS and DHCP

#### ► Task 1: Configure DNSSEC

1. On LON-DC1, in Server Manager, click **Tools**, and then click **DNS** on the drop-down list.
2. Expand **LON-DC1**, expand **Forward Lookup Zones**, and then select and right-click **Adatum.com**.
3. On the shortcut menu, click **DNSSEC > Sign the Zone**.
4. In the Zone Signing Wizard, click **Next**.
5. Select **Customize zone signing parameters**, and then click **Next**.
6. On the Key Master screen, ensure that **LON-DC1** is the **Key Master**. Click **Next**.
7. On the Key Signing Key (KSK) screen, click **Next**.
8. On the Key Signing Key (KSK) screen, click **Add**.
9. On the New Key Signing Key (KSK) screen, click **OK**.
10. On the Key Signing Key (KSK) screen, click **Next**.
11. On the Zone Signing Key (ZSK) screen, click **Next**.
12. On the Zone Signing Key (ZSK) screen, click **Add**.
13. On the New Zone Signing Key (ZSK) screen, click **OK**.
14. On the Zone Signing Key (ZSK) screen, click **Next**.
15. On the Next Secure (NSEC) screen, click **Next**.
16. On the Trust Anchors screen, check **Enable the distribution of trust anchors for this zone**. Click **Next**.
17. On the Signing and Polling Parameters screen, click **Next**.
18. On the DNS Security Extensions (DNSSEC) screen, click **Next**.
19. Click **Finish**.
20. Expand **Trust Points**, expand **com**, and click **Adatum**. Ensure that the DNSKEY resource records exist and that their status is valid.
21. Close the DNS Manager console.
22. In Server Manager, click **Tools**, and then on the drop-down list, click **Group Policy Management**.
23. Expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**, right-click the **Default Domain Policy**, and then click **Edit**.
24. In the Group Policy Management Editor, under Computer Configuration, expand **Policies**, expand **Windows Settings**, and then click the **Name Resolution Policy** folder.
25. To apply the rule to the suffix of the namespace, in the Create Rules section, in the **Suffix** field, type **Adatum.com**.
26. On the **DNSSEC** tab, click **Enable DNSSEC in this rule**.

27. Check **Require DNS clients to check that the name and address data has been validated by the DNS server**, and then click **Create**.
28. Close the Group Policy Management Editor and Group Policy Management console.

► **Task 2: Configure DHCP Name Protection**

1. In Server Manager, click **Tools**, and then on the drop-down list, click **DHCP**.
2. Expand **LON-DC1.adatum.com**.
3. Select and then right-click **IPv4**, and then click **Properties**.
4. Click the **DNS** tab.
5. In the **Name Protection** section, click **Configure**.
6. Check **Enable Name Protection**, and then click **OK**.
7. To close the **Properties** dialog box, click **OK**.

► **Task 3: Configure DHCP Failover**

1. On LON-SVR1, in Server Manager, click **Tools**, and then on the drop-down list, click **DHCP**. Note the server is authorized but no scopes are configured.
2. Switch to **LON-DC1**.
3. In the DHCP Management console right-click the **IPv4** node, and then click **Configure Failover**.
4. In the Configuration Failover Wizard, click **Next**.
5. On the Specify a partner server to use for failover screen, enter **172.16.0.21** in the **Partner Server** field, and then click **Next**.
6. On the Create a new failover relationship screen, in the **Relationship Name** field, type **Adatum**.
7. In the **Maximum Client Lead Time** field, set the hours to **zero**, and set the minutes to **15**.
8. Ensure the **Mode** field is set to **Load balance**.
9. Ensure the **Load Balance Percentage** is set to **50%**.
10. Check **State Switchover Interval**.
11. In the **Enable Message Authentication Shared Secret** field, type **Pa\$\$w0rd** and then click **Next and then click Finish**.
12. Click **Close**.
13. Switch to **LON-SVR1**. Notice that the IPv4 node is active.
14. Expand the **IPv4** node and expand the **Adatum Scope**.
15. Click the **Address Pool** node. Notice that the address pool is configured.
16. Click the **Scope Options** node. Notice that the scope options are configured.
17. Close the DHCP console on both LON-DC1 and LON-SVR1.

**Results:** After completing this exercise you will be able to configure DNSSEC, configure DHCP name protection, and configure and verify DHCP failover.

## Exercise 2: Configuring IP Address Management

### ► Task 1: Install the IPAM Feature

1. On LON-SVR2, in Server Manager, click **Add roles and features**.
2. In the Add Roles and Features Wizard, click **Next**.
3. On the Select installation type screen, click **Next**.
4. On the Select destination server screen, click **Next**.
5. On the Select server roles screen, click **Next**.
6. On the Select features screen, check **IP Address Management (IPAM) Server**.
7. In the **Add features that are required for IP Address Management (IPAM) Server** pop-up, click **Add Features**, and then click **Next**.
8. On the Confirm installation selections, click **Install**.
9. Close the wizard when completed.

### ► Task 2: Configure IPAM Related GPOs

1. On LON-SVR2, in the Server Manager, click **IPAM**.
2. In the IPAM Overview pane, after step 1 shows that LON-SVR2 is connected, click **Provision the IPAM server**.
3. In the Provision IPAM Wizard, click **Next**.
4. On the Select provisioning method screen, select the **Group Policy Based** method, type **IPAM** in the **GPO name prefix** field, and then click **Next**.
5. On the Confirm the Settings screen, click **Apply**.
6. When provisioning has completed, click **Close**.

### ► Task 3: Configure IP Management Server Discovery

1. On the IPAM Overview pane, click **Configure server discovery**.
2. To add the Adatum.com domain, in the **Configure Server Discovery** dialog box, click **Add**, and then click **OK**.
3. On the IPAM Overview pane, click **Start server discovery**.
4. In the yellow banner, to determine the discovery status, click the **More** link. Discovery will take a few minutes to complete.
5. To return to the IPAM pane, close the **Overview Tasks Details** dialog box.

### ► Task 4: Configure Managed Servers

1. From the IPAM Overview pane, click **Select or add servers to manage and verify IPAM access**.



**Note:** Notice that for LON-SVR1 and LON-DC1, the IPAM Access Status is Blocked. Scroll down to the Details View and note the status report. This is because the IPAM server has not yet been granted permission to manage LON-SVR1 or LON-DC1 by using Group Policy.

2. On the task bar click the **Windows PowerShell** icon.

3. Type the following command at the PowerShell prompt and then press Enter:

```
Invoke-IPAMGpoProvisioning -Domain Adatum.com -GpoPrefixName IPAM -IpamServerFqdn  
LON-SVR2.adatum.com
```

4. When you are prompted to confirm the action, press Enter. It will take a few moments to complete.
5. Return to Server Manager.
6. In the details pane of the IPAM Server Inventory, right-click **LON-DC1**, and then click **Edit Server**.
7. In the **Add or Edit Server** dialog box, set the **Manageability status** field to **Managed**, and then click **OK**.
8. Repeat steps 6 and 7 to configure **LON-SVR1** to be managed.
9. Switch to **LON-DC1**.
10. On the task bar click **Windows PowerShell**.
11. Type **gpupdate /force**, and then press Enter.
12. Switch to **LON-SVR1**.
13. On the task bar click **Windows PowerShell**.
14. Type **gpupdate /force**, and then press Enter.
15. Switch back to **LON-SVR2** and right-click **LON-DC1**, then click **Refresh Server Access Status**. This may take a few minutes to complete.
16. Repeat step 15 to refresh the status for **LON-SVR1**.
17. Refresh the page by clicking the **Refresh** icon on the top menu bar until status shows an IPAM Access Status Unblocked.
18. From the IPAM Overview pane, click **retrieve data from managed servers**. This action will take several moments to complete.

► **Task 5: Configure and Verify a New DHCP Scope with IPAM**

1. In the IPAM navigation pane, under **MONITOR AND MANAGE**, click **DNS and DHCP Servers**. **Refresh the console pane until all objects show Running.**
2. In the details pane, right-click the instance of **LON-DC1.Adatum.com** that holds the DHCP server role.
3. On the shortcut menu, click **Create DHCP Scope**.
4. In the **Create DHCP Scope** dialog box, in the **Scope Name** field, type **TestScope**.
5. Type **10.0.0.10** in the **Start IP address** field.
6. Type **10.0.0.100** in the **End IP address** field.
7. In the Create details pane click **Options**.
8. In the Configure options pane, click the drop-down arrow of the **Option** field, and then select option **003 Router**.
9. In the **Values** section click into the **IP Address** field and type **10.0.0.1**, click **Add to list**, and then click **OK**.
10. Switch to **LON-DC1**.
11. In the Server Manager toolbar, click **Tools** and then click **DHCP**.



12. In the DHCP console expand **LON-DC1.Adatum.com** and then expand **IPv4** and confirm the **TestScope** exists.
13. Right-click the **TestScope** and then click **Deactivate**. Click **Yes**.
14. Close the DHCP console.
15. On LON-SVR2, close all open windows.

**Results:** After completing this exercise you will be able to install and configure the IPAM feature, configure IPAM related GPOs, configure IP Management server discovery, configure managed servers, and configure and verify a new DHCP scope with IPAM.

### Exercise 3: Configuring NAP

#### ► Task 1: Configure Server and Client Certificate Requirements

1. On LON-SVR2, move the mouse to the lower right corner, click the **Search** icon on the flyout menu, type **MMC .EXE**, and press Enter.
2. In the Console1 window, click **File**, and then click **Add/Remove Snap-in**.
3. In the **Add or Remove Snap-ins** dialog box, click **Certificates** and then click **Add**.
4. In the **Certificates snap-in** dialog box, select **Computer account**, and then click **Next**.
5. In the **Select Computer** dialog box, click **Finish**, and then click **OK**.
6. In the console tree, expand **Certificates**, right-click **Personal**, point to **All Tasks**, and then click **Request New Certificate**.
7. In the **Certificate Enrollment** dialog box, click **Next**.
8. On the **Select Certificate Enrollment Policy** page, click **Active Directory Enrollment Policy** and then click **Next**.
9. Select the **Computer** check box and then click **Enroll**.
10. Verify the status of certificate installation as Succeeded and then click **Finish**.
11. Close the Console1 window. When you are prompted to save console settings, click **No**.
12. Log on to **LON-CL1** as **Adatum/Administrator** with a password of **Pa\$\$w0rd**.
13. Move the mouse to the lower right corner and then click the **Search** icon on the flyout menu, type **MMC**, and press Enter.
14. In the Console1 window click **File** and then click **Add/Remove Snap-in**.
15. In the **Add or Remove Snap-ins** dialog box click **Certificates** and then click **Add**.
16. In the **Certificates snap-in** dialog box select **Computer** account and then click **Next**.
17. In the **Select Computer** dialog box click **Finish** and then click **OK**.
18. In the console tree, expand **Certificates**, right-click **Personal**, point to **All Tasks**, and then click **Request New Certificate**.
19. In the **Certificate Enrollment** dialog box appears click **Next**.
20. On the **Select Certificate Enrollment Policy** page, click **Active Directory Enrollment Policy** and then click **Next**.
21. Select the **Computer** check box and then click **Enroll**.

22. Verify the status of certificate installation as Succeeded and then click **Finish**.
23. In the console tree, expand **Certificates**, right-click **Personal**, point to **All Tasks**, and then click **Request New Certificate**.
24. In the **Certificate Enrollment** dialog box appears click **Next**.
25. On the **Select Certificate Enrollment Policy** page, click **Active Directory Enrollment Policy** and then click **Next**.
26. Select the **Computer** check box, and then click **Enroll**.
27. Verify the status of certificate installation as Succeeded and then click **Finish**.
28. Close the Console1 window. When you are prompted to save console settings, click **No**.
29. Log on to **LON-CL1** as **Adatum/Administrator** with a password of **Pa\$\$w0rd**.
30. On the **Start** screen, type **MMC** and press Enter.
31. In the Console1 window click **File** and then click **Add/Remove Snap-in**.
32. In the **Add or Remove Snap-ins** dialog box, click **Certificates** and then click **Add**.
33. In the **Certificates snap-in** dialog box, select **Computer account** and then click **Next**.
34. In the **Select Computer** dialog box, click **Finish** and then click **OK**.
35. In the console tree, expand **Certificates**, right-click **Personal**, point to **All Tasks**, and then click **Request New Certificate**.
36. In the **Certificate Enrollment** dialog box appears click **Next**.
37. On the **Select Certificate Enrollment Policy** page, click **Active Directory Enrollment Policy** and then click **Next**.
38. Select the **Computer** check box and then click **Enroll**.
39. Verify the status of certificate installation as Succeeded and then click **Finish**.
40. Close the Console1 window. When you are prompted to save console settings, click **No**.

► **Task 2: Install the Network Policy Server Role**

1. On LON-SVR2, switch to Server Manager.
2. Click **Add roles and features**.
3. In the Add Roles and Features Wizard, click **Next**.
4. On the **Select installation type** page, click **Next**.
5. On the **Select destination server** page, click **Next**.
6. On the **Select server roles** page, check **Network Policy and Access Services**.
7. In the **Add Roles and Features Wizard** dialog box, click **Add Features** and then click **Next**.
8. On the **Select features** page, click **Next**.
9. On the **Network Policy and Access Services** page, click **Next**.
10. On the **Select role services** page, check **Network Policy Server**. Click **Next**.
11. On the **Confirm installation selections** page, click **Install**.
12. When the installation is succeeded click **Close**.

### ► Task 3: Configure Health Policies

1. On LON-SVR2, in Server Manager, click **Tools** and then click **Network Policy Server**.
2. Expand **Network Access Protection**, expand **System Health Validators**, expand **Windows Security Health Validator**, and then click **Settings**.
3. In the right pane under **Name**, double-click **Default Configuration**.
4. On the **Windows 8 Release Preview/Windows 7/Windows Vista** selection, clear all check boxes except the **A firewall is enabled for all network connections** check box, and then click **OK**.
5. Expand **Policies**.
6. Right-click **Health Policies** and then click **New**.
7. In the **Create New Health Policy** dialog box, under Policy name, type **Compliant**.
8. Under **Client SHV checks**, verify that **Client passes all SHV checks** is selected.
9. Under **SHVs used in this health policy**, select the **Windows Security Health Validator** check box, and then click **OK**.
10. Right-click **Health Policies**, and then click **New**.
11. In the **Create New Health Policy** dialog box, under Policy Name, type **Noncompliant**.
12. Under **Client SHV checks**, select **Client fails one or more SHV checks**.
13. Under **SHVs used in this health policy**, select the **Windows Security Health Validator** check box, and then click **OK**.

### ► Task 4: Configure Network Policies for Compliant and Noncompliant Computers

1. Under **Policies** click **Network Policies**.
2. Disable the two default policies found under **Policy Name** by right-clicking the policies and then clicking **Disable**.
3. Right-click **Network Policies** and then click **New**.
4. In the Specify Network Policy Name and Connection Type window, in the **Policy name** field, type **Compliant-Full-Access** and then click **Next**.
5. In the Specify Conditions window, click **Add**.
6. In the **Select condition** dialog box, scroll down and double-click **Health Policies**.
7. In the **Health Policies** dialog box, under **Health policies**, select **Compliant**, and then click **OK**.
8. In the Specify Conditions window, verify that **Health Policy** is specified under **Conditions with a value of Compliant** and then click **Next**.
9. In the Specify Access Permission window, verify that **Access granted** is selected.
10. Click **Next** three times.
11. In the Configure Settings window, click **NAP Enforcement**. Verify that **Allow full network access** is selected and then click **Next**.
12. In the Completing New Network Policy window, click **Finish**.
13. Right-click **Network Policies** and then click **New**.
14. In the Specify Network Policy Name and Connection Type window, in the **Policy name** field, type **Noncompliant-Restricted** and then click **Next**.
15. In the Specify Conditions window, click **Add**.

16. In the **Select condition** dialog box, scroll down and double-click **Health Policies**.
17. In the **Health Policies** dialog box, under **Health policies**, select **Noncompliant** and then click **OK**.
18. In the Specify Conditions window, under **Conditions**, verify that **Health Policy** is specified with a value of **Noncompliant** and then click **Next**.
19. In the Specify Access Permission window, verify that **Access granted** is selected.



**Note:** A setting of Access granted does not mean that noncompliant client computers are granted full network access. It specifies that the policy should continue to evaluate the client computers that match these conditions.

20. Click **Next** three times.
21. In the Configure Settings window, click **NAP Enforcement**. Select **Allow limited access** and clear the **Enable auto-remediation of client computers** check box.
22. In the Configure Settings window, click **IP Filters**.
23. Under **IPv4**, click **Input Filters** and then click **New**.
24. In the **Add IP Filter** dialog box, select **Destination network**. Type **172.16.0.10** next to **IP address** and then type **255.255.255.255** next to **Subnet mask**. This step ensures that traffic from noncompliant client computers can reach only LON-DC1.
25. Click **OK** to close the **Add IP Filter** dialog box and then select **Permit only the packets listed below** in the **Inbound Filters** dialog box and then click **OK**.
26. Under **IPv4**, click **Output Filters** and then click **New**.
27. In the **Add IP Filter** dialog box, select **Source network**. Type **172.16.0.10** next to **IP address** and then type **255.255.255.255** next to **Subnet mask**.
28. Click **OK** to close the **Add IP Filter** dialog box and then in the **Outbound Filters** dialog box select **Permit only the packets listed below**. This step ensures that only traffic from LON-DC1 can be sent to noncompliant client computers.
29. To close the **Outbound Filters** dialog box, click **OK**.
30. In the Configure Settings window click **Next** and then click **Finish**.

#### ► Task 5: Configure Connection Request Policies for VPN

1. Click **Connection Request Policies**.
2. Disable the default Connection Request policy named **Use Windows authentication for all users** by right-clicking the policy and then clicking **Disable**.
3. Disable the default RRAS policy by right-clicking the **Microsoft Routing and Remote Access Service Policy** and then click **Disable**.
4. Right-click **Connection Request Policies** and then click **New**.
5. In the Specify Connection Request Policy Name and Connection Type window, under **Policy name**, type **VPN Connections**.
6. Under **Type of network access server**, select **Remote Access Server (VPN-Dial up)** and then click **Next**.
7. In the Specify Conditions window, click **Add**.

8. In the Select Condition window, scroll down and double-click **Tunnel Type**, select **PPTP**, **SSTP**, and **L2TP**. Click **OK** and then click **Next**.
9. In the Specify Connection Request Forwarding window, verify that **Authenticate requests on this server** is selected and then click **Next**.
10. In the Specify Authentication Methods window, select **Override network policy authentication settings**.
11. Under **EAP Types**, click **Add**. In the **Add EAP** dialog box, under **Authentication methods**, click **Microsoft: Protected EAP (PEAP)** and then click **OK**.
12. Under **EAP Types**, click **Microsoft: Protected EAP (PEAP)** and then click **Edit**.
13. Verify that **Enforce Network Access Protection** is selected and then click **OK**.
14. Click **Next** two times and then click **Finish**.
15. Close the Network Policy Server.

**Results:** After completing this exercise you will be able to configure server and client computer certificate requirements, install the NPS server role, configure health policies, configure network policies, and configure connection request policies for VPN.

## Exercise 4: Verifying the NAP Deployment

### ► Task 1: Configure Security Center

1. Log on to **LON-CL1** as **Adatum\Administrator** with a password of **Pa\$\$w0rd**.
2. Move the mouse to the lower right corner and then click the **Search** icon on the flyout menu.
3. In the **Search** box, type **gpedit.msc**, click **Apps**, and press Enter.
4. In the Local Group Policy Editor console tree, expand **Local Computer Policy /Computer Configuration/Administrative Templates/Windows Components/Security Center**.
5. Double-click **Turn on Security Center (Domain PCs only)**, click **Enabled**, and then click **OK**.
6. Close the Local Group Policy Editor.

### ► Task 2: Enable a Client NAP enforcement method

1. ON LON-CL1, move the mouse to the lower right corner and then click the **Search** icon on the flyout menu.
2. In the **Search** box type **napclcfg.msc** and press Enter.
3. In the console tree, click **Enforcement Clients**.
4. In the details pane, right-click **EAP Quarantine Enforcement Client** and then click **Enable**.
5. Close the NAP Client Configuration console.
6. Move the mouse to the lower right corner and then click the **Search** icon on the flyout menu.
7. In the **Search** box type **Services.msc** and press Enter.
8. In the **Services** list, double-click **Network Access Protection Agent**.
9. In the **Network Access Protection Agent Properties** dialog box, change the **Startup** type to **Automatic** and then click **Start**.

10. Wait for the NAP Agent service to start and then click **OK**.
11. Close the Services console.

► **Task 3: Allow ping on LON-SVR2**

1. On LON-SVR2 click **Tools** in Server Manager, and then click **Windows Firewall with Advanced Security**.
2. Click **Inbound Rules**, right-click **Inbound Rules**, and then click **New Rule**.
3. Select **Custom** and then click **Next**.
4. Select **All programs** and then click **Next**.
5. In the **Protocol type** field, click the drop-down arrow and select **ICMPv4** and then click **Customize**.
6. Select **Specific ICMP types**, select the **Echo Request** check box, click **OK**, and then click **Next**.
7. Click **Next** to accept the default scope.
8. In the Action window, verify that **Allow the connection** is selected and then click **Next**.
9. Click **Next** to accept the default profile.
10. In the Name windows, type **Allow Ping** and then click **Finish**.
11. Close the Windows Firewall with Advanced Security console.

► **Task 4: Move the client to the Internet and establish a VPN connection**

1. On LON-CL1, move the mouse to the lower right corner and then click the **Search** icon on the flyout menu.
2. In the **Search** box type **Control Panel** and press Enter.
3. Click **Network and Internet**.
4. Click **Network and Sharing Center**.
5. Click **Change Adapter Settings**.
6. Right-click **Local Area Connection** and then click **Properties**.
7. Click **Internet Protocol Version 4 (TCP/IPv4)** and then click **Properties**.
8. Click **Use the following IP address**. Next to **IP address**, type **131.107.0.20**. Next to **Subnet mask**, type **255.255.0.0**. Remove the existing Default Gateway, and do not configure the Default gateway.
9. Click **OK** and then click **Close** to close the **Local Area Connection Properties** dialog box.
10. Close the Network Connections window.
11. In Hyper-V Manager, right-click **20417A-LON-CL1** and then click **Settings**.
12. Click **Legacy Network Adapter** and then under **Network** select **Private Network 2**, click **OK**.
13. On LON-CL1, move the mouse to the lower right corner and then click the **Search** icon on the popout menu.
14. In the **Search** box type **CMD** and press Enter.
15. At the command prompt, type **ping 131.107.0.1** and press Enter.
16. Verify that a response is received.
17. Close the command prompt.
18. Return to Control Panel and then click **Network and Internet**.

19. Click **Network and Sharing Center**.
20. Click **Set up a new connection or network**.
21. On the **Choose a connection option** page, click **Connect to a workplace** and then click **Next**.
22. On the **How do you want to connect** page, click **Use my Internet connection (VPN)**.
23. Click I'll set up an **Internet connection later**.
24. On the **Type the Internet address to connect to** page, next to **Internet address**, type **131.107.0.2**. Next to **Destination name**, type **Adatum VPN**.
25. Select the **Allow other people to use this connection** check box and then click **Create**.
26. In the Network And Sharing Center window, click **Change adapter settings**.
27. Right-click the **Adatum VPN** connection, click **Properties**, and then click the **Security** tab.
28. Under **Authentication**, click **Use Extensible Authentication Protocol (EAP)**.
29. In the Microsoft: Secured password (EAP-MSCHAP v2) (encryption enabled) list, click **Microsoft: Protected EAP (PEAP) (encryption enabled)** and then click **Properties**.
30. Ensure that the **Verify the server's identity by validating the certificate** check box is already selected. Clear the **Connect to these servers** check box, and then ensure that **Secured password (EAP-MSCHAP v2)** is already selected under **Select Authentication Method**. Clear the **Enable Fast Reconnect** check box, and then select the **Enforce Network Access Protection** check box.
31. To accept these settings, click **OK** two times.
32. In the Network Connections window, right-click the Adatum VPN connection and then click **Connect/Disconnect**.
33. In the **Networks** flyout menu, click **Adatum VPN** and then click **Connect**.
34. In the **Network Authentication** dialog box, type **Administrator** in the **User Name** field and type **Pa\$\$w0rd** in the **Password** field.
35. Click **OK** and then click **Connect**.

► **Task 5: To prepare for next module**

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20417A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20417A-LON-SVR1**, **20417A-LON-SVR2** and **20417A-LON-CL1**.

**Results:** After completing this exercise you will be able to configure Security Center, enable a client computer NAP enforcement method, allow Ping on LON-SVR2, and move the client computer to the Internet and establish a VPN connection.

**MCT USE ONLY. STUDENT USE PROHIBITED**



## Module 6: Implementing DirectAccess

# Lab: Implementing DirectAccess

### Exercise 1: Configuring the DirectAccess Infrastructure

#### ► Task 1: Configure the AD DS and DNS requirements

1. Create a security group for DirectAccess client computers by performing the following steps:
  - a. Switch to **LON-DC1**.
  - b. In the Server Manager console, in the upper-right corner, click **Tools**, and then click **Active Directory Users and Computers**.
  - c. In the Active Directory Users and Computers console tree, right-click **Adatum.com**, click **New**, and then click **Organizational Unit**.
  - d. In New Object – Organizational Unit window, in the **Name** box, type **DA\_Clients OU**, and then click **OK**.
  - e. In the Active Directory Users and Computers console tree, expand **Adatum.com**, right-click **DA\_Clients OU**, click **New**, and then click **Group**.
  - f. In the **New Object - Group** dialog box, under **Group name**, type **DA\_Clients**.
  - g. Under **Group scope**, select **Global**, under **Group type**, select **Security**, and then click **OK**.
  - h. In the details pane, double-click **DA\_Clients**.
  - i. In the **DA\_Clients Properties** dialog box, click the **Members** tab, and then click **Add**.
  - j. In the **Select Users, Contacts, Computers, Service Accounts, or Groups** dialog box, click **Object Types**, select the **Computers** check box, and then click **OK**.
  - k. Under **Enter the object names to select (examples)**, type **LON-SVR3**, and then click **OK**.
  - l. Verify that **LON-SVR3** is displayed below **Members**, and then click **OK**.
  - m. Close the Active Directory Users and Computers console.
2. Configure firewall rules for ICMPv6 traffic by performing the following steps:



**Note:** It is important to configure firewall rules for ICMPv6 traffic to enable subsequent testing of DirectAccess in the lab environment.

- a. In the Server Manager console, in the upper-right corner, click **Tools**, and then click **Group Policy Management**.
- b. In the console tree, expand **Forest: Adatum.com\Domains\adatum.com**.
- c. In the console tree, right-click **Default Domain Policy**, and then click **Edit**.
- d. In the console tree of the Group Policy Management Editor, navigate to **Computer Configuration\Policies\Windows Settings\Security Settings\Windows Firewall with Advanced Security\Windows Firewall with Advanced Security**.
- e. In the console tree, click **Inbound Rules**, right-click **Inbound Rules**, and then click **New Rule**.
- f. On the **Rule Type** page, click **Custom**, and then click **Next**.
- g. On the **Program** page, click **Next**.

- h. On the **Protocols and Ports** page, under **Protocol type**, click **ICMPv6**, and then click **Customize**.
  - i. In the **Customize ICMP Settings** dialog box, click **Specific ICMP types**, select **Echo Request**, and then click **OK**.
  - j. Click **Next**.
  - k. On the **Scope** page, click **Next**.
  - l. On the **Action** page, click **Next**.
  - m. On the **Profile** page, click **Next**.
  - n. On the **Name** page, in the **Name** box, type **Inbound ICMPv6 Echo Requests**, and then click **Finish**.
  - o. In the console tree, click **Outbound Rules**, right-click **Outbound Rules**, and then click **New Rule**.
  - p. On the **Rule Type** page, click **Custom**, and then click **Next**.
  - q. On the **Program** page, click **Next**.
  - r. On the **Protocols and Ports** page, under **Protocol type**, click **ICMPv6**, and then click **Customize**.
  - s. In the **Customize ICMP Settings** dialog box, click **Specific ICMP types**, select **Echo Request**, and then click **OK**.
  - t. Click **Next**.
  - u. On the **Scope** page, click **Next**.
  - v. On the **Action** page, click **Allow the connection**, and then click **Next**.
  - w. On the **Profile** page, click **Next**.
  - x. On the **Name** page, in the **Name** box, type **Outbound ICMPv6 Echo Requests**, and then click **Finish**.
  - y. Close the Group Policy Management Editor and Group Policy Management consoles.
3. Create required DNS records by performing the following steps:
  - a. In the Server Manager console, click **Tools**, and then click **DNS**.
  - b. In the console tree of DNS Manager, expand **LON-DC1\Forward Lookup Zones\adatum.com**.
  - c. Right-click **adatum.com** and then click **New Host (A or AAAA)**.
  - d. In the **Name** box, type **nls**. In the **IP address box**, type **172.16.0.21**. Click **Add Host** and then click **OK**.
  - e. In the **New Host** dialog box, in the **Name** box, type **CRL**. In the **IP address** box, type **172.16.0.22**, and then click **Add Host**.
  - f. In the DNS dialog box informing you that the record was created, click **OK**.
  - g. In the **New Host** dialog box, click **Done**.
  - h. Close the DNS Manager console.

4. Remove ISATAP from the DNS global query block list by performing the following steps:
  - a. Move the mouse pointer to the lower-right corner, select **search** on the right menu, and then type **cmd.exe** to launch the Command Prompt window.
  - b. In the Command Prompt window, type the following command and then press Enter:

```
dnscmd /config /globalqueryblocklist wpad
```

Ensure that **Command completed successfully** message appears.

- c. Close the Command Prompt window.
5. Configure the DNS suffix on LON-SVR2 by performing the following steps:
  - a. Switch to **LON-SVR2**.
  - b. Move the mouse to the lower right corner of the screen, click **Settings**, click **Control Panel**, and then click **View network status and tasks**.
  - c. In the Network and Sharing Center window, click **Change adapter settings**.
  - d. In the Network Connection window, right-click **Local Area Connection**, and then click **Properties**.
  - e. In the Local Area Network Properties window, double-click **Internet Protocol Version 4 (TCP/IPv4)**.
  - f. In the **Internet Protocol Version 4 (TCP/IPv4)** dialog box, click **Advanced**.
  - g. On the **DNS** tab, in the **DNS suffix for this connection** box, type **Adatum.com**, and then click **OK**.
  - h. In the **Internet Protocol Version 4 (TCP/IPv4)** dialog box, click **OK**.
  - i. In the **Local Area Connection Properties** dialog box, click **OK**.
  - j. Close the Network Connections window.

## ► Task 2: Configure certificate requirements

1. To configure the CRL distribution settings by performing the following steps:
  - a. On LON-DC1, in Server Manager, on the **Tools** menu, click **Certification Authority**.
  - b. In the details pane, right-click **Adatum-LON-DC1-CA**, and then click **Properties**.
  - c. In the **Adatum-LON-DC1-CA Properties** dialog box, click the **Extensions** tab.
  - d. On the **Extensions** tab, click **Add**. In the **Location** box, type **http://crl.adatum.com/crld/**.
  - e. Under **Variable**, click **<CAName>**, and then click **Insert**.
  - f. Under **Variable**, click **<CRLNameSuffix>**, and then click **Insert**.
  - g. Under **Variable**, click **<DeltaCRLAllowed>**, and then click **Insert**.
  - h. In the **Location** box, type **.crl** at the end of the **Location** string, and then click **OK**.
  - i. Select **Include in CRLs. Clients use this to find Delta CRL locations** and **Include in the CDP extension of issued certificates**, and then click **Apply**. Click **No** in the dialog box asking you to restart Active Directory Certificate Services.
  - j. Click **Add**.
  - k. In the **Location** box, type **\\lon-svr2\crl\dist\$**.
  - l. Under **Variable**, click **<CaName>**, and then click **Insert**.

- m. Under **Variable**, click **<CRLNameSuffix>**, and then click **Insert**.
  - n. Under **Variable**, click **<DeltaCRLAllowed>**, and then click **Insert**.
  - o. In the **Location** box, type **.crl** at the end of the string, and then click **OK**.
  - p. Select **Publish CRLs to this location** and **Publish Delta CRLs to this location**, and then click **OK**.
  - q. Click **Yes** to restart Active Directory Certificate Services.
2. Duplicate the web certificate template and configure appropriate permission by performing the following steps:
- a. In the Certification Authority console, expand **Adatum-LON-DC1-CA**, right-click **Certificate Templates**, and then select **Manage**.



**Note:** Users require the Enroll permission on the certificate.

- b. In the Certificate Templates console, in the content pane, right-click the **Web Server** template, and then click **Duplicate Template**.
  - c. Click the **General** tab and in the **Template display name** box, type **Adatum Web Server Certificate**.
  - d. Click the **Request Handling** tab and select **Allow private key to be exported**.
  - e. Click the **Security** tab and then click **Authenticated Users**.
  - f. In the Permissions for Authenticated Users window, under **Allow**, click **Enroll**, and then click **OK**.
  - g. Close the Certificate Templates console.
  - h. In the Certification Authority console, right-click **Certificate Templates**, and navigate to **New/Certificate Template to Issue**.
  - i. Select **Adatum Web Server Certificate**, and then click **OK**.
  - j. Close the Certification Authority console.
3. Configure computer certificate auto-enrollment by performing the following steps:
- a. On LON-DC1, switch to Server Manager, click **Tools** on the upper-right side of the window, and then click **Group Policy Management**.
  - b. In the console tree, expand **Forest: Adatum.com**, expand **Domains**, and then expand **Adatum.com**.
  - c. In the console tree, right-click **Default Domain Policy**, and then click **Edit**.
  - d. In the console tree of the Group Policy Management Editor, navigate to **Computer Configuration\Policies\Windows Settings\Security Settings\Public Key Policies**.
  - e. In the details pane, right-click **Automatic Certificate Request Settings**, point to **New**, and then click **Automatic Certificate Request**.
  - f. In the Automatic Certificate Request Setup Wizard, click **Next**.
  - g. On the **Certificate Template** page, click **Computer**, click **Next**, and then click **Finish**.
  - h. Close the Group Policy Management Editor and close the Group Policy Management console.

### ► Task 3: Configure the internal resources for DirectAccess

1. To request a certificate for LON-SVR1 by performing the following steps:
  - a. On LON-SVR1, move the mouse to the lower-right corner of the screen, select **Search**, type **cmd**, and then press Enter.
  - b. At the command prompt, type the following command and then press Enter.

**gpupdate /force**
  - c. At the command prompt, type the following command and then press Enter.

**mmc**
  - d. Click **File** and then click **Add/Remove Snap-in**.
  - e. Click **Certificates**, click **Add**, select **Computer account**, click **Next**, select **Local computer**, click **Finish**, and then click **OK**.
  - f. In the console tree of the Certificates snap-in, navigate to **Certificates (Local Computer) \Personal\Certificates**.
  - g. Right-click **Certificates**, point to **All Tasks**, and then click **Request New Certificate**.
  - h. Click **Next** twice.
  - i. On the **Request Certificates** page, click **Adatum Web Server Certificate**, and then click **More information is required to enroll for this certificate**.
  - j. On the **Subject** tab of the **Certificate Properties** dialog box, under **Subject name**, under **Type**, select **Common name**.
  - k. In the **Value** box, type **nls.adatum.com**, and then click **Add**.
  - l. Click **OK**, click **Enroll**, and then click **Finish**.
  - m. In the details pane of the Certificates snap-in, verify that a new certificate with the name **nls.adatum.com** was enrolled with **Intended Purposes of Server Authentication**.
  - n. Close the console window. When you are prompted to save settings, click **No**.
2. To change the HTTPS bindings, perform the following steps:
  - a. In Server Manager, click **Tools**, and then click **Internet Information Services (IIS) Manager**. At the Internet Information Services (IIS) Manager message box, click **No**.
  - b. In the console tree of Internet Information Services (IIS) Manager, navigate to **LON-SVR1/Sites**, and then click **Default Web site**.
  - c. In the Actions pane, click **Bindings**. Click **Add**.
  - d. In the **Add Site Bindings** dialog box, click **https**, in the **SSL Certificate**, click the certificate with the name **nls.adatum.com**, click **OK**, and then click **Close**.
  - e. Close the Internet Information Services (IIS) Manager console.

**► Task 4: Configure DirectAccess server**

1. Obtain required certificates for LON-SVR2 by performing the following steps:

- a. Switch to **LON-SVR2**.
- b. Open a command prompt and type the following command, and then press Enter:

```
gpupdate /force
```

- c. Move the mouse to the lower-right corner, select **Search**, type **mmc.exe**, and then press Enter.
  - d. Click **File** and then click **Add/Remove Snap-in**.
  - e. Click **Certificates**, click **Add**, click **Computer account**, click **Next**, select **Local computer**, click **Finish**, and then click **OK**.
  - f. In the console tree of the Certificates snap-in, navigate to **Certificates (Local Computer) \Personal\Certificates**.
  - g. Right-click **Certificates**, point to **All Tasks**, and then click **Request New Certificate**.
  - h. Click **Next** twice.
  - i. On the **Request Certificates** page, click **Adatum Web Server Certificate**, and then click **More information is required to enroll for this certificate**.
  - j. On the **Subject** tab of the **Certificate Properties** dialog box, under **Subject name**, under **Type**, select **Common name**.
  - k. In the **Value** box, type **131.107.0.2**, and then click **Add**.
  - l. Click **OK**, click **Enroll**, and then click **Finish**.
  - m. In the details pane of the Certificates snap-in, verify that a new certificate with the name **131.107.0.2** was issued with **Intended Purposes** of **Server Authentication**.
  - n. Right-click the certificate and then click **Properties**.
  - o. In the **Friendly Name** box, type **IP-HTTPS Certificate**, and then click **OK**.
  - p. Close the console window. If you are prompted to save settings, click **No**.
2. Create CRL distribution point on LON-SVR2 by performing the following steps:
    - a. Switch to Server Manager.
    - b. Click **Tools**, and then click **Internet Information Services (IIS) Manager**.
    - c. If the Internet Information Service Manager message box appears, click **No**.
    - d. In the console tree, browse to **LON-SVR2\Sites\Default Web Site**, right-click **Default Web Site**, and then click **Add Virtual Directory**.
    - e. In the **Add Virtual Directory** dialog box, in the **Alias box**, type **CRLD**. Next to **Physical path**, click the ellipsis button.
    - f. In the **Browse for Folder** dialog box, click **Local Disk (C:)**, and then click **Make New Folder**.
    - g. Type **CRLDist** and then press Enter. In the **Browse for Folder** dialog box, click **OK**.
    - h. In the **Add Virtual Directory** dialog box, click **OK**.
    - i. In the middle pane of the console, double-click **Directory Browsing**, and in the Actions pane, click **Enable**.
    - j. In the console tree, click the **CRLD** folder.

- k. In the middle pane of the console, double-click the **Configuration Editor** icon.
- l. Click the down-arrow of the **Section** drop-down list, and navigate to **system.webServer\security\requestFiltering**.
- m. In the middle pane of the console, double-click the **allowDoubleEscaping** entry to change the value from **False** to **True**.
- n. In the details pane, click **Apply**.
- o. Close Internet Information Services (IIS) Manager.

**Question:** Why do you make the CRL available on the Edge server?

**Answer:** You make the CRL available on the Edge Server so that the Internet DirectAccess clients can access the CRL.

3. Share and secure the CRL distribution point by performing the following steps:



**Note:** You perform this step to assign permissions to the CRL distribution point.

- a. On the taskbar, click **Windows Explorer**.
- b. Double-click **Local Disk (C:)**.
- c. In the details pane of Windows Explorer, right-click the **CRLDist** folder, and then click **Properties**.
- d. In the **CRLDist Properties** dialog box, click the **Sharing** tab, and then click **Advanced Sharing**.
- e. In the **Advanced Sharing** dialog box, select **Share this folder**.
- f. In the **Share name** box, add a dollar sign (\$) to the end so that the share name is **CRLDist\$**.
- g. In the **Advanced Sharing** dialog box, click **Permissions**.
- h. In the **Permissions for CRLDist\$** dialog box, click **Add**.
- i. In the **Select Users, Computers, Service Accounts, or Groups** dialog box, click **Object Types**.
- j. In the **Object Types** dialog box, select **Computers**, and then click **OK**.
- k. In the **Select Users, Computers, Service Accounts, or Groups** dialog box, in the **Enter the object names to select** box, type **LON-DC1**, and then click **Check Names**. Click **OK**.
- l. In the **Permissions for CRLDist\$** dialog box, in the **Group or user names** list, select **LON-DC1 (ADATUM\NYC-DC1\$)**. In the **Permissions for LON-DC1** area, under **Full control**, select **Allow**. Click **OK**.
- m. In the **Advanced Sharing** dialog box, click **OK**.
- n. In the **CRLDist Properties** dialog box, click the **Security** tab.
- o. On the **Security** tab, click **Edit**.
- p. In the **Permissions for CRLDist** dialog box, click **Add**.
- q. In the **Select Users, Computers, Service Accounts, or Groups** dialog box, click **Object Types**.
- r. In the **Object Types** dialog box, select **Computers**. Click **OK**.
- s. In the **Select Users, Computers, Service Accounts, or Groups** dialog box, in the **Enter the object names to select** box, type **LON-DC1**, click **Check Names**, and then click **OK**.

- t. In the **Permissions for CRLDist** dialog box, in the **Group or user names** list, select **LON-DC1 (ADATUM\LON-DC1\$)**. In the **Permissions for LON-DC1** area, under **Full control**, select **Allow**, and then click **OK**.
  - u. In the **CRLDist Properties** dialog box, click **Close**.
  - v. Close the Windows Explorer window.
4. Publish the CRL to LON-SVR2 by performing the following steps:



**Note:** This step makes the CRL available on the edge server for Internet-based DirectAccess clients.

- a. Switch to **LON-DC1**.
  - b. In Server Manager, click **Tools**, and then click **Certification Authority**.
  - c. In the console tree, expand **ADATUM-LON-DC1-CA**, right-click **Revoked Certificates**, point to **All Tasks**, and then click **Publish**.
  - d. In the **Publish CRL** dialog box, click **New CRL**, and then click **OK**.
  - e. On the taskbar, click **Windows Explorer**, type **\\LON-SVR2\CRLDist\$**, and then press Enter.
  - f. In the Windows Explorer window, notice the **Adatum-LON-DC1-CA** files.
  - g. Close the Windows Explorer window.
5. Complete DirectAccess setup wizard on LON-SVR2 by performing the following steps:



**Note:** This step configures LON-SVR2 as a DirectAccess server.

- a. On LON-SVR2, in Server Manager, on the **Tools** menu, click **Remote Access Management**.
- b. In the Remote Access Management console, click **Configuration**.
- c. On the **Enable DirectAccess Wizard**, click **Next**.
- d. Under **Select Groups**, in the details pane, click **Add**.
- e. In the **Select Group** dialog box, type **DA\_Clients**, click **OK**, and then click **Next**.
- f. In the **Network Topology**, verify that **Edge** is selected, and verify that **131.107.0.2** is the public name used by clients to connect to the Remote Access server. Click **Next**.
- g. On **Infrastructure Server Setup** page, click **Next**.
- h. On **Configure Remote Access** page, click **Next**.
- i. In **Summary**, click **Finish**, to apply DirectAccess Settings.
- j. When the configuration is complete, click **Close**.



**Note:** Because the server you already configured is a VPN server, you can only use getting started wizard which generate self-signed certificate for DirectAccess communication. Next steps will modify default DirectAccess settings to include already deployed certificates from the internal Certification Authority

- k. In the Remote Access Management console, under **Step 2**, click **Edit**.



- l. On the **Network Topology** page, verify that **Edge** is selected, and then type **131.107.0.2**
- m. Click **Next**.
- n. On the **Network Adapters** page, verify that **CN=131.107.0.2** is used as a certificate to authenticate IP-HTTPS connections, and then click **Next**.
- o. On the **Authentication** page, select **Use computer certificates**, click **Browse**, select **Adatum LON-DC1 CA**, click **OK**, and then **Next**.
- p. On the VPN Configuration page, click **Finish**.
- q. In the Remote Access Setup pane, under **Step 3**, click **Edit**.
- r. On the **Network Location Server** page, select the **The network location server is deployed on a remote web server (recommended)** and in the URL of the NLS, type **https://nls.adatum.com**, and then click **Validate**.
- s. Ensure that URL is validated.
- t. Click **Next**, and then on the **DNS** page, examine the values, and then click **Next**.
- u. In the **DNS Suffix Search List**, select **Next**.
- v. On the **Management** page, click **Finish**.
- w. Under Step 4, click **Edit**. On the **DirectAccess Application Server Setup** page, click **Finish**.
- x. Click **Finish** to apply the changes.
- y. In **Remote Access Review**, click **Apply**.
- z. Under **Applying Remote Access Setup Wizard Settings**, click **Close**.
6. Update Group Policy settings on **LON-SVR2** by performing the following steps:
  - a. Move the mouse pointer on the lower-right corner and on the menu bar, click **Search**, type **cmd**, and then press Enter.
  - b. At the command prompt, type the following commands and then press Enter.

```
gpupdate /force
Ipconfig
```



**Note:** Verify that **LON-SVR2** has an IPv6 address for **Tunnel adapter IPHTTPSInterface** starting with **2002**.

**Results:** After completing this exercise, you will have configured the DirectAccess infrastructure.

## Exercise 2: Configuring the DirectAccess Clients

### ► Task 1: Configure Group Policy to configure client settings for DirectAccess

1. Switch to **LON-SVR3**.
2. Restart **LON-SVR3** and then log back on as **Adatum\Administrator** with the password of **Pa\$\$w0rd**. This is to ensure that the LON-SVR3 computer connects to the domain as a member of the DA\_Clients security group.
3. Move the mouse pointer to the lower-right corner, select **Search** on the right menu, and then type **cmd** to open the Command Prompt window.

4. At the command prompt, type the following command and then press Enter:

```
gpupdate /force
```

5. At the command prompt, type the following command, and then press Enter:

```
gpresult /R
```

6. Verify that **DirectAccess Client Settings GPO** is displayed in the list of the Applied Policy objects for the Computer Settings.



**Note:** If the policy is not being applied, run the **gpupdate /force** command again. If the policy is still not being applied, restart the computer. After the computer restarts, log on as **Adatum\Administrator** and run the **Gpresult -R** command again.

### ► Task 2: Verify client computer certificate distribution

1. On LON-SVR3, move the mouse pointer to the lower-right corner, select **Search** on the right menu, type **mmc.exe**, and then press Enter
2. Click **File** and then click **Add/Remove Snap-in**.
3. Click **Certificates**, click **Add**, select **Computer account**, click **Next**, select **Local computer**, click **Finish**, and then click **OK**.
4. In the console tree of the Certificates snap-in, navigate to **Certificates (Local Computer) \Personal\Certificates**.
5. In the details pane, verify that a certificate with the name **Lon-SVR3.adatum.com** is present with **Intended Purposes of Client Authentication and Server Authentication**.
6. Close the console window. When you are prompted to save settings, click **No**.

**Question:** Why did you install a certificate on the client computer?

**Answer:** Without a certificate, the client cannot identify and authenticate itself to the DirectAccess server.

### ► Task 3: Verify IP address configuration

1. On LON-SVR3, switch to the **Start** screen and click the **Internet Explorer** tile.
2. In the Address bar, type **http://lon-svr1.adatum.com/** and then press Enter. The default IIS 8 web page for LON-SVR1 appears.
3. In the Address bar, type **https://nls.adatum.com/** and then press Enter. The default IIS 8 web page for LON-SVR1 appears.
4. Leave the Internet Explorer window open.
5. On the taskbar, click Windows Explorer, type **\\Lon-SVR1\Files**, and then press Enter. A folder window with the contents of the **Files** shared folder appears.
6. Close all open windows.

**Results:** After completing this exercise, you will have configured the DirectAccess clients.

## Exercise 3: Verifying the DirectAccess Configuration

### ► Task 1: Move the client computer to the Internet virtual network



**Note:** To verify the DirectAccess functionality, you must move the client computer to the Internet.

1. Switch to **LON-SVR3**.
2. On LON-SVR3, move the mouse pointer to the lower-right end of the screen, click **Settings**, select **Control Panel**, and then click **Network and Internet**.
3. Click **Network and Sharing Center**.
4. Click **Change Adapter Settings**.
5. Right-click **Local Area Connection** and then click **Properties**.
6. In the **Local Area Connection Properties** dialog box, double-click **Internet Protocol Version 4 (TCP/IPv4)**.
7. In the **Internet Protocol Version 4 (TCP/IPv4) Properties** dialog box, click **Use the following IP address**. Fill in the following information, and then click **OK**.
  - IP address: **131.107.0.10**
  - Subnet mask: **255.255.0.0**
  - Default gateway: **131.107.0.2**
8. In the **Local Area Connection Properties** dialog box, click **OK**.
9. In the Network Connections window, right-click **Local Area Connection**, and then click **Disable**.
10. In the Network Connections window, right-click **Local Area Connection**, and then click **Enable**.
11. In Hyper-V Manager, right-click **20417A-LON-SVR3** and then click **Settings**. Change the Legacy Network Adapter to be on the **Private Network 2** network. Click **OK**.

### ► Task 2: Verify connectivity to the DirectAccess server

1. On LON-SVR3, move the mouse pointer to the lower-right corner, select **Search** on the right menu, and then type **cmd** and then press Enter to open the command prompt.
2. At the command prompt, type the following command, and then press Enter:

```
ipconfig
```

3. Notice the IP address that start with **2002**. This is an IP-HTTPS address.
4. At the command prompt, type the following command, and then press Enter:

```
Netsh name show effectivepolicy
```

5. At the command prompt, type the following command, and then press Enter:

```
powershell
```

- At the Windows PowerShell command prompt, type the following command, and then press Enter:

```
Get-DAClientExperienceConfiguration
```



**Note:** Notice the DirectAccess client settings.

► **Task 3: Verify connectivity to the internal network resources**

- Switch to the Start screen and then click the **Internet Explorer** tile.
- In the Address bar, type **http://lon-svr1.adatum.com** and then press Enter. The default IIS 8 web page for LON-SVR1 appears.
- Leave the Internet Explorer window open.
- On the taskbar, click **Windows Explorer**, type **\\LON-SVR1\Files**, and then press Enter. A folder window with the contents of the Files shared folder appears.
- Switch to the Command Prompt window.
- At the command prompt, type the following command and then press Enter:

```
ping lon-dc1.adatum.com
```

Verify that you are receiving replies from lon-dc1.adatum.com.

- At the command prompt, type the following command, and then press Enter:

```
gpupdate /force
```

- Close all open windows.
- Switch to **LON-SVR2**.
- On the Start screen, click **Remote Access Management**.
- In the Console pane, click **Remote Client Status**.



**Note:** Notice that LON-SVR3 is connected via IPHttps. In the Connection Details pane, in the bottom-right of the screen, note the use of Kerberos for the Machine and the User.

- Close all open windows.

**Results:** After completing this exercise, you will have verified the DirectAccess configuration.

► **To prepare for the next module**

When you finish the lab, revert the virtual machines to their initial state. To do this, perform the following steps:

- On the host computer, start **Hyper-V Manager**.
- In the **Virtual Machines** list, right-click **20410A-LON-DC1**, and then click **Revert**.
- In the **Revert Virtual Machine** dialog box, click **Revert**.
- Repeat steps 2 and 3 for **20410A-LON-SVR1**, **20410A-LON-SVR2**, and **20410A-LON-SVR3**.

## Module 7: Implementing Failover Clustering

# Lab: Implementing Failover Clustering

### Exercise 1: Configuring a Failover Cluster

#### ► Task 1: Connect clients to the iSCSI targets

1. On LON-SVR3, in Server Manager, click **Tools**, and then click **iSCSI Initiator**.
2. In the **Microsoft iSCSI** dialog box, click **Yes**.
3. Click the **Discovery** tab.
4. Click **Discover Portal**.
5. In the **IP address or DNS name** box, type **172.16.0.21**, and then click **OK**.
6. Click the **Targets** tab.
7. Click **Refresh**.
8. In the **Targets** list, select **iqn.1991-05.com.microsoft:lon-svr1-target1-target**, and then click **Connect**.
9. Select **Add this connection to the list of Favorite Targets**, and then click **OK** two times.
10. On LON-SVR4, in Server Manager, click **Tools**, and then click **iSCSI Initiator**.
11. In the **Microsoft iSCSI** dialog box, click **Yes**.
12. Click the **Discovery** tab.
13. Click **Discover Portal**.
14. In the **IP address or DNS name** box, type **172.16.0.21**, and then click **OK**.
15. Click the **Targets** tab.
16. Click **Refresh**.
17. In the **Targets** list, select **iqn.1991-05.com.microsoft:lon-svr1-target1-target**, and then click **Connect**.
18. Select **Add this connection to the list of Favorite Targets**, and then click **OK** two times.
19. On LON-SVR3, in Server Manager, click **Tools**, and then click **Computer Management**.
20. Expand **Storage**, and then click **Disk Management**.
21. Right-click **Disk 1**, and then click **Online**.
22. Right-click **Disk 1**, and then click **Initialize disk**. In the **Initialize Disk** dialog box, click **OK**.
23. Right-click the unallocated space next to **Disk 1**, and then click **New Simple Volume**.
24. On the **Welcome** page, click **Next**.
25. On the **Specify Volume Size** page, click **Next**.
26. On the **Assign Drive Letter or Path** page, click **Next**.
27. On the **Format Partition** page, in the **Volume Label** box, type **Data**. Select the **Perform a quick format** check box, and then click **Next**.
28. Click **Finish**. (Note: If the Microsoft Windows window pops up with prompt to format the disk, click **Cancel**.)

29. Repeat steps 22 through 28 for **Disk 2 and Disk 3**. (Note: Use Data2 and Data3 for Volume Labels).
30. Close the Computer Management window.
31. On LON-SVR4, in Server Manager, click **Tools**, and then click **Computer Management**.
32. Expand **Storage**, and then click **Disk Management**.
33. Right-click **Disk Management**, and then click **Refresh**.
34. Right-click **Disk 1**, and then click **Online**.
35. Right-click **Disk 2**, and then click **Online**.
36. Right-click **Disk 3**, and then click **Online**.
37. Close the Computer Management window.

► **Task 2: Install the Failover Clustering feature**

1. On LON-SVR3, if it is not opened, click the Server Manager icon to open Server Manager.
2. Click **Add roles and features**.
3. On the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, click **Next**.
5. On the **Select destination server** page, make sure that **Select server from the server pool** is selected, and then click **Next**.
6. On the **Select server roles** page, click **Next**.
7. On the **Select features** page, in the **Features** list, click **Failover Clustering**. In the Add features that are required for Failover Clustering? window, click **Add Features**. Click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. When installation is complete (you get the message Installation succeeded on LON-SVRx), click **Close**.
10. Repeat steps 1 through 9 on **LON-SVR4**.

► **Task 3: Validate the servers for Failover Clustering**

1. On LON-SVR3, in the Server Manager, click **Tools**, and then click **Failover Cluster Manager**.
2. In the Actions pane of the Failover Cluster Manager, click **Validate Configuration**.
3. In the Validate a Configuration Wizard, click **Next**.
4. In the **Enter Name** box, type **LON-SVR3**, and then click **Add**.
5. In the **Enter Name** box, type **LON-SVR4**.
6. Click **Add**, and then click **Next**.
7. Verify that **Run all tests (recommended)** is selected, and then click **Next**.
8. On the **Confirmation** page, click **Next**.
9. Wait for the validation tests to finish (it might take up to 5 minutes), and then on the **Summary** page, click **View Report**.
10. Verify that all tests completed without errors. Some warnings are expected.
11. Close Internet Explorer.
12. On the **Summary** page, remove the check mark next to **Create the cluster now using the validated nodes**, click **Finish**.

#### ► Task 4: Create the Failover Cluster

1. On LON-SVR3, in Failover Cluster Manager, in the center pane, under **Management**, click **Create Cluster**.
2. In the Create Cluster Wizard on the **Before You Begin** page, read the information.
3. Click **Next**, in the **Enter server name** box, type **LON-SVR3**, and then click **Add**. Type **LON-SVR4**, and then click **Add**.
4. Verify the entries, and then click **Next**.
5. In **Access Point for Administering the Cluster**, in the **Cluster Name** box, type **Cluster1**.
6. Under **Address**, type **172.16.0.125**, and then click **Next**.
7. In the **Confirmation** dialog box, verify the information, and then click **Next**.
8. On the **Summary** page, click **Finish** to return to the Failover Cluster Manager.

**Results:** After this exercise, you will have installed and configured the Failover Clustering feature.

### Exercise 2: Deploying and Configuring a Highly-Available File Server

#### ► Task 1: Add the File Server application to the failover cluster

1. On LON-SVR3, in Server Manager, click **Dashboard** and then click **Add roles and features**.
2. On the **before your begin** page click **Next**.
3. On the **Select installation type** page click **Next**.
4. On the **Select destination server** page click **Next**.
5. On the **Select server roles** page, expand **File and Storage Services (Installed)**, expand **File and iSCSI services** and select **File Server**.
6. Click **Next** two times.
7. On the **Confirmation** page, click **Install**.
8. When **installation succeeded** message appears click **Close**.
9. Repeat steps 1-8 on **LON-SVR4**.
10. On LON-SVR3, in the Failover Cluster Manager expand **Cluster1.adatum.com**.
11. Expand **Storage**, and click **Disks**.
12. Make sure that three disks are present and online (with names Cluster Disk 1, Cluster Disk 2 and Cluster Disk 3).
13. Right-click **Roles**, and then select **Configure Role**.
14. On the **Before You Begin** page, click **Next**.
15. On the **Select Role** page, select **File Server**, and then click **Next**.
16. On the **File Server Type** page, click **File Server for general use**, and then click **Next**.
17. On the **Client Access Point** page, in the **Client Access Name** box, type **AdatumFS**, and in the **Address** box, type **172.16.0.130**, and then click **Next**.
18. On the **Select Storage** page, click **Cluster Disk 2**, and then click **Next**.

19. On the **Confirmation** page, click **Next**.
20. On the **Summary** page, click **Finish**.

► **Task 2: Add a shared folder to a highly-available file server**

1. On LON-SVR4, in the Server Manager console, click Tools and open **Failover Cluster Manager**.
2. Expand **Cluster1.Adatum.com**, and then click **Roles**.
3. Right-click **AdatumFS**, and then select **Add File Share**.
4. In the New Share Wizard, on the **Select the profile for this share** page, click **SMB Share – Quick**, and then click **Next**.
5. On the **Select the server and the path for this share** page, click **Next**.
6. On the **Specify share name** page, in the **Share name** box, type **Docs**, and then click **Next**.
7. On the **Configure share settings** page, review available options, and then click **Next**.
8. On the **Specify permissions to control access** page, click **Next**.
9. On the **Confirm selections** page, click **Create**.
10. On the View results page click **Close**.

► **Task 3: Configure failover and failback settings**

1. On LON-SVR4, in the Failover Cluster Manager, click **Roles**, right-click **AdatumFS**, and then click **Properties**.
2. Click the **Failover** tab and then click **Allow failback**.
3. Click **Failback between**, and set values to **4 and 5 hours**.
4. Click the **General** tab.
5. Select both **LON-SVR3** and **LON-SVR4** as preferred owners.
6. Move **LON-SVR4** up.
7. Click **OK**.

**Results:** After this exercise, you will have configured a highly-available file server.

### Exercise 3: Validate the Deployment of the Highly-Available File Server

► **Task 1: Validate the highly-available file server deployment**

1. On LON-DC1, open Windows Explorer, and in the Address bar, type **\\AdatumFS\**, and then press Enter.
2. Verify that you can access the location and that you can open the **Docs** folder. Create a test text document inside this folder.
3. On LON-SVR3, open the Failover Cluster Manager.
4. Expand **Cluster1.adatum.com**, and then click **Roles**. Note the current owner of AdatumFS. (Note: You can view the owner in the Owner node column. It will be either LON-SVR3 or LON-SVR4).
5. Right-click **AdatumFS**, and then click **Move**, and then click **Select Node**.
6. In the **Move Clustered Role** dialog box, click **OK**.



7. Verify that **AdatumFS** has moved to a new owner.
8. Switch to the **LON-DC1** computer and verify that you can still access the **\\AdatumFS\** location.

► **Task 2: Validate the failover and quorum configuration for the File Server role**

1. On LON-SVR3, in the Failover Cluster Manager, click **Roles**.
2. Verify the current owner for the AdatumFS role. (Note: You can view the owner in the Owner node column. It will be either LON-SVR3 or LON-SVR4).
3. Expand **Nodes**, and then select the node that is the current owner of the AdatumFS role.
4. Right-click the node, select **More Actions**, and then click **Stop Cluster Service**. Click **Yes** when prompted.
5. Verify that **AdatumFS** has moved to another node. To do this, click the other node and verify that AdatumFS is running.
6. Switch to the **LON-DC1** computer and verify that you can still access the **\\AdatumFS\** location.
7. Switch to the **LON-SVR3** computer, on the Failover Cluster Manager, and right-click the stopped node, select **More Actions**, and then click **Start Cluster Service**.
8. Expand **Storage** and then click **Disks**. In the center pane, right-click the disk that is assigned to **Disk Witness in Quorum** (Note: you can view this in the **Assigned to** column.)
9. Click **Take Offline**, and then click **Yes**.
10. Switch to **LON-DC1** and verify that you can still access the **\\AdatumFS\** location. By doing this, you verified that the cluster is still running even if the witness disk is offline.
11. Switch to the **LON-SVR3** computer and in Failover Cluster Manager, expand **Storage**, click **Disks**, right-click the disk that is in **Offline** status, and then click **Bring Online**.

**Results:** After this exercise, you will have tested the failover scenarios.

## Exercise 4: Configuring Cluster-Aware Updating on the Failover Cluster

► **Task 1: Configure Cluster-Aware Updating**

1. On LON-DC1, in Server Manager, click **Add roles and features**.
2. In the Add Roles and Features Wizard, on the **Before you begin** page, click **Next**.
3. On the **Select installation type** page, click **Next**.
4. On the **Select destination server** page, make sure that **Select server from the server pool** is selected, and then click **Next**.
5. On the **Select server roles** page, click **Next**.
6. On the **Select features** page, in the list of features, click **Failover Clustering**. In **Add features that are required for Failover Clustering?** dialog box, click **Add Features**. Click **Next**.
7. On the **Confirm installation selections** page, click **Install**.
8. When installation is complete, click **Close**.
9. Switch to **LON-SVR3**. Open **Server Manager**, click **Tools** and then click **Windows Firewall with Advanced Security**.
10. In Windows Firewall with Advanced Security window, click **Inbound Rules**.

11. In the rules list, find the rule **Inbound Rule for Remote Shutdown (RPC-EP-In)**. Right click the rule and select **Enable Rule**.
12. In the rules list, find the rule **Inbound Rule for Remote Shutdown (TCP-In)**. Right click the rule and select **Enable Rule**.
13. Close Windows Firewall with Advanced Security window.
14. Switch to **LON-SVR4** and repeat steps 9 to 13.
15. On LON-DC1, in the Server Manager dashboard, click **Tools**, and then click **Cluster-Aware Updating**.
16. In the Cluster-Aware Updating window, in the **Connect to a failover cluster** drop-down list, select **Cluster1**. Click **Connect**.
17. In the Cluster Actions pane, click **Preview updates for this cluster**.
18. In the Cluster1-Preview Updates window, click **Generate Update Preview List**. After several minutes, updates will be shown in the list. Review updates and then click **Close**.



**Note:** An Internet connection is required for this step to complete successfully. Make sure that MSL-TMG1 server is up and running and that you can access Internet from LON-DC1.

#### ► Task 2: Update the failover cluster and configure self-updating

1. On LON-DC1, in the Cluster-Aware Updating console, click **Apply updates to this cluster**.
2. On the **Getting Started** page, click **Next**.
3. On the **Advanced options** page, review the options for updating, and then click **Next**.
4. On the **Additional Update Options** page, click **Next**.
5. On the **Confirmation** page, click **Update**, and then click **Close**.
6. In the Cluster nodes pane, you can review the progress of updating. (Note: Remember that one node of the cluster is in Waiting state and the other node is restarting after it is updated).
7. Wait until the process is finished (Note: This may require a restart of both the nodes.). Process is finished when both nodes have Succeeded in Last Run status column.
8. Log on to **LON-SVR3** with the username as **Adatum\Administrator** and password as **Pa\$\$w0rd**.
9. On LON-SVR3, in the Server Manager, click **Tools**, and then click **Cluster-Aware Updating**.
10. In the **Cluster-Aware Updating** dialog box, in the **Connect to a failover cluster** drop-down list, select **Cluster1**. Click **Connect**.
11. Click the **Configure cluster self-updating options** in the Cluster Actions pane.
12. On the **Getting Started** page, click **Next**.
13. On the **Add CAU Clustered Role with Self-Updating Enabled** page, click **Add the CAU clustered role, with self-updating mode enabled, to this cluster**, and then click **Next**.
14. On the **Specify self-updating schedule** page, click **Weekly**, in the **Time of day** box, select **4:00 AM**, and then in the **Day of the week** box, select **Sunday**. Click **Next**.
15. On the **Advanced Options** page, click **Next**.
16. On the **Additional Update Options** page, click **Next**.

17. On the **Confirmation** page, click **Apply**.
18. After the clustered role is added successfully, click **Close**.

**Results:** After this exercise, you will have configured Cluster-Aware Updating.

► **To prepare for next module**

When you finish the lab, revert the virtual machines to their initial state. To do this, perform the following steps:

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20417A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20417A-LON-SVR1**, **20417A-LON-SVR3**, **MSL-TMG1**, and **20417A-LON-SVR4**.

**MCT USE ONLY. STUDENT USE PROHIBITED**

## Module 8: Implementing Hyper-V

# Lab: Implementing Server Virtualization with Hyper-V

### Exercise 1: Install the Hyper-V Server Role

#### ► Task 1: Configure network settings on LON-HOST1 and LON-HOST2

1. Restart the classroom computer, and in the **Windows Boot Manager**, select either **20417A-LON-HOST1** or **20417A-LON-HOST2**.

If you start LON-HOST1, your partner must start LON-HOST2.

2. Log onto the server with the **Adatum\Administrator** account and the password **Pa\$\$w0rd**.
3. In Server Manager, click **Local Server**.
4. In the Properties pane, click the **IPv4 address assigned by DHCP** link.
5. In the **Network Connections** dialog box, right-click the network object, and then click **Properties**.
6. In the **Properties** dialog box, click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
7. On the **General** tab, click **Use the following IP address**, and then configure the following:
  - LON-HOST1: **172.16.0.31**
  - LON-HOST2: **172.16.0.32**
  - Subnet mask: **255.255.0.0**
  - Default gateway: **172.16.0.1**
8. On the **General** tab, click **Use the following DNS server addresses**, and then configure the following:
  - Preferred DNS server: **172.16.0.10**
9. Click **OK** to close the Properties dialog box.
10. Click **OK** on the **Microsoft TCP/IP** dialog box.
11. Click **Close**.
12. Close the Network Connections dialog box.

#### ► Task 2: Install the Hyper-V server role

1. In the Server Manager console, on the **Manage** menu, click **Add Roles and Features**.
2. On the **Before you begin** page of the Add Roles and Features Wizard, click **Next**.
3. On the **Select installation type** page, select **Role-based or feature-based installation**, and then click **Next**.
4. On the **Select destination server** page, ensure that **LON-HOST1.Adatum.com** or **LON-HOST2.Adatum.com** is selected, and then click **Next**.
5. On the **Server Roles** page, select **Hyper-V**.
6. In the **Add Roles and Features Wizard** dialog box, click **Add Features**.
7. On the **Select Server Roles** page of the Add Roles and Features Wizard, click **Next**.

8. On the **Select features** page, click **Next**.
9. On the **Hyper-V** page, click **Next**.
10. On the **Create Virtual Switches** page, verify that no selections have been made, and then click **Next**.
11. On the **Virtual Machine Migration** page, click **Next**.
12. On the **Default Stores** page, review the location of **Default Stores**, and then click **Next**.
13. On the **Confirm Installation Selections** page, select **Restart the destination server automatically if required**.
14. In the **Add Roles and Features Wizard** dialog box, review the message about automatic restarts, and then click **Yes**.
15. On the **Confirm Installation Selections** page, click **Install**.
16. After a few minutes, the server will automatically restart. Ensure that you restart the machine by using the **Boot** menu, and then selecting **20417-LON-HOST1** or **20417-LON-HOST2**. The computer will restart several times.

► **Task 3: Complete Hyper-V role installation and verify settings**

1. Log on to **LON-HOST1** or **LON-HOST2** by using the username **Adatum\Administrator** and the password **Pa\$\$w0rd**.
2. When the installation of the Hyper-V tools complete, click **Close** to close the Add Roles and Features Wizard.
3. Click the **Tools** menu, and then click **Hyper-V Manager**.
4. In the Hyper-V Manager console, click the Hyper-V host server name (**LON-HOST1** or **LON-HOST2**).
5. In the Actions pane, click **Hyper-V Settings**.
6. In the **Hyper-V Settings** dialog box, click the **Keyboard** item. Verify that the **Keyboard** is set to use the **Use on the virtual machine** option.
7. In the **Hyper-V Settings** dialog box, click the **Virtual Hard Disks** item. Verify the location of the default folder is configured to use the **Virtual Hard Disk** folder, and then click **OK**.

**Question:** What additional features are required to support the Hyper-V role?

**Answer:** No additional features are required to support the Hyper-V role.

**Results:** After completing this exercise, you will have deployed the Hyper-V role to a physical server.

## Exercise 2: Configuring Virtual Networking

► **Task 1: Configure the external network**

1. In Hyper-V Manager, on the Actions pane, click **Virtual Switch Manager**.
2. In the **Virtual Switch Manager** dialog box, select **New virtual network switch**. Ensure that **External** is selected, and then click **Create Virtual Switch**.

3. In the **Virtual Switch Properties** area of the **Virtual Switch Manager** dialog box, specify the following information, and then click **OK**:
  - o Name: **Corporate Network**
  - o External Network: Mapped to the host computer's physical network adapter. Will vary depending on host computer
4. In the **Apply Networking Changes** dialog box, review the warning, and then click **Yes**.

► **Task 2: Create a private network**

1. In Hyper-V Manager, on the Actions pane, click **Virtual Switch Manager**.
2. Under **Virtual Switches**, select **New virtual network switch**.
3. Under **Create virtual switch**, select **Private**, and then click **Create Virtual Switch**.
4. In the **Virtual Switch Properties** section, configure the following settings, and then click **OK**:
  - o Name: **Private Network**
  - o Connection type: **Private network**

► **Task 3: Create an internal network**

1. In Hyper-V Manager, on the Actions pane, click **Virtual Switch Manager**.
2. Under **Virtual Switches**, select **New virtual network switch**.
3. Under **Create virtual switch**, select **Internal**, and then click **Create Virtual Switch**.
4. In the **Virtual Switch Properties** section, configure the following settings, and then click **OK**:
  - o Name: **Internal Network**
  - o Connection type: **Internal network**

**Results:** After completing this exercise, you will have configured virtual switch options on a physically deployed Windows Server 2012 server that is running the Hyper-V role.

## Exercise 3: Creating and Configuring a Virtual Machine

► **Task 1: Configure virtual machine storage**

1. On the taskbar, click **Windows Explorer**.
2. Click **Computer**, and then browse to the following location:  
**E:\Program Files\Microsoft Learning\Base**. (Note: The drive letter may depend upon the number of drives on the physical host machine)
3. Verify that the **Base12A-WS2012-RC.vhd** hard disk image file is present.
4. Click the **Home** tab, and then click the **New Folder** icon twice to create two new folders. Right-click each folder, and then rename each folders to each name listed below:
  - a. **LON-GUEST1**
  - b. **LON-GUEST2**
5. Close Windows Explorer.
6. Switch to the Hyper-V Manager.
7. In the Actions pane, click **New**, and then click **Hard Disk**.

8. On the **Before You Begin** page of the New Virtual Hard Disk Wizard, click **Next**.
9. On the **Choose Disk Format** page, select **VHD**, and then click **Next**.
10. On the **Choose Disk Type** page, select **Differencing**, and then click **Next**.
11. On the **Specify Name and Location** page, specify the following details, and then click **Next**:
  - a. Name: **LON-GUEST1.vhd**
  - b. Location: **E:\Program Files\Microsoft Learning\Base\LON-GUEST1\**
12. On the **Configure Disk** page, type the location: **E:\Program Files\Microsoft Learning\Base\Base12A-WS2012-RC.vhd**, and then click **Finish**.
13. On the taskbar, click the **PowerShell** icon.
14. At the PowerShell prompt, type the following command to import the Hyper-V module, and then press Enter:

```
Import-Module Hyper-V
```

15. At the PowerShell prompt, type the following command to create a new differencing disk to be used with LON-GUEST2, and then press Enter:

```
New-VHD "E:\Program Files\Microsoft Learning\Base\LON-GUEST2\LON-GUEST2.vhd"  
-ParentPath "E:\Program Files\Microsoft Learning\Base\Base12A-WS2012-RC.vhd"
```

16. Close the PowerShell window.
17. In the Actions pane of the Hyper-V Manager console, click **Inspect Disk**.
18. In the **Open** dialog box, browse to **E:\Program Files\Microsoft Learning\Base\LON-GUEST2\**, click **LON-GUEST2.vhd**, and then click **Open**.
19. In the **Virtual Hard Disk Properties** dialog box, verify that **LON-GUEST2.vhd** is configured as a differencing virtual hard disk with **E:\Program Files\Microsoft Learning\Base\Base12A-WS2012-RC.vhd** as a parent, and then click **Close**.

## ► Task 2: Create virtual machines

1. In the Hyper-V Manager, on the Actions pane, click **New** and then click **Virtual Machine**.
2. On the **Before You Begin** page of the New Virtual Machine Wizard, click **Next**.
3. On the **Specify Name and Location** page of the New Virtual Machine Wizard, select **Store the virtual machine in a different location**, enter the following values, and then click **Next**.
  - a. Name: **LON-GUEST1**
  - b. Location: **E:\Program Files\Microsoft Learning\Base\LON-GUEST1\**
4. On the **Assign Memory** page of the New Virtual Machine Wizard, enter a value of **1024 MB**, select the **Use Dynamic Memory for this virtual machine** option, and click **Next**.
5. On the **Configure Networking** page of the New Virtual Machine Wizard, choose **Private Network** and then click **Next**.
6. On the **Connect Virtual Hard Disk** page, choose **Use an existing virtual hard disk**. Click **Browse** and browse to **E:\Program Files\Microsoft Learning\Base\LON-GUEST1\lon-guest1.vhd**. Click **Open** and then click **Finish**.
7. On the **Taskbar**, click the **PowerShell** icon.



8. At the **PowerShell** prompt, enter the following command to import the Hyper-V module:

```
Import-Module Hyper-V
```

9. At the **PowerShell** prompt, enter the following command to create a new virtual machine named **LON-GUEST2**:

```
New-VM -Name LON-GUEST2 -MemoryStartupBytes 1024MB -VHDPATH "E:\Program  
Files\Microsoft Learning\Base\LON-GUEST2\LON-GUEST2.vhd" -SwitchName "Private  
Network"
```

10. Close the PowerShell window.
11. In the Hyper-V Manager console, click **LON-GUEST2**. In the Actions pane, under **LON-GUEST2**, click **Settings**.
12. On the **Settings for LON-GUEST2** dialog box, click **Automatic Start Action**, and then set the **Automatic Start Action** to **Nothing**.
13. On the **Settings for LON-GUEST2** dialog box, click **Automatic Stop Action**, and then set the **Automatic Stop Action** to **Shut down the guest operating system**.
14. Click **OK** to close the Settings for the LON-GUEST2 dialog box.

### ► Task 3: Configure VLANs and network bandwidth settings

1. In the Hyper-V Manager console, on the Actions pane, click **Virtual Switch Manager**.
2. Click **Internal Network**.
3. Select the **Enable virtual LAN identification for management operating system** check box.
4. In the **VLAN ID** box, type **4**, and then click **OK**.
5. Click **LON-GUEST2**, and click **Settings**.
6. Click **Network Adapter**.
7. Change the **Virtual switch** to **Internal Network**, and click **Enable virtual LAN identification**.
8. In the **VLAN identifier** box, type **4**.
9. Expand **Network Adapter**, click **Advanced Features**, enable the following options, and then click **OK**:
  - **Enable DHCP guard**
  - **Enable router advertisement guard**

**Question:** What kind of switch would you create if you added a new physical network adapter to the Hyper-V host and wanted to keep this separate from the existing networks you create during this exercise?

**Answer:** You should create an external switch. External switches map to external network adapters.

### ► Task 4: Import a virtual machine

1. In the Actions pane of the Hyper-V Manager console, click **Import Virtual Machine**.
2. On the **Before You Begin** page of the Import Virtual Machine wizard, click **Next**.

3. On the **Locate Folder** page, perform the following task, and then click **Next**:
  - If you are using LON-HOST1, type the path: **E:\Program Files\Microsoft Learning\20417\Drives\20417A-LON-DC1-B**
  - If you are using LON-HOST2, enter the path: **E:\Program Files\Microsoft Learning\20417\Drives\20417A-LON-SVR1-B**
4. On the **Select Virtual machine** page:
  - If you are using LON-HOST1, select **20417A-LON-DC1-B**.
  - If you are using LON-HOST2, select **20417A-LON-SVR1-B**.
5. On the **Choose Import Type** page, select **Register the virtual machine in-place (use the existing unique ID)**, and then click **Next**.
6. On the **Summary** page, click **Finish**.

► **Task 5: Configure virtual machine dynamic memory**

1. In the Hyper-V Manager console, right-click **LON-GUEST2**, and then click **Settings**.
2. In the **Settings for LON-GUEST2** dialog box, click **Memory**.
3. In the **Memory** page, configure the **Startup RAM** as **1024 MB**.
4. On the **Memory** page, select the **Enable Dynamic Memory** option.
5. Set the following dynamic memory settings:
  - Minimum RAM: **512 MB**
  - Maximum RAM: **2048 MB**
6. Click **OK** to close the Settings for LON-GUEST2 dialog box.

► **Task 6: Configure and test virtual machine snapshots**

1. If you are using **LON-HOST1**, start and connect to **20417A-LON-DC1-B**.
2. Log on to **LON-DC1** as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
3. If you are using **LON-HOST2**, start and connect to **20417A-LON-SVR1-B**.
4. Log on to **LON-SVR1** as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
5. Minimize the Server Manager console.
6. Right-click the desktop of the virtual machine, click **New**, and then click **Folder**. Name the folder **Sydney**.
7. Repeat step 6, and then create a second folder **Melbourne**.
8. Repeat step 6, and then create a third folder **Brisbane**.
9. On the Action menu of the Virtual Machine Connection window, click **Snapshot**.
10. In the **Snapshot Name** dialog box, in the **Name** box, type **Before Change**, and then click **Yes**.
11. Drag the **Sydney** folder to the Recycle Bin.
12. Drag the **Brisbane** folder to the Recycle Bin.
13. Right-click the **Recycle Bin**, and then click **Empty Recycle Bin**.
14. In the **Delete Multiple Items** dialog box, click **Yes**.
15. On the Action menu of the Virtual Machine Connection window, click **Revert**.

16. In the **Revert Virtual Machine** dialog box, click **Revert**.
17. Verify that the following folders are present on the desktop:
  - Sydney
  - Melbourne
  - Brisbane
18. Delete all three folders from the desktop.

**Question:** What state must the virtual machine be in to configure dynamic memory when using Windows Server 2008 R2 as a host? How is this different to Windows Server 2012 as a host?

**Answer:** The virtual machine must be powered off to configure dynamic memory. In Windows Server 2012, you can configure dynamic memory while the virtual machine is powered on.

**Results:** After completing this exercise, you will have deployed two separate virtual machines by using a sysprepped virtual hard-disk file to act as a parent disk for two differencing disks. You also will have imported a specially prepared virtual machine.

► **To prepare for the next module**

When you are finished the lab, leave the virtual machines running, as they are needed for the lab in Module 9.

**MCT USE ONLY. STUDENT USE PROHIBITED**

## Module 9: Implementing Failover Clustering with Hyper-V

### Lab: Implementing Failover Clustering with Hyper-V

#### Exercise 1: Configuring Hyper-V Replicas

##### ► Task 1: Import LON-CORE virtual machine on LON-HOST1

1. Log on to **LON-HOST1** as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On LON-HOST1 open the Hyper-V Manager console.
3. In the Actions pane, click **Import Virtual Machine**.
4. On the **Before You Begin** page in Import Virtual Machine Wizard, click **Next**.
5. On **Locate Folder** page click **Browse**.
6. Browse to folder **E:\Program Files\Microsoft Learning\20417\Drives\20417A-LON-CORE**. Click **Select Folder** and then click **Next**.



**Note:** The drive letter may be different based upon the number of drives on the physical host machine.

7. On **Select Virtual Machine** page, select **20417A-LON-CORE** and then click **Next**.
8. On the **Choose Import Type** page click **Next**.
9. On the **Summary** page click **Finish**.

##### ► Task 2: Configure a replica on both host machines

1. On LON-HOST2, open the Hyper-V Manager console.
2. In Hyper-V Manager, right-click **LON-HOST2** and select **Hyper-V Settings...**
3. In Hyper-V Settings for LON-HOST2, click **Replication Configuration**.
4. In Replication Configuration pane, click **Enable this computer as a Replica server**.
5. In the Authentication and ports section select **Use Kerberos (HTTP)**.
6. In the **Authorization and storage** section click **Allow replication from any authenticated server** and then click **Browse**.
7. Click on **Computer**, then double click **Local Disk (E)** and then click **New folder**. Type **VMReplica** for folder name and press Enter. Select **E:\VMReplica\** folder and then click **Select Folder**.
8. In Hyper-V Settings for LON-HOST2, click **OK**.
9. In the Settings window, read the notice and click **OK**.
10. Click to the **Start** screen and then click **Control Panel**.
11. In the Control Panel, click **System and Security**, and then click **Windows Firewall**.
12. Click **Advanced settings**.
13. Click **Inbound Rules**.

14. In the right pane, in the rule list, find the rule **Hyper-V Replica HTTP Listener (TCP-In)**. Right-click the rule and click **Enable Rule**.
15. Close the Windows Firewall with Advanced Security console and then close **Windows Firewall**.
16. Repeat steps 1-15 on **LON-HOST1**.

► **Task 3: Configure replication for LON-CORE virtual machine**

1. On LON-HOST1, open Hyper-V Manager. Click **LON-HOST1**, and then right-click **20417A-LON-CORE**.
2. Click **Enable Replication...**
3. On the **Before You Begin** page, click **Next**.
4. On the **Specify Replica Server** page, click **Browse**.
5. In the Select Computer window type **LON-HOST2** and then click **Check Names** and then click **OK**. Click **Next**.
6. On the **Specify Connection Parameters** page, review settings, and make sure that **Use Kerberos authentication (HTTP)** is selected. Click **Next**.
7. On the **Choose Replication VHDs** page, make sure that **20410A-LON-CORE.vhd** is selected and then click **Next**.
8. On the **Configure Recovery History** page, select **Only the latest recovery point** and then click **Next**.
9. On the **Choose Initial Replication Method** page, click **Send initial copy over the network** and select **Start replication immediately**, and then click **Next**.
10. On the **Completing the Enable Replication wizard** page, click **Finish**.
11. Wait 10-15 minutes. You can monitor the progress of initial replication in the **Status** column in Hyper-V Manager console. When it completes (progress reaches 100%) make sure that **20417A-LON-CORE** has appeared on LON-HOST2 in Hyper-V Manager.

► **Task 4: Validate a planned failover to the replica site**

1. On LON-HOST2 in Hyper-V Manager, right-click **20417A-LON-CORE**.
2. Select **Replication** and then click **View Replication Health**.
3. Review content of the window that appears and make sure that there are not errors.
4. Click **Close**.
5. On LON-HOST1, open Hyper-V Manager and verify that **20417A-LON-CORE** is turned off.
6. Right-click **20417A-LON-CORE**, select **Replication**, and then click **Planned Failover...**
7. In the Planned Failover window, make sure that option **Start the Replica virtual machine after failover** is selected and then click **Fail Over**.
8. In the Planned Failover window click **Close**.
9. On LON-HOST2, in Hyper-V Manager, make sure that **20417A-LON-CORE** is running.
10. On LON-HOST1, right-click **20417A-LON-CORE**, point to **Replication** and then click **Remove replication**.

11. In the **Remove replication** dialog box, click **Remove Replication**.
12. On LON-HOST2, right-click **20417A-LON-CORE** and select **Shut Down**. In the **Shut Down Machine** dialog box, click **Shut Down**.

**Results:** After completing this exercise you will have Hyper-V replica configured.

## Exercise 2: Configuring a Failover Cluster for Hyper-V

### ► Task 1: Connect to iSCSI target from both host machines

1. On LON-HOST1, open **Server Manager**, click **Tools**, and then click **iSCSI Initiator**. At the **Microsoft iSCSI** prompt, click **Yes**.
2. Click the **Discovery** tab.
3. Click **Discover Portal**.
4. In the **IP address or DNS name** box, type **172.16.0.21**, and then click **OK**.
5. Click the **Targets** tab.
6. Click **Refresh**.
7. In the **Targets** list, select **iqn.1991-05.com.microsoft:lon-svr1-target1-target**, and then click **Connect**.
8. Select **Add this connection to the list of Favorite Targets**, and then click **OK**.
9. Click **OK** to close iSCSI Initiator Properties.
10. On LON-HOST2, open Server Manager, click **Tools**, and then click **iSCSI Initiator**.
11. In the **Microsoft iSCSI** dialog box, click **Yes**.
12. Click the **Discovery** tab.
13. Click **Discover Portal**.
14. In the IP address or DNS name box, type **172.16.0.21**, and then click **OK**.
15. Click the **Targets** tab.
16. Click **Refresh**.
17. In the **Discovered targets** list, select **iqn.1991-05.com.microsoft:lon-svr1-target1-target**, and then click **Connect**.
18. Select **Add this connection to the list of Favorite Targets**, and then click **OK**. Click **OK** to close iSCSI Initiator Properties.
19. On LON-HOST2, in the Server Manager window, click **Tools**, and then click **Computer Management**.
20. Expand **Storage**, and then click **Disk Management**.
21. Right-click **Disk 2**, and then click **Online**.
22. Right-click **Disk 2**, and then click **Initialize Disk**. In the **Initialize Disk** dialog box, click **OK**.
23. Right-click the unallocated space next to **Disk 2**, and then click **New Simple Volume**.
24. On the **Welcome** page, click **Next**.
25. On the **Specify Volume Size** page, click **Next**.
26. On the **Assign Drive Letter or Path** page, click **Next**.

27. On the **Format Partition** page, in the **Volume label** box, type **ClusterDisk**. Select the **Perform a quick format** check box, and then click **Next**.
28. Click **Finish**.
29. Repeat steps 21 through 28 for Disk 3 and Disk 4. In step 27, provide name **ClusterVMs** for Disk 3 and **Quorum** for Disk 4.
30. On LON-HOST1 in Server Manager, click **Tools**, and then click **Computer Management**.
31. Expand **Storage**, and then click **Disk Management**.
32. Right-click **Disk Management**, and then click **Refresh**.
33. Right-click **Disk 2**, and then click **Online**.
34. Right-click **Disk 3**, and then click **Online**.
35. Right-click **Disk 4**, and then click **Online**.

► **Task 2: Configure failover clustering on both host machines**

1. On LON-HOST1, on the taskbar, click the **Server Manager** icon to open Server Manager.
2. From the **Dashboard**, click **Add roles and features**.
3. On the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, click **Next**.
5. On the **Select destination server** page, make sure that **Select server from the server pool** is selected, and then click **Next**.
6. On the **Select server roles** page, click **Next**.
7. On the **Select features** page, in the **Features** list, click **Failover Clustering**. In the **Add features that are required for failover clustering** prompt, click **Add Features**, and then click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. When installation is complete, click **Close**.
10. Repeat steps 1 through 9 on **LON-HOST2**.
11. On LON-HOST1, in the Server Manager console, click **Tools** and then click **Failover Cluster Manager**.
12. In Failover Cluster Manager, in the center pane, under **Management**, click **Create Cluster**.
13. In the Create Cluster Wizard on the **Before You Begin** page, read the information. Click **Next**.
14. In the **Enter server name** box, type **LON-HOST1**, and then click **Add**. Type **LON-HOST2**, and then click **Add**.
15. Verify the entries, and then click **Next**.
16. On the **Validation Warning** page, click **No. I don't require support from Microsoft for this cluster** and click **Next**.
17. In the **Access Point for Administering the Cluster** page, in the **Cluster Name** box, type **VMCluster**.
18. Under **Address**, in the **IP address name** box, type **172.16.0.126**, and then click **Next**.
19. In the **Confirmation** dialog box, verify the information, remove the checkmark next to **Add all eligible storage to the cluster**, and then click **Next**.
20. In the Create Cluster Wizard Summary page, click **Finish**.



### ► Task 3: Configure disks for failover cluster

1. On LON-HOST1, in the Failover Cluster Manager console, expand **VMCluster.Adatum.com**, expand **Storage** and right-click **Disks**.
2. Click **Add Disk**.
3. In the **Add Disks to Cluster** dialog box, verify that all disks are selected, and then click **OK**.
4. Verify that all disks appear available for cluster storage in **Failover Cluster Manager**.
5. Select the disk that displays the Volume name of **ClusterVMs**. Right-click the **ClusterVMs** disk and select **Add to Cluster Shared Volumes**.
6. Right-click **VMCluster.adatum.com**, select **More Actions** and then click **Configure Cluster Quorum Settings**. Click **Next**.
7. On the **Select Quorum Configuration Option** page, click **Use typical settings** and then click **Next**.
8. On the **Confirmation** page click **Next**.
9. On the **Summary** page, click **Finish**.

## Exercise 3: Configuring a Highly Available Virtual Machine

### ► Task 1: Move virtual machine storage to iSCSI target



**Note:** Make sure that LON-HOST1 is the owner of the ClusterVMs disk in Failover Cluster Manager. If it is not, then move the ClusterVMs resource to LON-HOST1 before doing this procedure.

- On LON-HOST1, open Windows Explorer and browse to **E:\Program Files\Microsoft Learning\20417\Drives\20410A-LON-CORE\Virtual Hard Disks** and move the **20417A-LON-CORE.vhd** virtual hard drive file to the **C:\ClusterStorage\Volume1** location.

### ► Task 2: Configure the virtual machine as Highly Available

1. In the Failover Cluster Manager console click **Roles** and then in the Actions pane, click **Virtual Machines**.
2. Click **New Virtual Machine**.
3. Select **LON-Host2** as the cluster node and then click **OK**.
4. In the New Virtual Machine Wizard, click **Next**.
5. On the **Specify Name and Location** page, type **TestClusterVM** for the **Name** and then click **Store the virtual machine in a different location** and then click **Browse**.
6. Browse to and select **C:\ClusterStorage\Volume1** and then click **Select Folder**.
7. Click **Next**.
8. On the **Assign Memory** page, type **1536** and then click **Next**.
9. On the **Configure Networking** page click select **Corporate Network** and then click **Next**.
10. On the **Connect Virtual Hard Disk** page click **Use an existing virtual hard disk** and then click **Browse**.
11. Locate **C:\ClusterStorage\Volume1** and select **20417A-LON-CORE.vhd** and then click **Open**.

12. Click **Next** and then click **Finish**.
13. On the **Summary** page of the High Availability Wizard click **Finish**.
14. Right-click the **TestClusterVM** and then click **Start**.
15. Make sure that the machine successfully starts.

► **Task 3: Perform a Live Migration for the virtual machine**

1. Open **Failover Cluster Manager** on LON-HOST2.
2. Expand **VMCluster.Adatum.com**, and then click **Roles**.
3. Right-click **TestClusterVM** and select **Move**, then select **Live Migration** and then click **Select Node....**
4. Click **LON-Host1** and then click **OK**.
5. Right-click **TestClusterVM** and then click **Connect**.
6. Make sure that you can access and operate virtual machine while it is migrating to another host.
7. Wait until migration is finished.

► **Task 4: Perform a Storage Migration for the virtual machine**

1. On Lon-host1, open Hyper-V Manager.
2. In the central pane click **LON-GUEST1**.
3. In the Actions pane, click **Start**. Wait until the virtual machine is fully started.
4. Switch back to Hyper-V Manager console, and in the Actions pane click **Move**.
5. On the **Before You Begin** page click **Next**.
6. On the **Choose Move Type** page select **Move the virtual machine's storage** and then click **Next**.
7. On the **Choose Options for Moving Storage** page, select **Move all of the virtual machine's data to a single location** and then click **Next**.
8. On the **Choose a new location for virtual machine** page, click **Browse**.
9. Locate **C:\** and then create a new folder called **Guest1**. Click **Select Folder**.
10. Click **Next**.
11. On the Summary page click **Finish**. Wait for move process to finish. While virtual machine is moving you can connect to it, and verify that it is fully operational.
12. Shut down all running virtual machines.

► **To prepare for next module**

1. Restart **LON-HOST1**.
2. When you are prompted with the boot menu select **Windows Server 2008 R2** and press Enter.
3. Log on to the host machine as directed by your instructor.
4. Repeat steps 1-3 on **LON-HOST2**.

## Module 10: Implementing Dynamic Access Control

# Lab: Implementing Dynamic Access Control

### Exercise 1: Planning the Dynamic Access Control Implementation and Preparing AD DS for Dynamic Access Control

#### ► Task 1: Plan the Dynamic Access Control Deployment Based on the Security and Business Requirements

Scenario requires the following:

1. Folders that belong to Research department can be accessed and modified only by employees that belong to Research department.
2. Files classified with classification High should be accessible only to Managers.
3. Managers should access confidential files only from workstations that belong to the ManagersWKS security group.



**Note:** You can meet these requirements by implementing claims, resource properties, and file classifications, used together in Dynamic Access Control. To implement this, you should first create appropriate claims for users and devices. User claim uses department as its source attribute, while device claim uses description as source attribute. After that, you should configure resource property for Research department. When you have these objects prepared, you should configure Central Access Rules and Central Access Policies to protect resources. At the same time, you should configure file classification for confidential documents. Finally, you should apply Central Access Policy to folders where files for Research and Managers are located.

4. As a solution for users that receive error messages, you should implemented Access Denied Assistance.

#### ► Task 2: Prepare AD DS to support Dynamic Access Control

1. On LON-DC1, in the Server Manager, click **Tools** and then click **Active Directory Users and Computers**.
2. In the Active Directory Users and Computers console, right-click **Adatum.com** and select **New**, and then click **Organizational Unit**.
3. In the New Object – Organizational Unit, in the **Name** field, type **Test** and then click **OK**.
4. Click the **Computers** container.
5. Press the Ctrl key and click the **LON-SVR1**, **LON-CL1** and **LON-CL2** computers. Right-click and select **Move...**
6. In the Move window, click **Test** and then click **OK**.
7. Close the Active Directory Users and Computers console.
8. On LON-DC1, in the Server Manager, click **Tools**, and then click **Group Policy Management**.
9. Expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**.
10. Right-click the **Managers** OU and then click **Block Inheritance**. This is to remove the block inheritance setting used in a later module in the course.
11. Click the **Group Policy Objects** container.

12. In the results pane, right-click **Default Domain Controllers Policy**, and then click **Edit**.
13. In the Group Policy Management Editor, under **Computer Configuration**, expand **Policies**, expand **Administrative Templates**, expand **System**, and then click **KDC**.
14. In the right pane, double-click **KDC support for claims, compound authentication and Kerberos armoring**.
15. In the KDC support for claims, compound authentication and Kerberos armoring window, select **Enabled**, and in the **Options** section, click the drop-down list and select **Supported**. Click **OK**.
16. Close the Group Policy Management Editor and Group Policy Management console.
17. Open Windows Power Shell, by clicking its icon on the task bar, and type **gpupdate /force** and press Enter. After Group Policy is updated, close Windows PowerShell.
18. On LON-DC1, in Server Manager, click **Tools**, and then click **Active Directory Users and Computers**.
19. Expand **Adatum.com**, right-click **Users**, click **New**, and then click **Group**.
20. Type **ManagersWKS** for the **Group name**, and then click **OK**.
21. Click the **Test** container.
22. Right-click **LON-CL1**, and then click **Properties**.
23. Click the **Member Of** tab and then click **Add**.
24. In the Select Groups window, type **ManagersWKS**. click **Check Names**, click **OK**, and then click **OK** again.
25. Click **Managers** organization unit.
26. Right-click **Aidan Delaney** and select **Properties**.
27. Click the **Organization** tab. Make sure that the **Department** field is populated with the value **Managers**. Click **Cancel**.
28. Click the **Research** organization unit.
29. Right-click **Allie Bellew** and select **Properties**.
30. Click the **Organization** tab. Make sure that the **Department** field is populated with the value **Research**. Click **Cancel**.

**Results:** After completing this exercise you will have design for Dynamic Access Control and you will have prepared AD DS for Dynamic Access Control implementation.

## Exercise 2: Configuring User and Device Claims

### ► Task 1: Review the Default Claim Types

1. On LON-DC1, in Server Manager, click **Tools** and then click **Active Directory Administrative Center**.
2. In the Active Directory Administrative Center console, in navigation pane, click **Dynamic Access Control**.
3. In the central pane double-click **Claim Types**.
4. Verify that there are no default claims defined.
5. In the navigation pane, click **Dynamic Access Control** and then double-click **Resource Properties**.

6. Review the default resource properties.



**Note:** Note that all properties are disabled by default.

7. In the navigation pane, click **Dynamic Access Control** and then double-click **Resource Property Lists**.
8. In the central pane right-click **Global Resource Property List**, and then click **Properties**.
9. In the **Global Resource Property List**, in the **Resource Properties**, section review available resource properties.
10. Click **Cancel**.

► **Task 2: Configure Claims for Users**

1. In the **Active Directory Administrative Center**, in the navigation pane, click **Dynamic Access Control**.
2. Double-click **Claim Types**.
3. In the Tasks pane, click **New** and then click **Claim Type**.
4. In the Create Claim Type window, in the **Source Attribute** section, select **department**.
5. In the **Display name** text box type **Company Department**.
6. Select both **User** and **Computer** check boxes.
7. Click **OK**.

► **Task 3: Configure Claims for Devices**

1. In the **Active Directory Administrative Center**, in the Tasks pane, click **New** and select **Claim Type**.
2. In the Create Claim Type window, in the **Source Attribute** section, select **description**.
3. Clear the **User** check box and select the **Computer** check box.
4. Click **OK**.

**Results:** After completing this exercise you will have configured user and device claims.

### Exercise 3: Configuring Resource Properties and File Classifications

► **Task 1: Configure Resource Property definitions**

1. In the **Active Directory Administrative Center**, click **Dynamic Access Control**.
2. In the central pane, double-click **Resource Properties**.
3. In the **Resource Properties** list, locate **Department**.
4. Right-click **Department**, and then click **Enable**.
5. In the **Resource Properties** list, locate **Confidentiality**.
6. Right-click **Confidentiality**, and then click **Enable**.
7. Make sure that both **Department** and **Confidentiality** properties are enabled in the list.
8. Double-click **Department**.

9. Scroll down to the **Suggested Values** section, and then click **Add**.
10. In the Add a suggested value window, type **Research** in both **Value** and **Display name** text boxes, and then click **OK** two times.
11. Click **Dynamic Access Control** and then double-click **Resource Property Lists**.
12. In the central pane, double-click **Global Resource Property List**.
13. Make sure that both **Department** and **Confidentiality** appear in **Resource Properties** list. If they do not, then click **Add** and add these two properties, and then click **OK** (or **Cancel** if you did not make any changes).
14. Close the Active Directory Administrative Center.

► **Task 2: Classify files**

1. On LON-SVR1, in Server Manager, click **Add roles and features**.
2. In the Add Roles and Features Wizard click **Next** three times.
3. On the **Select server roles** page, expand **File and Storage Services (Installed)**, expand **File and iSCSI Service (Installed)** and select **File Server Resource Manager**.
4. When prompted, click **Add Features**.
5. Click **Next** two times and then click **Install**. After installation finishes, click **Close**.
6. In Server Manager, click **Tools**, and then click **File Server Resource Manager**.
7. In the File Server Resource Manager console, expand **Classification Management**.
8. Select and then right-click **Classification Properties** and click **Refresh**.
9. Verify that **Confidentiality** and **Department** properties are in the list.
10. Click **Classification Rules**.
11. In the Actions pane, click **Create Classification Rule**.
12. In the Create Classification Rule window, enter **Set Confidentiality** for the **Rule** name.
13. Click the **Scope** tab. Click **Add**.
14. In the **Browse For Folder** dialog box, expand **Local Disk (C:)** and select the **Docs** folder, and then click **OK**.
15. Click the **Classification** tab.
16. Make sure that following settings are set:
  - Classification method: Content Classifier
  - Property: Confidentiality
  - Value: High
17. Click **Configure**.
18. In the **Classification Parameters** dialog box, click the **Regular expression** drop-down list and select **String**.
19. In the **Expression** field (next to the word String) type **secret**.
20. Click **OK**.
21. Click the **Evaluation Type** tab. Select **Re-evaluate existing property values**, and then click **Overwrite the existing value**.

22. Click **OK**.
23. In the File Server Resource Manager, in the Actions pane, click **Run Classification with all rules now**.
24. Select **Wait for classification to complete**, and then click **OK**.
25. After the classification is complete, you are presented with a report. Verify that two files were classified.



**Note:** You can see this in the Report Totals section.

26. Close the report.
27. Open Windows Explorer, and browse to the **C:\Docs** folder.
28. Right-click **Doc1.txt** and select **Properties**.
29. Click the **Classification** tab. Verify that **Confidentiality** is set to **High**.
30. Repeat steps 28 and 29 on files Doc2.txt and Doc3.txt.



**Note:** Doc2.txt should have the same confidentiality as Doc1.txt while Doc3.txt should have no value. This is because only Doc1 and Doc2 have the word *secret* in their content.

### ► Task 3: Assign properties to folder

1. On LON-SVR1, open Windows Explorer, and browse to Local Disk (C:).
2. Right-click the **Research** folder and then click **Properties**.
3. Click **Classification** tab.
4. Click **Department**.
5. In the **Value** section click **Research**. Click **Apply**.
6. Click **OK**.

**Results:** After this exercise, you will have configured resource properties and file classifications.

## Exercise 4: Configuring Central Access Rules and Policies

### ► Task 1: Configure Central Access Policy Rules

1. On LON-DC1, in Server Manager, click **Tools** and then click **Active Directory Administrative Center**.
2. In the Active Directory Administrative Center console, in the navigation pane, click **Dynamic Access Control**.
3. Double-click **Central Access Rules**.
4. In the Tasks pane, click **New**, and then click **Central Access Rule**.
5. In the **Central Access Rule** dialog box, type **Department Match** for the **Name**.
6. In the **Target Resources** section click **Edit**.
7. In the **Central Access Rule** dialog box, click **Add a condition**.

8. Set a condition as follows: **Resource-Department-Equals-Value-Research**, and then click **OK**.
9. In the **Permissions** section, click **Use the following permissions as current permissions**.
10. In the **Permissions** section, click **Edit**.
11. Remove permission for **Administrators**.
12. In Advanced Security Settings for Permissions, click **Add**.
13. In Permission Entry for Permissions, click **Select a principal**.
14. In the Select User, Computer, Service Account or Group window, type **Authenticated Users**, click **Check Names**, and then click **OK**.
15. In the **Basic permissions** section select **Modify, Read and Execute, Read and Write**.
16. Click **Add a condition**.
17. Click the **Group** drop-down list, and select **Company Department**.
18. On the **Value** drop-down list, and select **Resource**.
19. In the last drop-down box, select **Department**.



**Note:** As a result, you should have: User-Company Department-Equals-Resource-Department.

20. Click **OK** three times.
21. In the Tasks pane, click **New**, and then click **Central Access Rule**.
22. For the name of rule type **Access Confidential Docs**.
23. In the **Target Resources** section click **Edit**.
24. In the Central Access Rule window click **Add a condition**.
25. In the last drop-down box select **High**.



**Note:** You should have this expression as a result: Resource-Confidentiality-Equals-Value-High.

26. Click **OK**.
27. In the **Permissions** section, click **Use the following permissions as current permissions**.
28. In the **Permissions** section, click **Edit**.
29. Remove permission for **Administrators**.
30. In Advanced Security Settings for Permissions, click **Add**.
31. In the Permission Entry for Permissions, click **Select a principal**.
32. In the Select User, Computer, Service Account or Group window, type **Authenticated Users**, click **Check Names**, and then click **OK**.
33. In the Basic permissions section, select **Modify, Read and Execute, Read and Write**.
34. Click **Add a condition**.
35. Set first condition to:  
**User-Group-Member of each-Value-Managers**. Click **Add a condition**.



36. Set second condition to: **Device-Group-Member of each-Value-ManagersWKS**.



**Note:** If you can't find ManagersWKS in the last drop-down box, click **Add items**. Then in the Select User, Computer, Service Account or Group window, type **ManagersWKS** and click **Check Names**. Click **OK**.

37. Click **OK** three times.

### ► Task 2: Create Central Access Policy

1. On LON-DC1, in Active Directory Administrative Center, click **Dynamic Access Control**, and then double-click **Central Access Policies**.
2. In the Tasks pane, click **New**, and then click **Central Access Policy**.
3. For the **Name**, type **Protect confidential docs**.
4. Click **Add**.
5. Click the **Access Confidential Docs** rule, and then click >>.
6. Click **OK** twice.
7. In the Tasks pane, click **New**, and then click **Central Access Policy**.
8. For the **Name**, type **Department Match**.
9. Click **Add**.
10. Click the **Department Match** rule and then click >>.
11. Click **OK** twice.
12. Close the Active Directory Administrative Center.

### ► Task 3: Publish Central Access Policy with Group Policy

1. On LON-DC1, in Server Manager, click **Tools**, and then click **Group Policy Management**.
2. Under **Domains**, expand **Adatum.com**, and then right-click **Test** and click **Create a GPO in this domain, and link it here**.
3. Type **DAC Policy**, and then click **OK**.
4. Right-click **DAC Policy**, and then click **Edit**.
5. Browse to **Computer Configuration/Policies/Windows Settings/Security Settings/File System**, and then right-click **Central Access Policy**.
6. Click **Manage Central Access Policies**.
7. Click both **Department Match** and **Protect confidential docs**, and then click **Add**.
8. Click **OK**.
9. Close the Group Policy Management Editor.
10. Close the Group Policy Management console.

### ► Task 4: Apply Central Access Policy to resources

1. On LON-SVR1, start Windows PowerShell.
2. Type **gpupdate /force** and press Enter. Close the Command Prompt window.
3. Open Windows Explorer, browse to Drive C and right-click the **Docs** folder, and select **Properties**.

4. Click **Security** tab.
5. Click **Advanced**.
6. In the Advanced Security Settings for Docs window, click the **Central Policy** tab.
7. Click **Change**.
8. On the drop-down list, select **Protect confidential docs**.
9. Click **OK** two times.
10. Right-click the **Research** folder and select **Properties**.
11. Click **Security** tab.
12. Click **Advanced**.
13. In the Advanced Security Settings for Research window, click the **Central Policy** tab.
14. Click **Change**.
15. In drop-down box, select **Department Match**.
16. Click **OK** two times.

► **Task 5: Configure access denied remediation settings**

1. On LON-DC1, in Server Manager, click **Tools**, and then click **Group Policy Management**.
2. Expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**.
3. Click **Group Policy objects**.
4. Right-click **DAC Policy** and select **Edit**.
5. Under **Computer Configuration**, expand **Policies**, expand **Administrative Templates**, expand **System**, and then click **Access-Denied Assistance**.
6. In the right pane, double-click **Customize Message for Access Denied errors**.
7. In the Customize Message for Access Denied errors window, click **Enabled**.
8. In the **Display the following message to users who are denied access** text box, type: **You are denied access because of permission policy. Please request access.**
9. Select the **Enable users to request assistance** check box.
10. Review other options, do not make any changes, and then click **OK**.
11. In the right pane of Group Policy Management Editor, double-click **Enable access-denied assistance on client for all file types**.
12. Click **Enabled**, and then click **OK**.
13. Close the Group Policy Management Editor and close the Group Policy Management console.
14. Switch to LON-SVR1, open Windows PowerShell and type **gpupdate /force** and press Enter.

**Results:** After completing this exercise you will have configured central access rules and policies.

## Exercise 5: Validating and Remediating Access Control

### ► Task 1: Verify Dynamic Access Control functionality

1. Log on to **LON-CL1** as **Adatum\April** with password **Pa\$\$w0rd**.
2. Click **Desktop** and then open Windows Explorer by clicking its icon on the task bar.
3. In the address bar, type **\\LON-SVR1\Docs**, and then press Enter.
4. Try to open **Doc3**. You should be able to open that document.
5. In the address bar of Windows Explorer, type **\\LON-SVR1\Research** and press Enter.



**Note:** You should be unable to access folder.

6. Click **Request assistance**. Review options for sending messages, and then click **Close**.
7. Log off of LON-CL1.
8. Log on to **LON-CL1** as **Adatum\Allie** with the password of **Pa\$\$w0rd**.
9. Open Windows Explorer.
10. In the address bar, type **\\LON-SVR1\Research** and press Enter.



**Note:** You should be able to access this folder and open documents inside because Allie is in Research department.

11. Log off of LON-CL1.
12. Log on to **LON-CL1** as **Adatum\Aidan** with the password of **Pa\$\$w0rd**.
13. Open Windows Explorer.
14. In the address bar, type **\\LON-SVR1\Docs**.
15. You should be able to open all files in this folder.
16. Log off of LON-CL1.
17. Log on to **LON-CL2** as **Adatum\Aidan** with the password of **Pa\$\$w0rd**.
18. Open Windows Explorer
19. In the address bar, type **\\LON-SVR1\Docs**.



**Note:** You should be unable to see Doc1 and Doc2 since LON-CL2 is not permitted to view secret documents.

### ► Task 2: Configure staging for Dynamic Access Policy

1. On LON-DC1, in Server Manager, click **Tools** and then click **Group Policy Management**.
2. In the Group Policy Management console, expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**, and then click **Group Policy objects**.
3. Right-click **DAC Policy** and click **Edit**.
4. In the Group Policy Management Editor, browse to **Computer Configuration/Policies /Windows Settings/Security Settings/Advanced Audit Policy Configuration/Audit Policies**.

5. Select **Object Access**.
6. Double-click **Audit Central Access Policy Staging**. Select all three check boxes, and then click **OK**.
7. Double-click **Audit File System**. Select all three check boxes, then click **OK**.
8. Close the Group Policy Management Editor and the Group Policy Management console.

► **Task 3: Configure staging permissions**

1. On LON-DC1, open Server Manager, and then open Active Directory Administrative Center.
2. In the navigation pane, click **Dynamic Access Control**.
3. Double-click **Central Access Rules**.
4. Right-click **Department Match** and select **Properties**.
5. Scroll down to **Proposed Permissions**.
6. Click **Enable permission staging configuration**.
7. Click **Edit**.
8. Click **Authenticated Users**, and then click **Edit**.
9. Change the condition to: **User-Company Department-Equals-Value-Marketing**.
10. Click **OK** three times.
11. Switch to **LON-SVR1** and open Windows PowerShell.
12. Type **gpupdate /force** and press Enter.
13. Close the Windows PowerShell window.

► **Task 4: Verify staging**

1. Log on to **LON-CL1** as **Adatum\Adam** with the password of **Pa\$\$w0rd**.
2. Open Windows Explorer, and then in the address bar type **\\LON-SVR1\Research**. Attempt to open the folder. You will be unsuccessful. Click **Close**.
3. Switch to **LON-SVR1**.
4. In Server Manager, click **Tools** and select **Event Viewer**.
5. Expand **Windows Logs**, and then click **Security**.
6. Look for Events with ID 4818.
7. Read the content of these logs.

► **Task 5: Use effective permissions to test Dynamic Access Control**

1. On LON-SVR1, open Windows Explorer and locate the **C:\Research** folder.
2. Right-click the folder and click **Properties**.
3. Click **Security** tab.
4. Click **Advanced**, and then click **Effective Access**.
5. Click **select a user**.
6. In the Select User, Computer, Service Account, or Group window type **April**, and then click **Check Names**, and then click **OK**.
7. Click **View effective access**.

8. Review results. April should not have any access to this folder.
9. Click **Include a user claim**.
10. On the drop-down list, select **Company Department**.
11. In the **Value** text box type **Research**.
12. Click **View Effective access**. April should have access now.
13. Close all windows.

► **Task 6: To prepare for next module**

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20417A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20417A-LON-SVR1**, **20417A-LON-CL1** and **20417A-LON-CL2**.

**Results:** After this exercises you will have validated Dynamic Access Control functionality.

**MCT USE ONLY. STUDENT USE PROHIBITED**

# Module 11: Implementing Active Directory Domain Services

## Lab: Implementing AD DS

### Exercise 1: Deploying a Read-Only Domain Controller

#### ► Task 1: Add LON-SVR3 as a server to manage

1. Log on to **LON-DC1** as **Administrator** with a password of **Pa\$\$w0rd**.
2. In Server Manager Dashboard, click **Add other servers to manage**.
3. In the **Add Servers** dialog box, in the **Name (CN)** field, type **LON-SVR3**, and then click **Find Now**.
4. Select the **LON-SVR3** server in the details pane, and then click the **arrow** to move it to the Selected pane.
5. Click **OK**.

#### ► Task 2: Create a new Server Group

1. In the Server Manager Dashboard, click **Create a server group**.
2. In the **Create Server Group** dialog box, in the **Server group name** field, type **DCs**.
3. Select both **LON-SVR3** and **LON-DC1**, click the arrow to move them to the Selected pane, and then click **OK**.

#### ► Task 3: Install the RODC role remotely

1. In the Server Manager Dashboard, click **Add roles and features**.
2. In the Add Roles and Features Wizard, click **Next**.
3. On the **Select installation type** page, click **Next**.
4. On the **Select Destination Server** page, select **LON-SVR3.Adatum.com**, and then click **Next**.
5. On the **Select server role** page, click the check box for **Active Directory Domain Services**, click **Add Features** in the **Add features that are required for Active Directory Domain Services** dialog box, and then click **Next**.
6. On the **Select features** page, click **Next**.
7. On the **Active Directory Domain Services**, page click **Next**.
8. On the **Confirm installation selections** page, click the check box to **Restart the destination server automatically if required**, and then click **Install**. The installation will take several minutes.
9. When the installation is complete, click **Close**.
10. In Server Manager Dashboard, click the **notification** icon (the flag icon or yellow triangle) on the menu bar.
11. Locate the **Post-deployment Configuration** task, and then click **Promote this server to a domain controller**.
12. In the Active Directory Domain Services Configuration Wizard, ensure that **Add a domain controller to an existing domain** is selected.
13. In the **Supply the credentials to perform this operation** section, click **Change**.
14. In the **Windows Security** dialog box, type **Adatum\Administrator** in the user name field and in the password field type **Pa\$\$w0rd**.

15. Click **OK**, and then click **Next**.
16. On the **Domain Controller Options** page, select the check box for **Read only domain controller (RODC)**.
17. Type and confirm the Directory services Restore Mode (DSRM) password to be **Pa\$\$w0rd**, and then click **Next**.
18. On the **RODC Options** page, click **Next**.



**Note:** You will configure these options in the next exercise.

19. On the **Additional Options** page click **Next**.
20. On the **Paths** page click **Next**.
21. On the **Review Options** page click **Next**.
22. On the **Prerequisites Check** page, click **Install**.



**Note:** The installation will take several minutes and LON-SVR3 will automatically restart to complete the promotion.

23. When the promotion is completed click **Close**. Note that LON-SVR3 is restarting.

► **Task 4: Configure the Password Replication policy and administrative access**

1. On LON-DC1, in Server Manager, on the **Tools** menu, click **Active Directory Users and Computers**.
2. Expand **Adatum.com**, and then click the **Domain Controllers** OU.
3. In the details pane, right-click **LON-SVR3**, and then click **Properties**.
4. In the **LON-SVR3 Properties** dialog box, click the **Password Replication Policy** tab.
5. Click **Add**.
6. In the **Add Groups, Users and Computers** dialog box, click **Allow passwords for the account to replicate to this RODC**, and then click **OK**.
7. In the **Select Users, Computers, Services Accounts, or Groups** dialog box, type **Managers**, and then click **OK**.
8. Click the **Managed By** tab, and then click **Change**.
9. In the **Select User or Group** dialog box, type **IT**, and then click **OK**.
10. Click **OK** to close the LON-SVR3 Properties dialog box.

**Results:** After completing this exercise, you will have added LON-SVR3 as a server to manage, created a server group, deployed an RODC remotely, and configured the password replication policy and administrative assignments for the RODC.



## Exercise 2: Troubleshooting Group Policy

### ► Task 1: Troubleshoot Group Policy issues

1. Log on to **LON-CL1** as **Brad** with a password of **Pa\$\$word**. Brad is a member of the IT group.
2. At the **Start** screen, type **Control Panel**.
3. In the **Apps** results field click **Control Panel**.
4. In Control Panel under **Appearance and Personalization**, click **Change desktop background**.

**Question:** What is the result?

**Answer:** A message explains that this feature is disabled.

**Question:** Is this in line with company policy?

**Answer:** Yes, this is in line with company policy.

5. Close Control Panel.
6. Point to the lower right corner of the desktop, click the **Search** charm and in the **Apps** search field, type **Run**.
7. In the **Apps** results field click **Run**.
8. In the **Run** box type **Regedit**, and then click **OK**.

**Question:** What is the result?

**Answer:** A message explains that this feature is disabled.

**Question:** Is this in line with company policy?

**Answer:** No, this is against company policy.

9. To close the dialog box, click **OK**.
10. Point to the lower right corner of the desktop, click the **Search** charm and then in the **Apps** search field, type **Command Prompt**.
11. In the **Apps** results field, click **Command Prompt**.
12. In the Command Prompt window, type **GPREsult /R** and examine the results.

**Question:** What GPOs are being applied in User Settings?

**Answer:** The Prohibit Desktop Background policy and the Prohibit Registry Tools GPOs are being applied.

**Question:** Is this in line with company policy?

**Answer:** No, this is against company policy. The Prohibit Registry Tools policy should not be applied to an IT group user.

13. Sign out of LON-CL1.
14. Log on to **LON-CL1** as **Bill** with a password of **Pa\$\$word**. Bill is a member of the Managers group.
15. On the **Start** screen, type **Control Panel**.
16. In the **Apps** results field, click **Control Panel**.

17. In Control Panel under **Appearance and Personalization**, click **Change desktop background**.

**Question:** What is the result?

**Answer:** The Desktop Background dialog box appears and provides access to change the desktop background.

**Question:** Is this in line with company policy?

**Answer:** No, this is against company policy.

18. Close **Control Panel**.

19. Point to the lower right corner of the desktop, click the **Search** charm, and then type **Run**.

20. In the **Apps** results field, click **Run**.

21. In the **Run** box, type **Regedit**, and then click **OK**.

**Question:** What is the result?

**Answer:** The Registry Editor application starts.

**Question:** Is this in line with company policy?

**Answer:** No, this is against company policy.

22. Close the Registry Editor.

23. Point to the lower right corner of the desktop, click the **Search** charm, and type **Command Prompt** in the **Apps** search field.

24. Click **Command Prompt** in the **Apps** results field.

25. In the Command Prompt window, type **GPResult /R** and examine the results.

**Question:** What GPOs are being applied?

**Answer:** No GPOs are being applied.

**Question:** Is this correct?

**Answer:** No, both GPOs are supposed to be applied.

26. Sign Out of LON-CL1.

► **Task 2: Correct issues with Group Policy application**

1. Log on to **LON-DC1** as **Administrator** with a password of **Pa\$\$w0rd**.

2. In Server Manager, on the **Tools** menu, click **Group Policy Management**.

3. If required, expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**.

**Question:** What GPOs are linked to the Adatum.com domain?

**Answer:** Default Domain Policy, Prohibit Registry Tools and Prohibit Desktop Background. This confirms the policies are linked to the correct container.

**Question:** What is the current status of the Managers OU?

**Answer:** The Managers OU has blue circle with a white exclamation mark. This indicates the inheritance is being blocked. You must remove the inheritance block to resolve the issue with the Managers OU.

4. Right-click the **Managers** OU and clear the check mark next to **Block Inheritance**.

**Question:** How will you ensure that the Prohibit Registry Tools GPO will not be applied to the IT group users?

**Answer:** There are multiple ways that you could resolve this. For example, you could create a GPO that specifically reverses the Prevent access to registry editing tools setting and link it directly to the IT OU.

5. Expand the **Group Policy Objects** folder.
6. Click the **Prohibit Registry Tools** GPO.
7. In the details pane, click the **Delegation** tab.
8. Click **Advanced**.
9. In the **Prohibit Registry Tools Security Settings** dialog box, click **Add**.
10. In the **Select Users, Computers, Service Accounts, or Groups** dialog box type **IT**, and then click **OK**.
11. Click the **IT (Adatum\IT)** group in the **Security** list.
12. In the **Permissions for IT** section, locate the **Apply Group Policy** permission, and then click **Deny**.
13. Click **OK**.
14. If the **Windows Security** dialog box appears, click **Yes** to acknowledge the message.
15. Close the Group Policy Management console.

► **Task 3: Verify policies are being applied**

1. Log on to **LON-CL1** as **Bill** with a password of **Pa\$\$w0rd**.
2. On the Start screen, type **Command Prompt**.
3. In the **Apps** results field, click **Command Prompt**.
4. In the Command Prompt window, type **GPResult /R** and examine the results.

**Question:** What GPOs are being applied?

**Answer:** The Prohibit Desktop Background and the Prohibit Registry Tools.

**Question:** Is this correct?

**Answer:** Yes. The system is now in line with the company policy.

5. Sign Out of LON-CL1.
6. Log on to **LON-CL1** as **Brad** with a password of **Pa\$\$w0rd**.
7. On the Start screen, type **Command Prompt**.
8. In the **Apps** results field, click **Command Prompt**.
9. In the Command Prompt window, type **GPResult /R** and examine the results.

**Question:** What GPOs are being applied?

**Answer:** The Prohibit Desktop Background GPO is being applied.

**Question:** What GPOs are being filtered out?

**Answer:** Prohibit Registry Tools is being denied.

10. Sign Out of LON-CL1.

**Results:** After completing this exercise, you will be able to troubleshoot Group Policy issues, correct issues to apply Group Policy, and verify policies are being applied.

### Exercise 3: Implementing Service Accounts in AD DS

#### ► Task 1: Create and associate a Managed Service account

1. Log on to **LON-DC1** as **Administrator** with a password of **Pa\$\$w0rd**.
2. Right-click **Windows PowerShell** on the Taskbar and click **Run as Administrator**.
3. In the Windows PowerShell command window, type **Add-KdsRootKey -EffectiveTime ((get-date).addhours(-10))** at the prompt and press Enter.
4. Type **New-ADServiceAccount -Name Webservice -DNSHostName LON-DC1 -PrincipalsAllowedToRetrieveManagedPassword LON-DC1\$** and press Enter.
5. Type **Add-ADComputerServiceAccount -identity LON-DC1 -ServiceAccount Webservice** and press Enter.
6. Type **Get-ADServiceAccount -Filter \*** and press Enter to verify the account. Note the output of the command.
7. Type **Install-ADServiceAccount -Identity Webservice** and press Enter.
8. Minimize the Windows PowerShell command window.

#### ► Task 2: Configure the Web Server Application Pool to use the Group Managed Service account

1. On LON-DC1, in Server Manager, click the **Tools** menu and click **Internet Information Services (IIS) Manager**.
2. In the Internet Information Services (IIS) Manager console, expand **LON-DC1 (Adatum\Administrator)** and click **Application Pools**.
3. In the details pane, right-click the **DefaultAppPool** and click **Advanced Settings**.
4. In the **Advanced Settings** dialog box, click **Identity** and click the ellipses.
5. In the **Application Pool Identity** dialog box, click **Custom Account** and click **Set**.
6. In the **Set Credentials** dialog box, type **Adatum\Webservice\$** in the **User name:** field and click **OK** three times.
7. In the Actions pane, click **Stop** to stop the application pool.
8. Click **Start** to start the application pool.
9. Close the Internet Information Services (IIS) Manager.

**Results:** After completing this exercise, you will have created and associated a managed service account, installed a managed service account on a web server, and verified password change for a managed service account.

## Exercise 4: Maintaining AD DS

### ► Task 1: Create and view Active Directory snapshots

1. Switch to **LON-DC1**.
2. Move your mouse to the bottom right corner and click the **Search** charm.
3. In the **Apps** search box, type **CMD**.
4. In the Apps Results for CMD pane, right-click **Command Prompt** and then click **Run as administrator**.
5. In the command window, type **Ntdsutil** and then press Enter.
6. Type **Snapshot** and then press Enter.
7. Type **Activate instance ntds** and then press Enter.
8. Type **Create** and then press Enter.



**Note:** The GUID that is displayed is important for commands in later tasks. Make note of the GUID or, alternatively, copy it to the clipboard.

9. Mount the snapshot as a new instance of AD DS by running the following command: Mount {GUID} where {GUID} is the GUID returned by the create snapshot command.
10. Type **Quit** twice.
11. Expose the snapshot by typing **dsamain -dbpath c:\\$snap\_datetime\_volumec\$\windows\ntds\ntds.dit -ldapport 50000**, and then press Enter.



**Note:** Hint: Copy and paste the \$snap\_datetime from the previous command. (The port number can be any open, unique TCP port). Leave the Command Window open and the command running while you perform the next tasks.

12. In Server Manager, click the **Tools** menu and then click **Active Directory Users and Computers**.
13. Expand **Adatum.com** and then click **Research**.
14. In the details pane, right-click **Allie Bellew** and then click **Delete**. Click **Yes** to confirm in the message box.
15. Right click the **Active Directory Users and Computers** root node and then click **Change Domain Controller**.
16. Click **<Type a Directory Server name[:port] here>** and type **LON-DC1:50000** and then press Enter.
17. Click **OK**.
18. Expand **Adatum.com** and click **Research**.



**Note:** Notice that the user Allie Bellew exists in the snapshot because it was taken before the user was deleted.

19. Close Active Directory Users and Computers and close the command window.

► **Task 2: Enable the Active Directory recycle bin**

1. In Server Manager, on the **Tools** menu, click **Active Directory Administrative Center**.
2. In the navigation pane, click **Adatum (local)**.
3. In the Tasks pane, click **Enable Recycle Bin**.
4. In the **Enable Recycle Bin Confirmation** dialog box, click **OK**.
5. In the **Active Directory Administrative Center** dialog box, click **OK**.
6. On the menu bar, click the **Refresh** icon.



**Note:** Notice a Deleted Object container now appears.

► **Task 3: Delete a test user**

1. In the center pane, double-click the **Managers** OU.
2. Ensure that the **Aidan Delaney** user account is selected, and then in the tasks pane, click **Delete**.
3. In the **Delete Confirmation** dialog box, click **Yes**.
4. Click **Adatum (local)** in the navigation pane to return to the main tree.

► **Task 4: Restore the deleted user**

1. In the center pane, double-click the **Deleted Objects** folder.
2. In the Tasks pane, click **Restore**. In the navigation pane under **Adatum (local)**, click **Managers**.



**Note:** Note that the Aidan Delaney account is restored.

► **Task 5: To prepare for the next module**

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20417A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20417A-LON-CL1** and **20417A-LON-SVR3**.

**Results:** After completing this exercise, you will have created and viewed Active Directory snapshots, enabled the Active Directory Recycle Bin, deleted a user as a test, and used the Active Directory Administrative Center to restore a deleted user account.

## Module 12: Implementing Active Directory Federation Services

### Lab: Implementing AD FS

#### Exercise 1: Configuring AD FS Prerequisites

##### ► Task 1: Configure DNS forwarders

1. On LON-DC1, in Server Manager, click **Tools**, and then click **DNS**.
2. Expand **LON-DC1**, and click **Conditional Forwarders**.
3. Right-click **Conditional Forwarders**, and click **New Conditional Forwarder**.
4. In the **DNS Domain** box, type **TreyResearch.com**.
5. Click in the **IP address** column, and then type **172.16.10.10**. Press Enter, and then click **OK**.
6. Close the DNS Manager.
7. On MUN-DC1, in Server Manager, click **Tools**, and then click **DNS**.
8. Expand **MUN-DC1**, and then click **Conditional Forwarders**.
9. Right-click **Conditional Forwarders**, and then click **New Conditional Forwarder**.
10. In the **DNS Domain** box, type **Adatum.com**.
11. Click in the **IP address** column, and then type **172.16.0.10**. Press Enter, and then click **OK**.
12. Close the DNS Manager.

##### ► Task 2: Exchange root certificates to enable certificate trusts

1. On LON-DC1, access the **Search** page.
2. In the **Search** box, type **\\MUN-DC1.treyresearch.com\certenroll**, and then press Enter.
3. In the CertEnroll window, right-click the **MUN-DC1.TreyResearch.com\_TreyResearch-MUN-DC1-CA.crt** file, and then click **Copy**.
4. In the left pane, click **Documents**, and then paste the file into the **Documents** folder.
5. Open a Windows PowerShell® command prompt, type **MMC** and then press Enter.
6. In the Console1 window, click **File**, and click **Add/Remove Snap-in**.
7. Click **Group Policy Management Editor**, and then click **Add**.
8. In **Select Group Policy Object**, click **Browse**.
9. Click **Default Domain Policy**, and then click **OK**.
10. Click **Finish**, and then click **OK**.
11. Double-click **Default Domain Policy**. In the console tree, expand the following path: **Computer Configuration > Policies > Windows Settings > Security Settings > Public Key Policies > Trusted Root Certification Authorities**.
12. Right-click **Trusted Root Certification Authorities**, and then click **Import**.
13. On the Welcome to the Certificate Import Wizard page, click **Next**.
14. On the **File to Import** page, click **Browse**.

15. In the Open window, click **MUN-DC1.TreyResearch.com\_TreyResearch-MUN-DC1-CA.crt**, click **Open**, and then click **Next**.
16. On the **Certificate Store** page, verify that **Place all certificates in the following store** is selected, verify that the **Trusted Root Certification Authorities** store is listed, and then click **Next**.
17. On the **Completing the Certificate Import Wizard** page, click **Finish**, and then click **OK**.
18. Close the Group Policy Management Editor without saving changes.
19. On MUN-DC1, access the **Search** page.
20. In the **Search** box, type **\\LON-DC1.adatum.com\certenroll**, and then press Enter.
21. In the CertEnroll window, right-click the **LON-DC1.Adatum.com\_Adatum-LON-DC1-CA.crt** file, and then click **Copy**.
22. In the left pane, click **Documents**, and then paste the file into the **Documents** folder.
23. Open a Windows PowerShell command prompt, type **MMC**, and then press Enter.
24. In the Console1 window, click **File**, and then click **Add/Remove Snap-in**.
25. Click **Certificates**, and click **Add**.
26. Click **Computer Account**, and then click **Next**.
27. Verify that **Local computer** is selected, click **Finish**, and then click **OK**.
28. Expand **Certificates**, and then click **Trusted Root Certification Authorities**.
29. Right-click **Trusted Root Certification Authorities**, point to **All Tasks**, and then click **Import**.
30. On the **Welcome to the Certificate Import Wizard** page, click **Next**.
31. On the **File to Import** page, click **Browse**.
32. In the open window, click **LON-DC1.Adatum.com\_Adatum-LON-DC1-CA.crt**, click **Open**, and then click **Next**.
33. On the **Certificate Store** page, verify that **Place all certificates in the following store** is selected, verify that the **Trusted Root Certification Authorities** store is listed, and then click **Next**.
34. On the **Completing the Certificate Import Wizard** page, click **Finish**, and then click **OK**.
35. Close Console1 without saving changes.

► **Task 3: Request and install a certificate for the web server**

1. On LON-SVR1, in Server Manager, click **Tools**, and then click **Internet Information Services (IIS) Manager**.
2. In the console tree, click **LON-SVR1 (Adatum\Administrator)**. Click **No** to dismiss the message.
3. In middle pane, double-click **Server Certificates**.
4. In the Actions pane, click **Create Domain Certificate**.
5. On the **Distinguished Name Properties** page, enter the settings as listed below, and then click **Next**:
  - Common name: LON-SVR1.adatum.com
  - Organization: A. Datum
  - Organization unit: IT
  - City/locality: London



- State/province: England
  - Country/region: GB
6. On the **Online Certification Authority** page, in **Specify Online Certification Authority**, click **Select** to search for a CA server in the domain.
  7. Select **Adatum-LON-DC1-CA**, and then click **OK**.
  8. In **Friendly name**, type **LON-SVR1.adatum.com**, and then click **Finish**.

► **Task 4: Bind the certificate to the claims aware application on the web server and verify application access**

1. On LON-SVR1, in Internet Information Services (IIS) Manager, expand **Sites**, click **Default Web Site**, and then in the Actions pane, click **Bindings**.
2. In the **Site Bindings** dialog box, click **Add**.
3. In the **Add Site Binding** dialog box, under **Type** select **https**, and under **Port**, verify that **443** is selected
4. In the **SSL Certificate** drop-down list, click **LON-SVR1.adatum.com**, and then click **OK**.
5. Click **Close**, and then close Internet Information Services (IIS) Manager.
6. On LON-DC1, open Internet Explorer.
7. Connect to **https://lon-svr1.adatum.com/adatumtestapp**.
8. Verify that you can connect to the site, but that you receive a 401 access denied error. This is expected because you have not yet configured AD FS for authentication.
9. Close Internet Explorer.

**Results:** In this exercise, you configured DNS forwarding to enable name resolution between A. Datum and Trey Research, and you exchanged root certificates between the two organizations. You also installed and configured a web certificate on the application server.

## Exercise 2: Installing and Configuring AD FS

► **Task 1: Install and configure AD FS 2.0**

1. On the LON-DC1, in Server Manager, click **Manage**, and then click **Add Roles and Features**.
2. On the **Before you begin** page, click **Next**.
3. On the **Select installation type** page, click **Next**.
4. On the **Select destination server** page, click **Next**.
5. On the **Select server roles** page, select the **Active Directory Federation Services** check box, click **Add Features**, and then click **Next**.
6. On the **Select features** page, click **Next**.
7. On the **Active Directory Federation Services (AD FS)** page, click **Next**.
8. On the **Select role services** page, click **Next**.
9. On the **Confirm installation selections** page, click **Install**, and then wait for the installation to finish. Do not close the window.

### ► Task 2: Create a stand-alone Federation Server by using the AD FS Federation Server Configuration Wizard

1. On the **Installation progress** page, click **Run the AD FS Management snap-in**.
2. In the Overview pane, click the **AD FS Federation Server Configuration Wizard** link.
3. On the **Welcome** page, ensure that **Create a new Federation Service** is selected, and then click **Next**.
4. On the **Select Stand-Alone or Farm Deployment** page, click **Stand-alone federation server**, and then click **Next**.
5. On the **Specify the Federation Service Name** page, ensure that the **SSL certificate** selected is **LON-DC1.Adatum.com**, the **Port** is **443**, and the **Federation Service name** is **LON-DC1.Adatum.com**. Click **Next**.
6. On the **Ready to Apply Settings** page, verify that the correct configuration settings are listed, and then click **Next**.
7. Wait for the configuration to finish, and then click **Close**.

### ► Task 3: Verify that FederationMetaData.xml is present and contains valid data

1. Log on to the **LON-CL1** virtual machine as **Adatum\Brad** using the password **Pa\$\$w0rd**.
2. Click the Desktop tile, and then open Internet Explorer.
3. Click the **Settings** icon in the top-right corner, and then click **Internet options**.
4. On the **Security** tab, click **Local intranet**.
5. Click **Sites**, and then clear the **Automatically detect intranet network** check box.
6. Click **Advanced**, and in the **Add this website to the zone** box, type **https://lon-dc1.adatum.com**, and then click **Add**.
7. Type **https://lon-svr1.adatum.com**, click **Add**, and then click **Close**.
8. Click **OK** twice.
9. Connect to **https://lon-dc1.adatum.com/federationmetadata/2007-06/federationmetadata.xml**.
10. Verify that the xml file opens successfully, and then scroll through its contents.
11. Close Internet Explorer.

**Results:** In this exercise, you installed and configured the AD FS server role, and then verified a successful installation by viewing the Federation Meta Data .xml contents.

## Exercise 3: Configure AD FS for a Single Organization

### ► Task 1: Configure a Token Signing Certificate for LON-DC1.Adatum.com

1. On the LON-DC1 virtual machine, in Server Manager, click **Tools**, and then click **Windows PowerShell**.
2. At the prompt, type **set-ADFSProperties -AutoCertificateRollover \$False**, and then press Enter. This step is required so that you can modify the certificates that AD FS uses.
3. Close the Windows PowerShell window.

4. Click **Tools**, and click **AD FS Management**.
5. In the AD FS console, in the left pane, expand **Service**, and then click **Certificates**.
6. Right-click **Certificates**, and then click **Add Token-Signing Certificate**.
7. In the **Select a token signing certificate** dialog box, click **LON-DC1.Adatum.com**, and then click **OK**.
8. In the AD FS Management warning, click **OK**.



**Note:** Verify that the certificate has a subject of CN=LON-DC1.Adatum.com. If no name is listed under the **Subject** when you add the certificate, delete the certificate, and then add the next certificate in the list.

9. Right-click the newly added certificate, and then click **Set as Primary**. Note the warning message, and then click **Yes**.
10. Select the certificate that has just been superseded, right-click the certificate, and then click **Delete**. Click **Yes** to confirm the deletion.

### ► Task 2: Configure the Active Directory Claims Provider Trust

1. In the AD FS console, expand **Trust Relationships**, and then click claims provider **Trusts**.
2. In the middle pane, right-click **Active Directory**, and then click **Edit Claim Rules**.
3. In the Edit Claims Rules for Active Directory window, on the **Acceptance Transform Rules** tab, click **Add Rule**.
4. The Add Transform Claim Rule Wizard appears.
5. On the **Select Rule Template** page, under **Claim rule template**, select **Send LDAP Attributes as Claims**, and then click **Next**.
6. On the **Configure Rule** page, in the **Claim rule name** box, type **Outbound LDAP Attributes Rule**.
7. In the **Attribute store** drop-down list, select **Active Directory**.
8. In the **Mapping of LDAP attributes to outgoing claim types** section, select the following values for the LDAP Attribute and the Outgoing Claim Type:
  - E-Mail-Addresses = **E-Mail Address**
  - User-Principal-Name = **UPN**
  - Display-Name = **Name**
9. Click **Finish**, and then click **OK**.

### ► Task 3: Configure the claims application to trust incoming claims by running the WIF Federation Utility

1. On LON-SVR1, click to the **Start** screen, and then click **Windows Identity Foundation Federation Utility**.
2. On the **Welcome to the Federation Utility wizard** page, in **Application configuration location**, type **C:\inetpub\wwwroot\AdatumTestApp\web.config** for the location of the web.config file of the WIF sample application.
3. In **Application URI**, type **https://lon-svr1.adatum.com/AdatumTestApp/** to indicate the path to the sample application that will trust the incoming claims from the federation server. Click **Next** to continue.

4. On the **Security Token Service** page, select **Use an existing STS**, type **https://lon-dc1.adatum.com/federationmetadata/2007-06/federationmetadata.xml** for the STS WS-Federation metadata document location, and then click **Next** to continue. In the warning, click **Yes**.
5. On the **Security token encryption** page, select **No encryption**, and then click **Next**.
6. On the **Offered claims** page, review the claims that will be offered by the federation server, and then click **Next**.
7. On the **Summary** page, review the changes that will be made to the sample application by the Federation Utility Wizard, scroll through the items to understand what each item is doing, and then click **Finish**.
8. Click **OK**.

► **Task 4: Configure a relying party trust for the claims aware application**

1. On LON-DC1, in the AD FS Management console, click **AD FS**.
2. In the middle pane, click **Required: Add a trusted relying party**.
3. On the **Welcome** page of the Add relying party Trust Wizard, click **Start**.
4. On the **Select Data Source** page, select **Import data about the relying party published online or on a local network**, and then type **https://lon-svr1.adatum.com/adatumtestapp**.
5. Click **Next** to continue.



**Note:** This action prompts the wizard to check for the MetaData of the application that the web server role hosts.

6. On the **Specify Display Name** page, in the **Display name** box, type **ADatum Test App**, and then click **Next**.
7. On the **Choose Issuance Authorization Rules** page, ensure that the **Permit all users to access this relying party** is selected, and then click **Next**.
8. On the **Ready to Add Trust** page, review the relying party trust settings, and then click **Next**.
9. On the **Finish** page, click **Close**. The Edit Claim Rules for ADatum Test App window opens.

► **Task 5: Configure claim rules for the relying party trust**

1. In the Edit Claim Rules for WIF Sample Claims App window, on the **Issuance Transform Rules** tab, click **Add Rule**. The Add Transform Claim Rule Wizard opens.
2. On the **Select Rule Template** page, under **Claim rule template**, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.



**Note:** This action passes an incoming claim through to the user by means of Windows Integrated Authentication.

3. On the **Configure Rule** page, in **Claim rule name**, type **Pass through Windows Account name rule**. In the **Incoming claim type** drop-down list, select **Windows account name**, and then click **Finish**.
4. Click **Add Rule**.

5. On the **Select Rule Template** page, under **Claim rule template**, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
6. On the **Configure Rule** page, in **Claim rule name**, type **Pass through E-mail Address rule**. In the **Incoming claim type** drop-down list, select **E-mail Address**, and then click **Finish**.
7. Click **Add Rule**.
8. On the **Select Rule Template** page, under **Claim rule template**, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
9. On the **Configure Rule** page, in **Claim rule name**, type **Pass through UPN rule**. In the **Incoming claim type** drop-down list, select **UPN**, and then click **Finish**.
10. Click **Add Rule**.
11. On the **Select Rule Template** page, under **Claim rule template**, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
12. On the **Configure Rule** page, in **Claim rule name**, type **Pass through Name rule**. In the **Incoming claim type** drop-down list, select **Name**, and then click **Finish**.
13. Click **Apply**, and then click **OK**.

► **Task 6: Test the access to the claims aware application**

1. On LON-CL1, open Internet Explorer.
2. Connect to **https://lon-svr1.adatum.com/AdatumTestApp/**.



**Note:** Note: Ensure that you type the trailing "/"

3. If you are prompted for credentials, type **Adatum\Brad** with password **Pa\$\$w0rd**, and then press Enter. The page renders, and then you see the claims that were processed to allow access to the web site.

**Results:** After this exercise, you configured a token signing certificate and configured a claims provider trust for Adatum.com. You also configured the sample application to trust incoming claims and configured a relying party trust and associated claim rules. You also tested access to the sample WIF application in a single organization scenario.

## Exercise 4: Configure AD FS for Federated Business Partners

► **Task 1: Add a claims provider trust for the TreyResearch.com AD FS server**

1. On LON-DC1, if required, in Server Manager, click **Tools**, and click **AD FS Management**.
2. In the AD FS console, expand **Trust Relationships**, and then click claims provider **Trusts**.
3. In the Actions pane, click **Add claims provider Trust**.
4. On the **Welcome** page, click **Start**.
5. On the **Select Data Source** page, select **Import data about the claims provider published online or on a local network**, type **https://mun-dc1.treyresearch.com**, and then click **Next**.
6. On the **Specify Display Name** page, click **Next**.
7. On the **Ready to Add Trust** page, review the claims provider trust settings, and then click **Next** to save the configuration.

8. On the **Finish** page, click **Close** to close the wizard. The Edit Claim Rules for mun-dc1.treyresearch.com window appears.
9. On the **Acceptance Transform Rules** tab, click **Add Rule**.
10. In the **Claim rule template** list, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
11. In the **Claim rule name** box, type **Pass through Windows account name rule**.
12. In the **Incoming claim type** drop-down list, select **Windows account name**.
13. Select **Pass through all claim values**, and then click **Finish**. Click **Yes**.
14. Click **OK**, and then close the AD FS console.
15. On LON-DC1, in Server Manager, click **Tools**, and then click **Windows PowerShell**.
16. At the prompt, type the following command, and then press Enter:

```
Set-ADFSClaimsProviderTrust -TargetName "mun-dc1.treyresearch.com" -  
SigningCertificateRevocationCheck None
```

17. Close the Windows PowerShell window.

► **Task 2: Configure a relying party trust on MUN-DC1 for A. Datum's claim aware application**

1. On the MUN-DC1, in Server Manager, click **Tools**, and then click **AD FS Management**.
2. In the AD FS console, on the **Overview** page, click **Required: Add a trusted relying party**.
3. On the **Welcome** page, click **Start**.
4. On the **Select Data Source** page, select **Import data about the relying party published online or on a local network**, type **https://lon-dc1.adatum.com**, and then click **Next**.
5. On the **Specify Display Name** page, in the **Display name** box, type **Adatum TestApp**, and then click **Next**.
6. On the **Choose Issuance Authorization Rules** page, select **Permit all users to access this relying party**, and then click **Next**.
7. On the **Ready to Add Trust** page, review the relying party trust settings, and then click **Next** to save the configuration.
8. On the **Finish** page, click **Close** to close the wizard. The Edit Claim Rules for Adatum TestApp window appears.
9. On the **Issuance Transform Rules** tab, click **Add Rule**.
10. In the **Claim rule template** list, select **Pass Through or Filter an Incoming claim**, and then click **Next**.
11. In the **Claim rule name** box, type **Pass through Windows account name rule**.
12. In the **Incoming Claim type** drop-down list, select **Windows account name**.
13. Select **Pass through all claim values**, and then click **Finish**.
14. Click **OK**, and then close the AD FS console.

► **Task 3: Verify access to the A. Datum Test Application for Trey Research users**

1. On MUN-DC1, open Internet Explorer, and connect to **https://lon-svr1.adatum.com/adatumtestapp/**.



**Note:** The logon process has changed, and you must now select an authority that can authorize and validate the access request. The **Home Realm Discovery** page (the Sign In page) appears, and you must select an authority.

2. On the **Sign In** page, select **mun-dc1.treyresearch.com**, and then click **Continue to Sign in**.
3. When prompted for credentials, type **TreyResearch\April** with password **Pa\$\$w0rd**, and then press Enter. You should be able to access the application.
4. Close Internet Explorer.
5. Open Internet Explorer, and then connect to **https://lon-svr1.adatum.com/adatumtestapp/** again.
6. When prompted for credentials, type **TreyResearch\April** with password **Pa\$\$w0rd**, and then press Enter. You should be able to access the application.
7. Close Internet Explorer.



**Note:** You are not prompted for a home realm again. Once users have selected a home realm and been authenticated by a realm authority, they are issued with an \_LSRealm cookie by the relying party Federation Server. The default lifetime for the cookie is 30 days. Therefore, for us to log on multiple times, we should delete that cookie after each logon attempt to return to a clean state.

#### ► Task 4: Configure claim rules for the claim provider trust and the relying party trust to allow access only for a certain group

1. On MUN-DC1, in the AD FS console, expand **Trust Relationships**, and then click relying party **Trusts**.
2. Select **Adatum TestApp**, and in the Actions pane, click **Edit Claim Rules**.
3. On the Edit Claim Rules for Adatum TestApp window, on the **Issuance Transform Rules** tab, click **Add Rule**.
4. On the **Select Rule Template** page, under **Claim rule template**, select **Send Group Membership as a Claim**, and then click **Next**.
5. On the **Configure Rule** page, in **Claim rule name**, type **Permit Production Group Rule**.
6. Beside **User's Group**, click **Browse**, type **Production** and click **OK**.
7. Under **Outgoing claim type**, click **Group**.
8. Under **Outgoing claim value**, type **Production**, click **Finish** and then click **OK**.
9. On LON-DC1, if required, open the AD FS Management console.
10. In the AD FS console, expand **Trust Relationships**, and then click **Claim Provider Trusts**.
11. Select **mun-dc1.treyresearch.com**, and in the Actions pane, click **Edit Claim Rules**.
12. On the **Acceptance Transform Rules** tab, click **Add Rule**.
13. On the **Select Rule Template** page, under **Claim rule template**, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
14. On the **Configure Rule** page, in **Claim rule name**, type **Send Production Group Rule**.
15. In the **Incoming claim type** drop down list, click **Group**, and click **Finish**. Click **Yes** and then click **OK**.
16. In the AD FS console, under **Trust Relationships**, click relying party **Trusts**.



17. Select the **Adatum Test App**, and in the Actions pane, click **Edit Claim Rules**.
18. On the **Issuance Transform Rules** tab, click **Add Rule**.
19. Under **Claim rule template**, click **Pass Through or Filter an Incoming Claim**, and then click **Next**.
20. Under **Claim rule name**, type **Send TreyResearch Group Name Rule**.
21. In the **Incoming claim type** drop down list, click **Group**. Click **Finish**.
22. On the Edit Claim Rules for Adatum Test App window, on the **Issuance Authorization Rules** tab, select the rule named **Permit Access to All Users**, and click **Remove Rule**. Click **Yes** to confirm. With no rules, no users are permitted access.
23. On the **Issuance Authorization Rules** tab, click **Add Rule**.
24. On the **Select Rule Template** page, under **Claim rule template**, select **Permit or Deny Users Based on an Incoming Claim**, and then click **Next**.
25. On the **Configure Rule** page, in **Claim rule name** type **Permit TreyResearch Production Group Rule**, in the **Incoming claim type** drop-down list, select **Group**. In **Incoming claim value**, type **Production**, select the option to **Permit access to users with this incoming claim**, and then click **Finish**.
26. On the **Issuance Authorization Rules** tab, click **Add Rule**.
27. On the **Select Rule Template** page, under **Claim rule template**, select **Permit or Deny Users Based on an Incoming Claim**, and then click **Next**.
28. On the **Configure Rule** page, in **Claim rule name** type **Temp**, in the **Incoming claim type** drop-down list, select **UPN**. In **Incoming claim value**, type **@adatum.com**, select the option to **Permit access to users with this incoming claim**, and then click **Finish**.
29. Click the **Temp** rule, and click **Edit Rule**.
30. In the **Edit Rule –Temp** dialog box, click **View Rule Language**.
31. Press Ctrl + C to copy the rule language to the clipboard. Click **OK**.
32. Click **Cancel**.
33. Click the **Temp** rule, click **Remove Rule**, and then click **Yes**.
34. On the **Issuance Authorization Rules** tab, click **Add Rule**.
35. On the **Select Rule Template** page, under **Claim rule template**, select **Send Claims Using a Custom Rule**, and then click **Next**.
36. On the **Configure Rule** page, type **ADatum User Access Rule** as the **Claim rule name**.
37. Click in the **Custom rule** box, and then press Ctrl+V to paste the clipboard contents into the box. Edit the first URL to match the following text, and then click **Finish**:

```
c:[Type == "http://schemas.xmlsoap.org/ws/2005/05/identity/claims/upn", Value =~
"^(?:i).+@adatum\.com$"]=> issue(Type =
"http://schemas.microsoft.com/authorization/claims/permit", Value = "PermitUsersWithClaim");
```



**Note:** This rule enables access to anyone who presents a claim that includes the UPN of @adatum.com. The Value line in the first URL defines the attribute that must be matched in the claim. In this line, ^ indicates the beginning of the string to match, (?:i) means that the text is case insensitive, .+ means that one or more characters will be added, and \$ means the end of the string.



38. Click **OK** to close the property page and save the changes to the relying party trust.

► **Task 5: Verify restrictions and accessibility to the claims aware application**

1. On MUN-DC1, open Internet Explorer, connect to On MUN-DC1, launch Internet Explorer, and then connect to **https://lon-svr1.adatum.com/adatumtestapp/**.
2. When prompted for credentials, type **TreyResearch\April** with the password **Pa\$\$w0rd**, and then press Enter.



**Note:** April is not a member of the Production group, so she should not be able to access the application.

3. Close Internet Explorer.
4. Open Internet Explorer, click the **Settings** icon in the top-right corner, and then click **Internet options**.
5. Under **Browsing history**, click **Delete**, click **Delete** again, and then click **OK**.
6. Connect to **https://lon-svr1.adatum.com/adatumtestapp/**.
7. Select **mun-dc1.treyresearch.com** on the **Sign In** page, and then click **Continue to Sign in**.
8. When prompted for credentials, type **TreyResearch\Morgan** with the password **Pa\$\$w0rd**, and then press Enter.



**Note:** Morgan is a member of the Production group, so she should be able to access the application.

9. Close Internet Explorer.

► **Task 6: To shut down the virtual machines**

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20417A-MUN-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20417A-LON-CL1**, **20417A-LON-SVR1** and **20417A-LON-DC1**.

**Results:** In this exercise, you configured a claims provider trust for Trey Research on Adatum.com and a relying party trust for Adatum on TreyResearch.com. You verified access to the A. Datum claim-aware application. Then you configured the application to restrict access from TreyResearc.com to specific groups, and you verified appropriate access.

**MCT USE ONLY. STUDENT USE PROHIBITED**