

OFFICIAL MICROSOFT LEARNING PRODUCT

20410C

Installing and Configuring 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, e-mail 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.

© 2014 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: 20410C

Part Number: X19-09958

Released: 01/2014

MCT USE ONLY. STUDENT USE PROHIBITED

MICROSOFT LICENSE TERMS MICROSOFT INSTRUCTOR-LED COURSEWARE

These license terms are an agreement between Microsoft Corporation (or based on where you live, one of its affiliates) and you. Please read them. They apply to your use of the content accompanying this agreement which includes the media on which you received it, if any. These license terms also apply to Trainer Content and any updates and supplements for the Licensed Content unless other terms accompany those items. If so, those terms apply.

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

If you comply with these license terms, you have the rights below for each license you acquire.

1. DEFINITIONS.

- a. "Authorized Learning Center" means a Microsoft IT Academy Program Member, Microsoft Learning Competency Member, or such other entity as Microsoft may designate from time to time.
- b. "Authorized Training Session" means the instructor-led training class using Microsoft Instructor-Led Courseware conducted by a Trainer at or through an Authorized Learning Center.
- c. "Classroom Device" means one (1) dedicated, secure computer that an Authorized Learning Center owns or controls that is located at an Authorized Learning Center's training facilities that meets or exceeds the hardware level specified for the particular Microsoft Instructor-Led Courseware.
- d. "End User" means an individual who is (i) duly enrolled in and attending 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 content accompanying this agreement which may include the Microsoft Instructor-Led Courseware or Trainer Content.
- 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, and (ii) currently certified as a Microsoft Certified Trainer under the Microsoft Certification Program.
- g. "Microsoft Instructor-Led Courseware" means the Microsoft-branded instructor-led training course that educates IT professionals and developers on Microsoft technologies. A Microsoft Instructor-Led Courseware title may be branded as MOC, Microsoft Dynamics or Microsoft Business Group courseware.
- h. "Microsoft IT Academy Program Member" means an active member of the Microsoft IT Academy Program.
- i. "Microsoft Learning Competency Member" means an active member of the Microsoft Partner Network program in good standing that currently holds the Learning Competency status.
- j. "MOC" means the "Official Microsoft Learning Product" instructor-led courseware known as Microsoft Official Course that educates IT professionals and developers on Microsoft technologies.
- k. "MPN Member" means an active Microsoft Partner Network program member in good standing.

- l. "Personal Device" means one (1) personal computer, device, workstation or other digital electronic device that you personally own or control that meets or exceeds the hardware level specified for the particular Microsoft Instructor-Led Courseware.
- m. "Private Training Session" means the instructor-led training classes provided by MPN Members for corporate customers to teach a predefined learning objective using Microsoft Instructor-Led Courseware. 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.
- n. "Trainer" means (i) an academically accredited educator engaged by a Microsoft IT Academy Program Member to teach an Authorized Training Session, and/or (ii) a MCT.
- o. "Trainer Content" means the trainer version of the Microsoft Instructor-Led Courseware and additional supplemental content designated solely for Trainers' use to teach a training session using the Microsoft Instructor-Led Courseware. Trainer Content may include Microsoft PowerPoint presentations, trainer preparation guide, train the trainer materials, Microsoft One Note packs, classroom setup guide and Pre-release course feedback form. To clarify, Trainer Content does not include any software, virtual hard disks or virtual machines.

2. 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 five separate sets of use rights. Only one set of rights apply to you.

a. If you are a Microsoft IT Academy Program Member:

- i. Each license acquired on behalf of yourself may only be used to review one (1) copy of the Microsoft Instructor-Led Courseware in the form provided to you. If the Microsoft Instructor-Led Courseware is in digital format, you may install one (1) copy on up to three (3) Personal Devices. You may not install the Microsoft Instructor-Led Courseware on a device you do not own or control.
- ii. For each license you acquire on behalf of an End User or Trainer, you may either:
 - 1. distribute one (1) hard copy version of the Microsoft Instructor-Led Courseware to one (1) End User who is enrolled in the Authorized Training Session, and only immediately prior to the commencement of the Authorized Training Session that is the subject matter of the Microsoft Instructor-Led Courseware being provided, **or**
 - 2. provide one (1) End User with the unique redemption code and instructions on how they can access one (1) digital version of the Microsoft Instructor-Led Courseware, **or**
 - 3. provide one (1) Trainer with the unique redemption code and instructions on how they can access one (1) Trainer Content,

provided you comply with the following:
- iii. you will only provide access to the Licensed Content to those individuals who have acquired a valid license to the Licensed Content,
- iv. you will ensure each End User attending an Authorized Training Session has their own valid licensed copy of the Microsoft Instructor-Led Courseware that is the subject of the Authorized Training Session,
- v. you will ensure that each End User provided with the hard-copy version of the Microsoft Instructor-Led Courseware will be presented with a copy of this agreement and each End User will agree that their use of the Microsoft Instructor-Led Courseware will be subject to the terms in this agreement prior to providing them with the Microsoft Instructor-Led Courseware. Each individual will be required to denote their acceptance of this agreement in a manner that is enforceable under local law prior to their accessing the Microsoft Instructor-Led Courseware,
- vi. you will ensure that each Trainer teaching an Authorized Training Session has their own valid licensed copy of the Trainer Content that is the subject of the Authorized Training Session,

- vii. you will only use qualified Trainers who have in-depth knowledge of and experience with the Microsoft technology that is the subject of the Microsoft Instructor-Led Courseware being taught for all your Authorized Training Sessions,
- viii. you will only deliver a maximum of 15 hours of training per week for each Authorized Training Session that uses a MOC title, and
- ix. you acknowledge that Trainers that are not MCTs will not have access to all of the trainer resources for the Microsoft Instructor-Led Courseware.

b. If you are a Microsoft Learning Competency Member:

- i. Each license acquired on behalf of yourself may only be used to review one (1) copy of the Microsoft Instructor-Led Courseware in the form provided to you. If the Microsoft Instructor-Led Courseware is in digital format, you may install one (1) copy on up to three (3) Personal Devices. You may not install the Microsoft Instructor-Led Courseware on a device you do not own or control.
- ii. For each license you acquire on behalf of an End User or Trainer, you may either:
 - 1. distribute one (1) hard copy version of the Microsoft Instructor-Led Courseware to one (1) End User attending the Authorized Training Session and only immediately prior to the commencement of the Authorized Training Session that is the subject matter of the Microsoft Instructor-Led Courseware provided, **or**
 - 2. provide one (1) End User attending the Authorized Training Session with the unique redemption code and instructions on how they can access one (1) digital version of the Microsoft Instructor-Led Courseware, **or**
 - 3. you will provide one (1) Trainer with the unique redemption code and instructions on how they can access one (1) Trainer Content,**provided you comply with the following:**
- iii. you will only provide access to the Licensed Content to those individuals who have acquired a valid license to the Licensed Content,
- iv. you will ensure that each End User attending an Authorized Training Session has their own valid licensed copy of the Microsoft Instructor-Led Courseware that is the subject of the Authorized Training Session,
- v. you will ensure that each End User provided with a hard-copy version of the Microsoft Instructor-Led Courseware will be presented with a copy of this agreement and each End User will agree that their use of the Microsoft Instructor-Led Courseware will be subject to the terms in this agreement prior to providing them with the Microsoft Instructor-Led Courseware. Each individual will be required to denote their acceptance of this agreement in a manner that is enforceable under local law prior to their accessing the Microsoft Instructor-Led Courseware,
- vi. you will ensure that each Trainer teaching an Authorized Training Session has their own valid licensed copy of the Trainer Content that is the subject of the Authorized Training Session,
- vii. you will only use qualified Trainers who hold the applicable Microsoft Certification credential that is the subject of the Microsoft Instructor-Led Courseware being taught for your Authorized Training Sessions,
- viii. you will only use qualified MCTs who also hold the applicable Microsoft Certification credential that is the subject of the MOC title being taught for all your Authorized Training Sessions using MOC,
- ix. you will only provide access to the Microsoft Instructor-Led Courseware to End Users, and
- x. you will only provide access to the Trainer Content to Trainers.

c. If you are a MPN Member:

- i. Each license acquired on behalf of yourself may only be used to review one (1) copy of the Microsoft Instructor-Led Courseware in the form provided to you. If the Microsoft Instructor-Led Courseware is in digital format, you may install one (1) copy on up to three (3) Personal Devices. You may not install the Microsoft Instructor-Led Courseware on a device you do not own or control.
- ii. For each license you acquire on behalf of an End User or Trainer, you may either:
 1. distribute one (1) hard copy version of the Microsoft Instructor-Led Courseware to one (1) End User attending the Private Training Session, and only immediately prior to the commencement of the Private Training Session that is the subject matter of the Microsoft Instructor-Led Courseware being provided, **or**
 2. provide one (1) End User who is attending the Private Training Session with the unique redemption code and instructions on how they can access one (1) digital version of the Microsoft Instructor-Led Courseware, **or**
 3. you will provide one (1) Trainer who is teaching the Private Training Session with the unique redemption code and instructions on how they can access one (1) Trainer Content,**provided you comply with the following:**
- iii. you will only provide access to the Licensed Content to those individuals who have acquired a valid license to the Licensed Content,
- iv. you will ensure that each End User attending an Private Training Session has their own valid licensed copy of the Microsoft Instructor-Led Courseware that is the subject of the Private Training Session,
- v. you will ensure that each End User provided with a hard copy version of the Microsoft Instructor-Led Courseware will be presented with a copy of this agreement and each End User will agree that their use of the Microsoft Instructor-Led Courseware will be subject to the terms in this agreement prior to providing them with the Microsoft Instructor-Led Courseware. Each individual will be required to denote their acceptance of this agreement in a manner that is enforceable under local law prior to their accessing the Microsoft Instructor-Led Courseware,
- vi. you will ensure that each Trainer teaching an Private Training Session has their own valid licensed copy of the Trainer Content that is the subject of the Private Training Session,
- vii. you will only use qualified Trainers who hold the applicable Microsoft Certification credential that is the subject of the Microsoft Instructor-Led Courseware being taught for all your Private Training Sessions,
- viii. you will only use qualified MCTs who hold the applicable Microsoft Certification credential that is the subject of the MOC title being taught for all your Private Training Sessions using MOC,
- ix. you will only provide access to the Microsoft Instructor-Led Courseware to End Users, and
- x. you will only provide access to the Trainer Content to Trainers.

d. If you are an End User:

For each license you acquire, you may use the Microsoft Instructor-Led Courseware solely for your personal training use. If the Microsoft Instructor-Led Courseware is in digital format, you may access the Microsoft Instructor-Led Courseware online using the unique redemption code provided to you by the training provider and install and use one (1) copy of the Microsoft Instructor-Led Courseware on up to three (3) Personal Devices. You may also print one (1) copy of the Microsoft Instructor-Led Courseware. You may not install the Microsoft Instructor-Led Courseware on a device you do not own or control.

e. If you are a Trainer.

- i. For each license you acquire, you may install and use one (1) copy of the Trainer Content in the form provided to you on one (1) Personal Device solely to prepare and deliver an Authorized Training Session or Private Training Session, and install one (1) additional copy on another Personal Device as a backup copy, which may be used only to reinstall the Trainer Content. You may not install or use a copy of the Trainer Content on a device you do not own or control. You may also print one (1) copy of the Trainer Content solely to prepare for and deliver an Authorized Training Session or Private Training Session.

- ii. You may customize the written portions of the Trainer Content that are logically associated with instruction of a training session in accordance with the most recent version of the MCT agreement. If you elect to exercise the foregoing rights, you agree to comply with the following: (i) customizations may only be used for teaching Authorized Training Sessions and Private Training Sessions, and (ii) all customizations will comply 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 is licensed as a single unit and you may not separate their components and install them on different devices.

2.3 Redistribution of Licensed Content. Except as expressly provided in the use rights above, you may not distribute any 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 Notices. The Licensed Content may include third party code that Microsoft, not the third party, licenses to you under this agreement. Notices, if any, for the third party code are included for your information only.

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 your use of that respective component and supplements the terms described in this agreement.

3. LICENSED CONTENT BASED ON PRE-RELEASE TECHNOLOGY. If the Licensed Content's subject matter is based on a pre-release version of Microsoft technology ("Pre-release"), then in addition to the other provisions in this agreement, these terms also apply:

- a. **Pre-Release Licensed Content.** This Licensed Content subject matter is on the Pre-release version of the Microsoft technology. The technology may not work the way a final version of the technology will and we may change the technology for the final version. We also may not release a final version. Licensed Content based on the final version of the technology may not contain the same information as the Licensed Content based on the Pre-release version. Microsoft is under no obligation to provide you with any further content, including any Licensed Content based on the final version of the technology.
- 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 technology, 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 technology, technologies, or products to third parties because we include your feedback in them. These rights survive this agreement.
- c. **Pre-release Term.** If you are an Microsoft IT Academy Program Member, Microsoft Learning Competency Member, MPN Member or Trainer, you will cease using all copies of the Licensed Content on the Pre-release technology upon (i) the date which Microsoft informs you is the end date for using the Licensed Content on the Pre-release technology, or (ii) sixty (60) days after the commercial release of the technology that is the subject of the Licensed Content, whichever is earliest ("**Pre-release term**"). Upon expiration or termination of the Pre-release term, you will irretrievably delete and destroy all copies of the Licensed Content in your possession or under your control.

- 4. 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:
- access or allow any individual to access the Licensed Content if they have not acquired a valid license for the Licensed Content,
 - alter, remove or obscure any copyright or other protective notices (including watermarks), branding or identifications contained in the Licensed Content,
 - modify or create a derivative work of any Licensed Content,
 - publicly display, or make the Licensed Content available for others to access or use,
 - copy, print, install, sell, publish, transmit, lend, adapt, reuse, link to or post, make available or distribute the Licensed Content to any third party,
 - work around any technical limitations in the Licensed Content, or
 - 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.
- 5. 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.
- 6. 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.
- 7. SUPPORT SERVICES.** Because the Licensed Content is "as is", we may not provide support services for it.
- 8. 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 termination of this agreement for any reason, you will immediately stop all use of and delete and destroy all copies of the Licensed Content in your possession or under your control.
- 9. 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.
- 10. ENTIRE AGREEMENT.** This agreement, and any additional terms for the Trainer Content, updates and supplements are the entire agreement for the Licensed Content, updates and supplements.
- 11. 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.

- 12. 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.
- 13. DISCLAIMER OF WARRANTY. THE LICENSED CONTENT IS LICENSED "AS-IS" AND "AS AVAILABLE." YOU BEAR THE RISK OF USING IT. MICROSOFT AND ITS RESPECTIVE AFFILIATES GIVES NO EXPRESS WARRANTIES, GUARANTEES, OR CONDITIONS. YOU MAY HAVE ADDITIONAL CONSUMER RIGHTS UNDER YOUR LOCAL LAWS WHICH THIS AGREEMENT CANNOT CHANGE. TO THE EXTENT PERMITTED UNDER YOUR LOCAL LAWS, MICROSOFT AND ITS RESPECTIVE AFFILIATES EXCLUDES ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT.**
- 14. LIMITATION ON AND EXCLUSION OF REMEDIES AND DAMAGES. YOU CAN RECOVER FROM MICROSOFT, ITS RESPECTIVE AFFILIATES AND ITS SUPPLIERS ONLY DIRECT DAMAGES UP TO US\$5.00. YOU CANNOT RECOVER ANY OTHER DAMAGES, INCLUDING CONSEQUENTIAL, LOST PROFITS, SPECIAL, INDIRECT OR INCIDENTAL DAMAGES.**

This limitation applies to

- anything related to the Licensed Content, services, 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 July 2013

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 performance¹. 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

Microsoft | Learning

¹ IDC, Value of Certification: Team Certification and Organizational Performance, November 2006

Acknowledgments

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

Andrew J. Warren - Lead Subject Matter Expert

Andrew Warren has more than 25 years of experience in the IT industry, many of which he has spent teaching and writing. He has been involved as the Subject Matter Expert for many Windows Server® 2008 courses and the technical lead on a number of other courses. He also has been involved in developing TechNet sessions on Microsoft® Exchange Server® 2007. Based in the United Kingdom, he runs his own IT training and education consultancy.

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 a Microsoft Certified Trainer (MCT) since 1997.

Dave Franklyn - Subject Matter Expert

David M. Franklyn, MCT, Microsoft Certified Solutions Expert (MCSE), Microsoft Certified IT Professional (MCITP), Microsoft Most Valuable Professional (MVP) Windows Expert-It Pro, is a Senior Information Technology Trainer and Consultant at Auburn University in Montgomery, Alabama and the owner of DaveMCT, Inc. LLC. He is also Adjunct Faculty with MyITStudy.com. He is an Eastern USA Regional Lead MCT. Dave has been a Microsoft MVP since 2011 and has been teaching at Auburn University since 1998. Working with computers since 1976, Dave started out in the mainframe world and moved early into the networking arena. Before joining Auburn University, Dave spent 22 years in the US Air Force as an electronic communications and computer systems specialist, retiring in 1998. Dave is president of the Montgomery Windows IT Professional Group, and a guest speaker at many events involving Microsoft products.

Vladimir Meloski - Content Developer

Vladimir is a MCT, an MVP on Exchange Server, and consultant, providing unified communications and infrastructure solutions based on Exchange Server, Lync Server, Windows Server and Microsoft® System Center. Vladimir has 17 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 Microsoft Official Curriculum courses.

Stan Reimer - Subject Matter Expert

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® Domain Services (AD DS) 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 ten years, Stan has been writing courseware for Microsoft Learning, specializing in Active Directory and Exchange Server courses. Stan has been a MCT for 13 years.

Telmo Sampaio - Content Developer

Telmo Sampaio, who has a Bachelor of Science degree in Computer Science, also is an MCT, MCSE, Microsoft Certified Solutions Developer (MCSD), and an MCT Regional Lead. He is the “Chief Geek” for MCTrainer.NET and TechKnowLogical. Telmo specializes in System Center, Microsoft SharePoint®, Microsoft SQL Server®, and .NET, and has worked for IBM, Microsoft, and several start-ups during the past 20 years. He is very active in the MCT community, and travels the world providing consulting services and attending training engagements. His home base is Miami, Florida. Telmo has passed more than 100 Microsoft exams since his first certification in 1996.

David Susemihl - Content Developer

David Susemihl has worked as consultant, trainer, and courseware developer since 1996. David has extensive experience consulting on Microsoft Systems Management Server and Microsoft System Center Configuration Manager 2007, as well as Active Directory, Exchange Server, and Terminal Server/Citrix deployments. David has developed courseware development for Microsoft and Hewlett-Packard, and delivered those courses successfully in Europe, Central America, and across North America. For the last several years, David has been writing courseware for Microsoft Learning, and consulting on infrastructure transitions in Michigan.

Brian Svidergol - Subject Matter Expert

Brian Svidergol specializes in Microsoft infrastructure and cloud-based solutions built around Windows, Active Directory, Exchange Server, System Center, virtualization, and Microsoft Desktop Optimization Pack (MDOP). He holds a variety of Microsoft and industry certifications. Brian authored the *Active Directory Cookbook*, 4th edition. He has also worked as a Subject Matter Expert and Technical Reviewer on many Microsoft Official Curriculum courses, Microsoft certification exams, and authored or reviewed related training content.

Orin Thomas - Subject Matter Expert

Orin Thomas is an MVP, an MCT, and has a variety of 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 1990's. He regularly speaks at events such as TechED in Australia, and around the world on Windows Server, Windows client operating systems, System Center, and security topics. Orin founded and runs the Melbourne System Center Users Group.

Brian Langan - Technical Reviewer

Brian Langan is president and founder of Langan Enterprises Inc., a consulting/training/security firm established in 1995. He has worked in the industry for over 20 years becoming an MCT in 1996, and holds a number of Microsoft certifications on clients, servers, messaging and System Center products. He has written courses on many different topics including Windows Troubleshooting and Security courses for Global Knowledge and other training companies.

Contents

Module 1: Deploying and Managing Windows Server 2012 R2

Lesson 1: Windows Server 2012 R2 Overview	1-2
Lesson 2: Installing Windows Server 2012 R2	1-13
Lesson 3: Post-Installation Configuration of Windows Server 2012 R2	1-20
Lesson 4: Overview of Windows Server 2012 R2 Management	1-31
Lesson 5: Introduction to Windows PowerShell	1-38
Lab: Deploying and Managing Windows Server 2012 R2	1-46

Module 2: Introduction to Active Directory Domain Services

Lesson 1: Overview of AD DS	2-2
Lesson 2: Overview of Domain Controllers	2-12
Lesson 3: Installing a Domain Controller	2-18
Lab: Installing Domain Controllers	2-26

Module 3: Managing Active Directory Domain Services Objects

Lesson 1: Managing User Accounts	3-2
Lesson 2: Managing Groups	3-11
Lesson 3: Managing Computer Accounts	3-20
Lesson 4: Delegating Administration	3-27
Lab: Managing Active Directory Domain Services Objects	3-31

Module 4: Automating Active Directory Domain Services Administration

Lesson 1: Using Command-line Tools for AD DS Administration	4-2
Lesson 2: Using Windows PowerShell for AD DS Administration	4-8
Lesson 3: Performing Bulk Operations with Windows PowerShell	4-14
Lab: Automating AD DS Administration by Using Windows PowerShell	4-25

Module 5: Implementing IPv4

Lesson 1: Overview of TCP/IP	5-2
Lesson 2: Understanding IPv4 Addressing	5-6
Lesson 3: Subnetting and Supernetting	5-11
Lesson 4: Configuring and Troubleshooting IPv4	5-16
Lab: Implementing IPv4	5-26

Module 6: Implementing Dynamic Host Configuration Protocol

Lesson 1: Installing a DHCP Server Role	6-2
Lesson 2: Configuring DHCP Scopes	6-9
Lesson 3: Managing a DHCP Database	6-15
Lesson 4: Securing and Monitoring DHCP	6-19
Lab: Implementing DHCP	6-24

Module 7: Implementing DNS

Lesson 1: Name Resolution for Windows Clients and Servers	7-2
Lesson 2: Installing a DNS Server	7-15
Lesson 3: Managing DNS Zones	7-21
Lab: Implementing DNS	7-25

Module 8: Implementing IPv6

Lesson 1: Overview of IPv6	8-2
Lesson 2: IPv6 Addressing	8-7
Lesson 3: Coexistence with IPv4	8-13
Lesson 4: IPv6 Transition Technologies	8-17
Lab: Implementing IPv6	8-22

Module 9: Implementing Local Storage

Lesson 1: Overview of Storage	9-2
Lesson 2: Managing Disks and Volumes	9-12
Lesson 3: Implementing Storage Spaces	9-23
Lab: Implementing Local Storage	9-29

Module 10: Implementing File and Print Services

Lesson 1: Securing Files and Folders	10-2
Lesson 2: Protecting Shared Files and Folders by Using Shadow Copies	10-15
Lesson 3: Configuring Work Folders	10-18
Lesson 4: Configuring Network Printing	10-25
Lab: Implementing File and Print Services	10-31

Module 11: Implementing Group Policy

Lesson 1: Overview of Group Policy	11-2
Lesson 2: Group Policy Processing	11-10
Lesson 3: Implementing a Central Store for Administrative Templates	11-16
Lab: Implementing Group Policy	11-20

Module 12: Securing Windows Servers Using Group Policy Objects

Lesson 1: Windows Operating Systems Security Overview	12-2
Lesson 2: Configuring Security Settings	12-6
Lab A: Increasing Security for Server Resources	12-17
Lesson 3: Restricting Software	12-24
Lesson 4: Configuring Windows Firewall with Advanced Security	12-28
Lab B: Configuring AppLocker and Windows Firewall	12-33

Module 13: Implementing Server Virtualization with Hyper-V

Lesson 1: Overview of Virtualization Technologies	13-2
Lesson 2: Implementing Hyper-V	13-8
Lesson 3: Managing Virtual Machine Storage	13-17
Lesson 4: Managing Virtual Networks	13-26
Lab: Implementing Server Virtualization with Hyper-V	13-32

Lab Answer Keys

Module 1 Lab: Deploying and Managing Windows Server 2012 R2	L1-1
Module 2 Lab: Installing Domain Controllers	L2-9
Module 3 Lab: Managing Active Directory Domain Services Objects	L3-13
Module 4 Lab: Automating AD DS Administration by Using Windows PowerShell	L4-21
Module 5 Lab: Implementing IPv4	L5-25
Module 6 Lab: Implementing DHCP	L6-29
Module 7 Lab: Implementing DNS	L7-35
Module 8 Lab: Implementing IPv6	L8-43
Module 9 Lab: Implementing Local Storage	L9-47
Module 10 Lab: Implementing File and Print Services	L10-53
Module 11 Lab: Implementing Group Policy	L11-61
Module 12 Lab A: Increasing Security for Server Resources	L12-65
Module 12 Lab B: Configuring AppLocker and Windows Firewall	L12-71
Module 13 Lab: Implementing Server Virtualization with Hyper-V	L13-77

About This Course

This section provides you with a brief description of the course—*20410C: Installing and Configuring Windows Server® 2012*—audience, suggested prerequisites, and course objectives.

Course Description

Note: This release (C) Microsoft® Official Curriculum (MOC) version of course 20410 has been developed on the final release version of Windows server 2012 R2 software.

This course is part one of a three-part series that provides the skills and knowledge necessary to implement a core Windows Server 2012 infrastructure in an existing enterprise environment. The three courses collectively cover implementing, managing, maintaining, and provisioning services and infrastructure in a Windows Server 2012 environment. While there is some cross-over in skills and tasks across the courses, this course focuses on the initial implementation and configuration of core services including Active Directory® Domain Services (AD DS), networking services, and Microsoft® Hyper-V® Server 2012 R2 configuration.

Audience

This course is intended for Information Technology (IT) Professionals who have some knowledge and experience working with Windows operating systems and who want to acquire the skills and knowledge necessary to install and perform the initial configuration of a Windows Server 2012 or Windows Server 2012 R2 server in an existing Windows server environment. Candidates typically interested in attending this course are:

- Windows Server Administrators who are relatively new to Windows Server administration and related technologies and who are looking to learn more about Windows Server 2012 or Windows Server 2012 R2.
- IT Professionals who are experienced in other non-Microsoft technologies, who meet the pre-requisites and want to cross-train on Windows Server 2012.
- IT Professionals who want to take the exam, "70-410: Installing and Configuring Windows Server 2012" exam.
- IT professional wishing to take the Microsoft Certified Solutions Expert (MCSE) exams in Datacenter, Desktop Infrastructure, Messaging, Collaboration and Communications, who will also be interested in taking this course as they prepare for the Microsoft Certified Solutions Associate (MCSA) exams, which are a pre-requisite for their individual specialties.

Student Prerequisites

This course requires that students meet the following prerequisites:

- Have an understanding of networking fundamentals
- Understand basic Active Directory concepts
- Have an awareness and understanding of security best practices
- Have basic knowledge of server hardware
- Experience working with and configuring Windows client operating systems such as Windows 7 or Windows 8

Students would also benefit from having some previous Windows Server operating system experience.

Course Objectives

After completing this course, students will be able to:

- Install and Configure Windows Server 2012.
- Describe AD DS.
- Manage Active Directory objects.
- Automate Active Directory administration.
- Implement IPv4.
- Implement Dynamic Host Configuration Protocol (DHCP).
- Implement Domain Name System (DNS).
- Implement IPv6.
- Implement local storage.
- Share files and printers.
- Implement Group Policy.
- Use Group Policy Objects (GPOs) to secure Windows Servers.
- Implement server virtualization using Hyper-V.

Course Outline

This section provides an outline of the course:

Module 1, Deploying and Managing Windows Server 2012

This module starts the course by discussing installation of Windows Server 2012. This is not the most commonly performed task that is discussed in the course, but it provides a logical starting point for students to begin working with Windows Server 2012 or Windows Server 2012 R2.

Module 2, Introduction to Active Directory Domain Services

AD DS is a core part of network management in an enterprise environment. It is introduced early in the course so that students will be able to use it to perform other tasks—such as creating users and groups—in later modules. In this module, students will install a domain controller.

Module 3, Managing Active Directory Domain Services Objects

This module discusses creating and managing specific Active Directory objects such as users, groups, and computer accounts. This is a core part of what a beginning server administrator does on a daily basis. Some of these tasks are also delegated to help desk staff.

Module 4, Automating Active Directory Domain Services Administration

This module expands on the knowledge gained in Module 3 by providing students with methods for automating the creation and management of Active Directory objects. This is a relatively advanced topic, but logically flows after Module 3.

Module 5, Implementing IPv4

This module begins a new thread of learning in the course. Configuring and understanding IPv4 is fundamental to working as a system administrator.

Module 6, Implementing Dynamic Host Configuration Protocol

This module discusses how to use DHCP to distribute IPv4 address information.

Module 7, Implementing Domain Name System

This module includes how DNS converts names to IP addresses and why this is important in an Active Directory environment.

Module 8, Implementing IPv6

This module introduces IPv6 configuration, which is likely to be new content for the students. Module 8 is separated from Module 5 because both are theory-intensive and may overwhelm students if they are sequential. IPv6 knowledge is not required for Modules 6 and 7.

Module 9, Implementing Local Storage

This module includes content on storage configuration for Windows Server 2012. This is prerequisite information for Module 10, which discusses creating and securing file shares.

Module 10, Implementing File and Print Services

This module discusses file shares and printing at the same time because both are commonly used network services. Security for file shares and printing uses the knowledge about user accounts and groups covered in Modules 2 and 3.

Module 11, Implementing Group Policy

This module builds on the information students have already learned about AD DS to introduce the creation and management of GPOs.

Module 12, Securing Windows Servers Using Group Policy Objects

This module then covers specific Group Policy settings that can be used to increase security. The settings include security policies, application restriction policies, and Windows Firewall rules.

Module 13, Implementing Server Virtualization with Hyper-V

The final module discusses how to configure Hyper-V and how to create virtual machines. This module is last because the lab has the potential to negatively impact the virtual machines already deployed on the student machines.

Exam/Course Mapping

This course, *20410C: Installing and Configuring Windows Server® 2012*, has a direct mapping of its content to the objective domain for the Microsoft exam *70-410: Installing and Configuring Windows Server 2012*.

The table below 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 utilize 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-70-410.aspx>, under Skills Measured.

Exam Objective Domain: Exam 70-410: Installing and Configuring Windows Server 2012		Course Content		
1. Install and Configure Servers (17%)		Module	Lesson	Lab
1.1. 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
1.2. 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; install and configure PowerShell Desired State Configuration (DSC)	Mod 1	Lesson 2/3/4	Mod 1 Ex 1/2/3/4
1.3. 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; configure storage pools and disk pools; create storage pools by using disk enclosures	Mod 9	Lesson 2/3	Mod 9 Lab Ex 1/2/3
2. Configure Server Roles and Features (17%)				
2.1. Configure file and share access.	This objective may include but is not limited to: Create and configure shares; configure share permissions; configure offline files; configure NTFS permissions; configure access-based enumeration (ABE); configure Volume Shadow Copy Service (VSS); configure NTFS quotas; create and configure Work Folders	Mod 10	Lesson 1/2/3	Mod 10 Lab Ex 1/2/3
2.2. Configure print and document services.	This objective may include but is not limited to: Configure the Easy Print print driver; configure Enterprise Print Management; configure drivers; configure printer pooling; configure print priorities; configure printer permissions	Mod 10	Lesson 4	Mod 10 Lab Ex 4

Exam Objective Domain: Exam 70-410: Installing and Configuring Windows Server 2012		Course Content		
2.3. 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; manage non-domain joined servers	Mod 1	Lesson 3/4/5	Mod 1 Lab Ex 4
		Mod 12	Lesson 4	Mod 12 Lab B Ex 2
3. Configure Hyper-V (18%)				
3.1. 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; create and configure Generation 1 and 2 virtual machines; configure and use extended session mode; Configure RemoteFX	Mod 13	Lesson 2	Mod 13 Ex 3
3.2. 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 checkpoints; implement a virtual Fibre Channel adapter; configure storage Quality of Service	Mod 13	Lesson 2/3	Mod 13 Ex 3/4
3.3. Create and configure virtual networks.	This objective may include but is not limited to: configure Hyper-V virtual switches; optimize network performance; configure MAC addresses; configure network isolation; configure synthetic and legacy virtual network adapters; configure NIC teaming in virtual machines	Mod 13	Lesson 4	Mod 13 Lab Ex 2
4. Deploy and Configure Core Network Services (17%)				
4.1. Configure IPv4 and IPv6 addressing.	This objective may include but is not limited to: Configure IP address options; configure IPv4 or IPv6 subnetting; configure supernetting; configure interoperability between IPv4 and IPv6; configure ISATAP; configure Teredo	Mod 1	Lesson 3	Mod 1 Lab Ex 1/2
		Mod 5	Lesson 2/3/4	Mod 5 Lab Ex 1/2
		Mod 8	Lesson 3/4	Mod 8 Lab Ex 1/2
4.2. Deploy and configure Dynamic Host Configuration Protocol (DHCP) service.	This objective may include but is not limited to: Create and configure scopes; configure a DHCP reservation; configure DHCP options; configure client and server for PXE boot; configure DHCP relay agent; authorize DHCP server	Mod 6	Lesson 1/2/3/4	Mod 6 Ex 1/2
4.3. Deploy and configure DNS service.	This objective may include but is not limited to: Configure Active Directory integration of primary zones; configure forwarders; configure Root Hints; manage DNS cache; create A and PTR resource records	Mod 7	Lesson 1/2/3	Mod 7 Ex 1/2/3

Exam Objective Domain: Exam 70-410: Installing and Configuring Windows Server 2012		Course Content		
5. Install and Administer Active Directory (14%)				
5.1. 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; deploy Active Directory IaaS in Azure	Mod 2	Lesson 1/2/3	Mod 2 Lab Ex 1/2
5.2. Create and manage Active Directory users and computers.	This objective may include but is not limited to: Automate the creation of Active Directory accounts; create, copy, configure, and delete users and computers; configure templates; perform bulk Active Directory operations; configure user rights; offline domain join; manage inactive and disabled accounts	Mod 3	Lesson 1/2/3	Mod 3 Lab Ex 2/3
		Mod 4	Lesson 1/2/3	Mod 4 Lab Ex 1/2/3
5.3. Create and manage Active Directory groups and organizational units (OUs).	This objective may include but is not limited to: Configure group nesting; convert groups including security, distribution, universal, domain local, and domain global; manage group membership using Group Policy; enumerate group membership; delegate the creation and management of Active Directory objects; manage default Active Directory containers; create, copy, configure, and delete groups and OUs	Mod 3	Lesson 2/4	Mod 3 Lab Ex 1/2/3
		Mod 4	Lesson 1/2	Mod 4 Lab Ex 1/2/3
6. Create and Manage Group Policy (16%)				
6.1. Create Group Policy objects (GPOs).	This objective may include but is not limited to: Configure a Central Store; manage starter GPOs; configure GPO links; configure multiple local group policies	Mod 11	Lesson 1/2/3	Mod 11 Lab Ex 1/2
6.2. Configure security policies.	This objective may include but is not limited to: Configure User Rights Assignment; configure Security Options settings; configure Security templates; configure Audit Policy; configure Local Users and Groups; configure User Account Control (UAC)	Mod 12	Lesson 2	Mod 12 Lab A Ex 1/2/3
6.3. Configure application restriction policies.	This objective may include but is not limited to: Configure rule enforcement; configure AppLocker rules; configure Software Restriction Policies	Mod 12	Lesson 3	Mod 12 Lab B Ex 1
6.4. Configure Windows Firewall.	This objective may include but is not limited to: Configure rules for multiple profiles using Group Policy; configure connection security rules; configure Windows Firewall to allow or deny applications, scopes, ports, and users; configure authenticated firewall exceptions; import and export settings	Mod 12	Lesson 4	Mod 12 Lab B Ex 2



Note: 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:

- Real-world, hands-on experience Installing and configuring a Windows Server 2012 Infrastructure
- Windows 7 or Windows 8 client configuration experience
- 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-70-410.aspx>, under Preparation options.

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/course.aspx?ID=20410C>, under Overview, Audience Profile.

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.

You may be accessing either a printed course hand book or digital courseware material via the Arvato Skillpipe reader. Your Microsoft Certified Trainer will provide specific details but both contain the following:

- **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 is needed.



Course Companion Content on the <http://www.microsoft.com/learning/companionmoc> Site: Searchable, easy-to-navigate digital content with integrated premium online resources designed to supplement the Course Handbook.

- **Modules:** Include companion content, such as questions and answers, detailed demo steps and additional reading links, for each lesson. Additionally, they include Lab Review questions and answers and Module Reviews and Takeaways sections, which contain the review questions and answers, best practices, common issues and troubleshooting tips with answers, and real-world issues and scenarios with answers.

- **Resources:** Include well-categorized additional resources that give you immediate access to the most up-to-date premium content on TechNet, Microsoft Developer Network®, and Microsoft Press®.



Student Course files on the <http://www.microsoft.com/learning/companionmoc> Site

- **Course evaluation** At the end of the course, you will 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 e-mail to support@microsoft.com. To inquire about the Microsoft Certification Program, send e-mail to 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 virtual machines built in 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. The following table shows the role of each virtual machine used in this course.

Virtual machine	Role
20410C-LON-DC1	A domain controller running Windows Server 2012 R2 in the Adatum.com domain.
20410C-LON-SVR1	A member server running Windows Server 2012 R2 in the Adatum.com domain.
20410C-LON-SVR2	A member server running Windows Server 2012 R2 in the Adatum.com domain. This server will be located on a second subnet.
20410C-LON-SVR3	A blank virtual machine on which students will install Windows Server 2012 R2.
20410C-LON-HOST1	A bootable VHD for running Windows Server 2012 R2 as the host for Hyper-V.
20410C-LON-CORE	A standalone server running Windows Server 2012 R2 Server Core.
20410C-LON-RTR	A router that is used for network activities that require a separate subnet. Also running Windows Server 2012 R2.
20410C-LON-CL1	A client computer running Windows 8.1 and Microsoft® Office 2013 in the Adatum.com domain.
20410C-LON-CL2	A client computer running Windows 8.1 and Office 2013 in the Adatum.com domain that is located in a second subnet.

Software Configuration

The following software has been installed in the course:

- Microsoft Message Analyzer is installed on LON-SVR1.

Classroom Setup

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

You may be accessing the lab virtual machines in either in a hosted online environment with a web browser or by using Hyper-V on a local machine. The labs and virtual machines are the same in both scenarios however there may be some slight variations because of hosting requirements. Any discrepancies will be called out in the Lab Notes on the hosted lab platform.

You Microsoft Certified Trainer will provide details about your specific lab environment.

Course Hardware Level

Where labs are being run locally, 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 7 with 16 gigabytes (GB) of random access memory (RAM)

Navigation in Windows Server 2012

If you are not familiar with the user interface in Windows Server 2012 R2 or Windows 8.1, then the following information will help orient you to the new interface.

- *Sign in* and *Sign out* replace *Log in* and *Log off*.
- Administrative tools are found in the **Tools** menu of Server Manager.
- Get to the Start screen, Settings, and Search as follows:
 - To get to the Start screen, in the lower-left corner of the screen, click the **Start** button. This provides access to some applications.
 - To get to Settings, point your mouse to the lower-right corner of the screen, and then click the **Settings** charm when it appears. Settings include Control Panel and Power.
 - To get to Search, point your mouse to the lower-right corner of the screen, and then click the **Search** charm when it appears. This allows you to search applications, settings, and files.

You may also find the following shortcut keys useful:

- Windows logo key: Opens the **Start** screen
- Windows logo key +I: Opens Settings
- Windows logo key +R: Opens Run
- Windows logo key +C: Displays the selection of charms

Module 1

Deploying and Managing Windows Server 2012

Contents:

Module Overview	1-1
Lesson 1: Windows Server 2012 R2 Overview	1-2
Lesson 2: Installing Windows Server 2012 R2	1-13
Lesson 3: Post-Installation Configuration of Windows Server 2012 R2	1-20
Lesson 4: Overview of Windows Server 2012 R2 Management	1-31
Lesson 5: Introduction to Windows PowerShell	1-38
Lab: Deploying and Managing Windows Server 2012 R2	1-46
Module Review and Takeaways	1-54

Module Overview

Understanding the capabilities of a new Windows Server® 2012 R2 operating system enables you to leverage that operating system effectively. If you do not understand the capabilities of your new Windows Server 2012 R2 operating system, you might end up using it the same way that you used the previous operating system, which would forego the advantages of the new system. By understanding how to utilize your new Windows Server 2012 R2 operating system fully, and by understanding the tools that are available to manage that functionality, you can provide your organization with more value.

This module introduces the new Windows Server 2012 R2 administrative interface. In this module, you will learn about the different roles and features that are available with the Windows Server 2012 R2 operating system. You will also learn about the different installation options that you can use when you install Windows Server 2012 R2.

This module discusses the configuration steps that you can perform both during installation and after deployment to ensure that the servers can begin functioning in its assigned role. You will also learn how to use Windows PowerShell® to perform common administrative tasks in Windows Server 2012 R2.

Objectives

After completing this module, you will be able to:

- Describe Windows Server 2012 R2.
- Install Windows Server 2012 R2.
- Perform post-installation configuration of Windows Server 2012 R2.
- Describe the management tools available in Windows Server 2012 R2.
- Perform basic administrative tasks using Windows PowerShell.

Lesson 1

Windows Server 2012 R2 Overview

Before deploying Windows Server 2012 R2, you need to understand how each of the Windows Server 2012 R2 editions might benefit your organization's servers. You also need to know whether a particular hardware configuration is appropriate for Windows Server 2012 R2, whether a virtual deployment might be more suitable than a physical deployment, and which installation source allows you to deploy Windows Server 2012 R2 in an efficient manner. If you do not have an understanding of these issues, you could end up costing your organization time and money by making a choice that you must later correct.

This lesson provides an overview of the various Windows Server 2012 R2 editions, installation options, roles, and features. Using this information, you will be able to determine which Windows Server 2012 R2 edition and installation options are right for your organization.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the different Windows Server 2012 R2 editions.
- Describe the purpose and function of a Windows Server 2012 R2 Server Core installation.
- Explain the function of Windows Server 2012 R2 roles.
- Explain the purpose of various Windows Server 2012 R2 features.

Windows Server 2012 R2 Editions

There are several different editions of Windows Server 2012 R2 from which to choose. These editions allow organizations to select a version of Windows Server 2012 R2 that best meets their needs, rather than pay for features they do not require.

When deploying a server for a specific role, systems administrators can save substantially by selecting the appropriate edition.

The following table lists the Windows Server 2012 R2 editions.


Windows Server 2012 editions:


- Windows Server 2012 Standard
- Windows Server 2012 Datacenter
- Windows Server 2012 Foundation
- 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
The Windows Server 2012 R2 Standard operating system	Provides all the roles and features available on the Windows Server 2012 R2 platform. Supports up to 64 sockets and up to 4 terabytes (TB) of random access memory (RAM). Includes two virtual machine licenses.
The Windows Server 2012 R2 Datacenter operating system	Provides all the roles and features that are available on the Windows Server 2012 R2 platform. Includes unlimited virtual machine licenses for virtual machines that are run on the same hardware. Supports 64 sockets, up to 640 processor cores, and up to 4 TB of RAM.

(continued from previous page)

Edition	Description
The Windows Server 2012 R2 Foundation operating system	Designed for small business owners, it allows only 15 users, cannot be joined to a domain, and includes limited server roles. Supports one processor core and up to 32 gigabytes (GB) of RAM.
The Windows Server 2012 R2 Essentials operating system	<p>Next edition of Small Business Server. It is now available in two forms:</p> <ul style="list-style-type: none"> • As an installable server role in an existing domain. • As a core Windows Server edition on a virtual machine (using a wizard). <p>It cannot function as a Hyper-V®, Failover Clustering, Server Core, or Remote Desktop Services server. It has limits for 25 users and 50 devices. Supports two processor cores and 64 GB of RAM.</p> <p>There are many new features and improvements for Windows Server 2012 R2 Essentials R2, including client deployment, user management, storage and data protection, and Office 365 integration.</p>
Microsoft Hyper-V Server 2012	Standalone Hyper-V platform for virtual machines with no UI. No licensing cost (free) for host operating system, but virtual machines are licensed normally. Supports 64 sockets and 4 TB of RAM. Supports domain join. Does not support other Windows Server 2012 R2 roles, other than limited file services features.
The Windows Storage Server®2012 Workgroup operating system	Entry-level unified storage appliance. Limited to 50 users, one processor core, 32 GB of RAM. Supports domain join.
The Windows Storage Server 2012 Standard operating system	Supports 64 sockets, but is licensed on a two-socket, incrementing basis. Supports 4 TB of RAM. Includes two virtual machine licenses. Supports domain join. 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).
The Windows MultiPoint Server 2012 Standard operating system	Supports multiple users accessing the same host computer directly using separate mouse, keyboard, and monitors. Limited to one socket, 32 GB of RAM, and a maximum of 12 sessions. Supports some roles, including DNS and DHCP Server roles, but does not support others, including AD DS, AD CS, and AD FS. Does not support domain join.
The Windows MultiPoint Server 2012 Premium operating system	Supports multiple users accessing the same host computer directly using separate mouse, keyboard, and monitors. Limited to two sockets, 4 TB of RAM, and a maximum of 22 sessions. Supports some roles, including DNS and DHCP Server roles, but does not support others, including AD DS, AD CS, and AD FS. Supports domain join.

 **Reference Links:** For detailed information on the new features in Windows Server 2012 R2 Essentials, see "What's New in Windows Server 2012 R2 Essentials" at <http://go.microsoft.com/fwlink/?LinkID=331071>.

 **Additional Reading:** For more information about the differences between Windows Server 2012 R2 editions, download the Windows Server 2012 R2 Products and Editions Comparison chart at <http://go.microsoft.com/fwlink/?LinkID=331070>.



Reference Links: Many features have been removed or deprecated in Windows Server 2012 R2. For more information, go to: Features Removed or Deprecated in Windows Server 2012 R2 Preview at <http://go.microsoft.com/fwlink/?LinkID=331069>.

What Is Server Core?

Server Core is an installation option for Windows Server 2012 R2 that can contain variations of the Graphical User Interface (GUI) depending on the requirements of the server roles to be installed. You can manage Server Core locally by using Windows PowerShell or a command-line interface, rather than by using GUI-based tools, or remotely by using one of the remote management options. Remote management is covered later in this module.

A Windows Server 2012 R2 Server Core installation offers fewer components and administrative management options than the full installation of Windows Server 2012 R2.

Server Core installation is the default installation option when installing Windows Server 2012 R2. Server Core has the following advantages over a traditional Windows Server 2012 R2 deployment:

- Reduced update requirements. Because Server Core installs fewer components, its deployment requires you to install fewer software updates. This reduces the number of monthly reboots required and the amount of time required for an administrator to service Server Core.
- Reduced hardware footprint. Server Core computers require less RAM and less hard disk space. When virtualized, this means that you can deploy more servers on the same host.

Increasing numbers of Microsoft server programs are designed to run on computers with Server Core–installed operating systems. For example, you can install SQL Server 2012 on computers that are running the Server Core–installed version of Windows Server 2012 R2.

You can switch from Server Core to the graphical version of Windows Server 2012 R2 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 R2 installation files:

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

You can also use Windows Update or the installation DVD as the installation file source. Installing the graphical components gives you the option of performing administrative tasks using the graphical tools.

Once 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 R2 to Server Core by removing the following components of the User Interfaces and Infrastructure feature:

- Graphical Management Tools and Infrastructure. This contains a minimal server interface to provide some server management user interface tools such as Server Manager and Administrative Tools.
- Server Graphical Shell. This contains the full GUI including Internet Explorer and File Explorer and other user interface components. This has a larger footprint than the Graphical Management Tools and Infrastructure option.

Server Core:

- Is a more secure, less resource-intensive installation option
 - Can be converted to full graphical shell version of Windows Server 2012
 - Is the default installation option for Windows Server 2012
 - Is managed locally using sconfig.cmd
- Note that if you enable remote management, you rarely will need to sign in locally



Note: Be careful when removing graphical features, because servers might have other components installed that are dependent upon those features.

When connected locally, you can use the tools that are listed in the following table to manage Server Core deployments of Windows Server 2012 R2.

Tool	Function
Cmd.exe	Allows you to run traditional command-line tools such as ping.exe, ipconfig.exe, and netsh.exe.
PowerShell.exe	Launches a Windows PowerShell session on the Server Core deployment. You can then perform Windows PowerShell tasks normally. Windows Server 2012 R2 comes with Windows PowerShell version 4.0 installed.
Sconfig.cmd	A command-line menu-driven administrative tool that allows you to perform most common server administrative tasks.
Notepad.exe	Allows you to use the Notepad.exe text editor within the Server Core environment.
Regedt32.exe	Provides registry access within the Server Core environment.
Msiinfo32.exe	Allows you to view system information about the Server Core deployment.
Taskmgr.exe	Launches the Task Manager.
SCregEdit.wsf	Used to enable Remote Desktop on the Server Core deployment.



Note: If you accidentally close the command window on a computer that is running Server Core, you can recover the command window by performing the following steps:

1. Press the Ctrl+Alt+Del keys, and then click **Task Manager**.
2. From the **File** menu, click **New Task (Run...)**, and then type **cmd.exe**.

Server Core supports most Windows Server 2012 R2 roles and features. However, you cannot install the following roles on a computer running Server Core:

- AD FS
- Application Server
- Network Policy and Access Services (NPAS)
- Windows Deployment Services

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



Note: You can check which roles on Server Core are available and which are not by running the query **Get-WindowsFeature | where-object {\$_.InstallState -eq "Removed"}**.

You can use the following tools to remotely manage a computer that is running the Server Core installation option:

- **Server Manager.** You can add a server that is running Server Core to Server Manager that is on a server which is running a full installation of Windows. You can then use Server Manager to manage the server roles running on the Server Core computer. You can configure Remote Desktop by using `Sconfig.cmd`.
- **Remote Windows PowerShell.** You can use Remote Windows PowerShell to run Windows PowerShell commands or scripts against correctly-configured remote servers if the script is hosted on the local server. With Remote Windows PowerShell, you also can locally load Windows PowerShell modules, such as Server Manager, and execute the cmdlets available in that module against appropriately-configured remote servers.
- **Remote Desktop.** You can connect to a computer that is running the Server Core installation option by using Remote Desktop. Configure Remote Desktop by using `Sconfig.cmd`.
- **Remote Management Consoles.** For most server roles, you can add a computer that is running the Server Core installation option to a management console that is running on another computer.

Windows Server 2012 R2 Roles

To correctly plan how you are going to use Windows Server 2012 R2 to support your organization's requirements, you need to be fully aware of what roles are available as part of the operating system. Each version of Windows Server comes with a different set of roles. As new versions of Windows Server are released, some roles are enhanced and others are deprecated. For the most part, the roles that are available in Windows Server 2012 R2 are familiar to IT professionals that have managed Windows Server 2008 and Windows Server 2003.

Functions

- Web Server
- Domain Controller
- Certificate Server

Roles

- Roles are made up of role services components that provide additional functionality associated with the role
- In Server Manager 2012, console servers with a similar role are grouped together
- Role deployment also includes the configuration of dependencies

Windows Server 2012 R2 supports the server roles that are listed in the following table.

Role	Function	Changes in Windows Server 2012 R2
AD CS	Allows you to deploy certification authorities and related role services.	
AD DS	A centralized store of information about network objects, including user and computer accounts. Used for authentication and authorization.	Windows Server 2003 domain and functional levels of AD DS and the File Replication Service have been deprecated in Windows Server 2012 R2
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.	

(continued from previous page)

Role	Function	Changes in Windows Server 2012 R2
Active Directory Rights Management Services (AD RMS)	Allows you to apply rights management policies to prevent unauthorized access to sensitive documents.	
Application Server	Supports centralized management and hosting of high-performance distributed business applications, such as those built with Microsoft .NET Framework 4.5.	Deprecated in Windows Server 2012 R2
DHCP Server	Provisions client computers on the network with temporary IP addresses.	
DNS Server	Provides name resolution for TCP/IP networks.	
Fax Server	Supports sending and receiving of faxes. Also allows you to manage fax resource on the network.	
File and Storage Services	Supports the management of shared folders storage, distributed file system (DFS), and network storage.	
Hyper-V	Enables you to host virtual machines on computers that are running Windows Server 2012.	
Network Access Protection (NAP)	A mechanism to create and enforce policies that describe software and security update requirements before the requesting computer is allowed to access the LAN. A computer that is not in compliance can be provided with ways to remediate its configuration to bring it into compliance.	Deprecated in Windows Server 2012 R2
Print and Document Services	Supports centralized management of document tasks, including network scanners and networked printers.	
Remote Access	Supports Seamless Connectivity, Always On, and Always Managed features based on the Windows 7 DirectAccess feature. Also supports remote access through virtual private network (VPN) and dial-up connections.	
Remote Desktop Services (RDS)	Supports access to virtual desktops, session-based desktops, and RemoteApp programs.	

(continued from previous page)

Role	Function	Changes in Windows Server 2012 R2
Volume Activation Services	Allows you to automate and simplify the management of volume license keys and volume key activation. Allows you to manage a Key Management Service (KMS) 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.	Internet Information Service (IIS) 6.0 Manager has been deprecated in Windows Server 2012 R2
Windows Deployment Services	Allows you to deploy server operating systems to clients over the network.	Windows PowerShell cmdlets have been added and cmdlet scripting is supported in Windows Server 2012 R2
Windows Server Update Services (WSUS)	Provides a method of deploying updates for Microsoft products to network computers.	

When you deploy a role, Windows Server 2012 R2 automatically configures aspects of the server's configuration (such as firewall settings), to support the role. Windows Server 2012 R2 also automatically and simultaneously deploys role dependencies. For example, when you install the WSUS role, the Web Server (IIS) role components that are required to support the WSUS role are also installed automatically.

You add and remove roles using the Add Roles and Features Wizard, which is available from the Windows Server 2012 R2 Server Manager console. If you are using Server Core, then you can also add and remove roles using the **Install-WindowsFeature** and **Remove-WindowsFeature** Windows PowerShell cmdlets.

Question: Which roles are often co-located on the same server?

What Are the Windows Server 2012 R2 Features?

Windows Server 2012 R2 *features* are independent components that often support role services or support the server directly. For example, Windows Server Backup is a feature because it only provides backup support for the local server. It is not a resource that other servers on the network can use.

Features:

- Are components that support the server such as Windows Server Backup or Failover clustering
- Usually do not provide a service directly to clients on the network

Keep in mind the following points:

- Roles can have features as dependencies
- Features on Demand are features that need to be installed using a mounted image as a source

Windows Server 2012 R2 includes the features that are listed in the following table.

Feature	Description	Changes in Windows Server 2012 R2
.NET Framework 3.5 Features	Installs .NET Framework 3.5 technologies.	
.NET Framework 4.5 Features	Installs .NET Framework 4.5 technologies. This feature is installed by default.	
Background Intelligent Transfer Service (BITS)	Allows asynchronous transfer of files to ensure that other network applications are not adversely impacted.	
Windows BitLocker® Drive Encryption	Supports full-disk and full-volume encryption, and startup environment protection.	
BitLocker network unlock	Provides a network-based key protector that can unlock locked BitLocker-protected domain-joined operating systems.	
Windows BranchCache®	Allows the server to function as either a hosted cache server or a BranchCache content server for BranchCache clients.	
Client for NFS	Provides access to files stored on network file system (NFS) servers.	
Data Center Bridging	Allows you to enforce bandwidth allocation on Converged Network Adapters.	
Enhanced Storage	Provides support for additional functionality available in Enhanced Storage Access (IEEE 1667 protocol) device, including data access restrictions.	
Failover Clustering	A high availability feature that allows Windows Server 2012 R2 to participate in failover clustering.	
Group Policy Management	An administrative management tool for administering Group Policy across an enterprise.	
Ink and Handwriting Services	Allows use of Ink Support and Handwriting Recognition.	
Internet Printing Client	Supports use of Internet Printing Protocol.	
IP Address Management (IPAM) Server	Centralized management of IP address and namespace infrastructure.	
Internet SCSI (iSCSI) Target Storage Provider	Provides iSCSI target and disk management services to Windows Server 2012 R2.	
Internet Storage name Service (iSNS) Server service	Supports discovery services of iSCSI storage area networks (SANs).	

(continued from previous page)

Feature	Description	Changes in Windows Server 2012 R2
Line Printer Remote (LPR) Port Monitor	Allows computer to send print jobs to printers that are shared using the Line Printer Daemon (LPD) service.	Deprecated in Windows Server 2012 R2
Management Open Data Protocol (OData) IIS Extension	Allows you to expose Windows PowerShell cmdlets through an OData-based web service running on the Internet Information Services (IIS) platform.	
Media Foundation	Supports media file infrastructure.	
Message Queuing	Supports message delivery between applications.	
Multipath input/output (I/O)	Supports multiple data paths to storage devices.	
Network Load Balancing (NLB)	Allows traffic to be distributed in a load-balanced manner across multiple servers that host the same stateless applications.	
Peer Name Resolution Protocol (PNRP)	Name resolution protocol that allows applications to resolve names on the computer.	
Quality Windows Audio Video Experience	Supports audio and video streaming applications on IP home networks.	
Remote Access Server (RAS) Connection Manager Administration Kit	Allows you to create connection manager profiles that simplify remote access configuration deployment to client computers.	
Remote Assistance	Allows remote support through invitations.	
Remote Differential Compression (RDC)	Transfers the differences between files over a network, minimizing bandwidth utilization.	
Remote Server Administration Tools	Collection of consoles and tools for remotely managing roles and features on other servers.	
Remote Procedure Call (RPC) over HTTP Proxy	Relays RPC traffic over HTTP as an alternative to VPN connections.	
Simple TCP/IP Services	Supports basic TCP/IP services, including Quote of the Day.	
Simple Mail Transfer Protocol (SMTP) Server	Supports transfer of email messages.	Deprecated in Windows Server 2012 R2
Simple Network Management Protocol (SNMP) Service	Includes SNMP agents that are used with the network management services.	

(continued from previous page)

Feature	Description	Changes in Windows Server 2012 R2
Subsystem for UNIX-based Applications	Supports Portable Operating System Interface for UNIX (POSIX)–compliant UNIX-based applications.	
Telnet Client	Allows outbound connections to Telnet servers and other Transmission Control Protocol (TCP)–based services.	
Telnet Server	Allows clients to connect to the server using the Telnet protocol.	Deprecated in Windows Server 2012 R2
Trivial File Transfer Protocol (TFTP) Client	Allows you to access TFTP servers.	
User Interfaces and Infrastructure	Contains the components necessary to support the graphical interface installation option on Windows Server 2012 R2. On graphical installations, this feature is installed by default.	
Windows Biometric Framework (WBF)	Allows use of fingerprint devices for authentication.	
Windows Feedback Forwarder	Supports sending of feedback to Microsoft when joining a Customer Experience Improvement Program (CEIP).	
Windows Identity Foundation 3.5	Set of .NET Framework classes that support implementing claims based identity on .NET applications.	Deprecated in Windows Server 2012 R2
Windows Internal Database	Relational data store that can only be used by Windows roles and features such as WSUS.	
Windows PowerShell	Task-based command-line shell and scripting language used to administer computers running Windows operating systems. This feature is installed by default.	Version 4.0 is installed in Windows Server 2012 R2
Windows PowerShell Web Access	Allows remote management of computers by running Windows PowerShell sessions in a web browser.	
Windows Process Activation service (WAS)	Allows applications hosting Windows Communication Foundation (WCF) services that to not use HTTP protocols to use features of IIS.	
Windows Search service	Allows fast searches of files hosted on a server for clients compatible with the Windows Search service.	

(continued from previous page)

Feature	Description	Changes in Windows Server 2012 R2
Windows Server Backup	Backup and recovery software for Windows Server 2012 R2.	
Windows Server Migration Tools	Collection of Windows PowerShell cmdlets that assist in the migration of server roles, operating system settings, files, and shares from computers running previous versions of Windows Server operating systems to Windows Server 2012 R2.	
Windows Standards-Based Storage Management	Set of Application Programming Interfaces (APIs) that allow the discovery, management, and monitoring of storage devices that use standards such as Storage Management Initiative Specification (SMI-S).	
Windows System Resource Manager (WSRM)	Allows you to control the allocation of CPU and memory resources.	Removed in Windows Server 2012 R2
Windows TIFF IFilter	Supports Optical Character Recognition on Tagged Image File Format (TIFF) 6.0-compliant files.	
WinRM IIS Extension	Windows Remote Management for IIS.	
Windows Internet Naming Service (WINS) Server	Supports name resolution for NetBIOS names.	
Wireless local area network (LAN) Service	Allows the server to use a wireless network interface.	
Windows on Windows (WoW) 64 Support	Supports running 32-bit apps on Server Core installations. This feature is installed by default.	
XPS Viewer	Supports the viewing and signing of documents in XPS formats.	

Features on Demand

Features on Demand enables you to add and remove role and feature files, also known as *feature payload*, from the Windows Server 2012 R2 operating system to conserve space. You can install roles and features when the feature payload is not present by using a remote source, such as a mounted image of the full operating system. If an installation source is not present but an Internet connection is, source files will be downloaded from Windows Update. The advantage of a Features on Demand installation is that it requires less hard disk space than a traditional installation. The disadvantage is that if you want to add a role or feature, you must have access to a mounted installation source. This is something that is not necessary if you perform an installation of Windows Server 2012 R2 with the graphical features enabled.

Question: Which feature do you need to install to support NetBIOS name resolution for client computers running a Microsoft Windows NT® 4.0 operating system workstation?

Lesson 2

Installing Windows Server 2012 R2

When preparing to install Windows Server 2012 R2, you need to understand whether a particular hardware configuration is appropriate. You also need to know whether a Server Core deployment might be more suitable than a full GUI deployment, and which installation source allows you to deploy Windows Server 2012 R2 in an efficient manner.

In this lesson you will learn about the process of installing Windows Server 2012 R2, including the methods that you can use to install the operating system, the different installation options, the minimum system requirements, and the decisions that you need to make when using the Installation Wizard.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the different methods that you can use to install Windows Server 2012 R2.
- Identify the different installation types that you can choose when installing the Windows Server 2012 R2.
- Determine whether to upgrade or migrate to Windows Server 2012 R2.
- Determine whether a computer or virtual machine meets the minimum hardware requirements necessary to install Windows Server 2012 R2.
- Describe the decisions that you need to make when performing a Windows Server 2012 R2 installation.
- Describe how to migrate server roles and features.

Installation Methods

Microsoft distributes Windows Server 2012 R2 on optical media and in an .iso (ISO) image format. ISO format is becoming more common as organizations acquire software over the Internet rather than obtaining physical removable media.

Once you have obtained the Windows Server 2012 R2 operating system from Microsoft, you can then use your own method to deploy the operating system. You can install Windows Server 2012 R2 by using a variety of methods, including the following:

- Optical Media
 - Advantages include:
 - Traditional method of deployment
 - Disadvantages include:
 - Requires that the computer have access to a DVD-ROM drive.
 - Is typically slower than USB media.
 - You cannot update the installation image without replacing the media.
 - You can only perform one installation per DVD-ROM at a time.



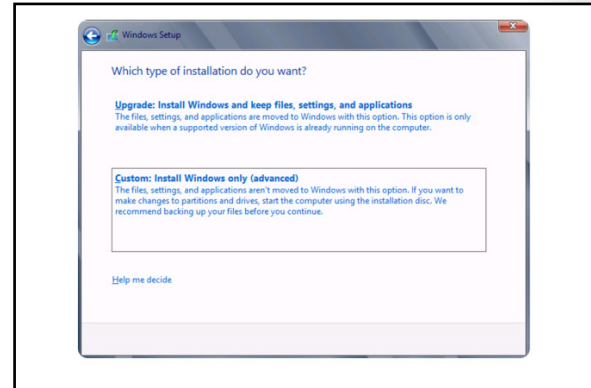
- USB Media
 - Advantages include:
 - All computers with USB drives allow boot from USB media.
 - The image can be updated as new software updates and drivers become available.
 - The answer file can be stored on a USB drive, minimizing the amount of interaction that the administrator must perform.
 - Disadvantages include:
 - Requires the administrator perform special steps to prepare USB media from an ISO file.
- Mounted ISO image
 - Advantages include:
 - With virtualization software, you can mount the ISO image directly and install Windows Server 2012 R2 on the virtual machine.
- Network Share
 - Advantages include:
 - It is possible to boot a server off a boot device (DVD or USB drive) and install from installation files that are hosted on a network share.
 - Disadvantages include:
 - This method is much slower than using Windows Deployment Services. If you already have access to a DVD or USB media, it is simpler to use those tools for operating system deployment.
- Windows Deployment Services
 - Advantages include:
 - You can deploy Windows Server 2012 R2 from .wim image files or specially prepared VHD files.
 - You can use the Windows Automated Installation Kit (AIK) to configure lite-touch deployment.
 - Clients perform a Preboot eXecution Environment (PXE) boot to contact the Windows Deployment Services server, and the operating system image is transmitted to the server over the network.
 - Windows Deployment Services allows multiple concurrent installations of Windows Server 2012 R2 using multicast network transmissions.
- System Center Configuration Manager
 - Advantages include:
 - Configuration Manager allows you to fully automate the deployment of Windows Server 2012 R2 to new servers that do not have an operating system installed. This process is called *Zero Touch deployment*.
- Virtual Machine Manager Templates
 - Advantages include:
 - Windows Server 2012 R2 is typically deployed in private cloud scenarios from preconfigured virtual machine templates. You can configure multiple components of the System Center suite to allow self-service deployment of Windows Server 2012 R2 virtual machines.

Question: What is another method that you can use to deploy Windows Server 2012 R2?

Installation Types

How you deploy Windows Server 2012 R2 on a specific server depends on the circumstances of that installation. Installing on a server that is running Windows Server 2008 R2 requires different actions than installing on a server running an x86 edition of Windows Server 2003.

When you are performing an installation of the Windows Server 2012 R2 operating system, you can choose one of the options in the following table.



Installation option	Description
Fresh installation	Allows you to perform a fresh install on a new disk or volume. Fresh installations are the most frequently used, and take the shortest amount of time. You can also use this option to configure Windows Server 2012 R2 to perform a dual boot if you want to keep the existing operating system.
Upgrade	An upgrade preserves the files, settings, and applications that are installed already on the original server. You perform an upgrade when you want to keep all of these items, and want to continue to use the same server hardware. You can only upgrade to an equivalent or newer edition of Windows Server 2012 from x64 versions of Windows Server 2008 and Windows Server 2008 R2. You can also upgrade from Windows Server 2012 to Windows Server 2012 R2. You launch an upgrade by running setup.exe from within the original Windows Server operating system.
Migration	Use migration when migrating to Windows Server 2012 R2 from x86 and x64 versions of Windows Server 2003, Windows Server 2003 R2, or Windows Server 2008. You can use the Windows Server Migration Tools feature in Windows Server 2012 R2 to transfer files and settings.

When you perform a fresh installation, you can deploy Windows Server 2012 R2 to an unpartitioned disk, or to an existing volume. You can also install Windows Server 2012 R2 to a specially-prepared VHD file in a "boot from VHD" or "VHD native boot" scenario. You might come across the use of both terms, or variations of them, to refer to this scenario. Boot from VHD requires special preparation, and is not an option that you can choose when performing a typical installation using the Windows Setup Wizard.

Choosing Whether to Upgrade or Migrate

Once the decision has been made to move to a new server operating system, and you have determined that your apps are compatible with the new operating system, you must choose whether to perform an in-place upgrade of the operating system on the existing hardware, or perform a clean install on new hardware and migrate the roles, apps and data.

If you choose an in-place upgrade, Setup performs compatibility checks to verify that all the components can be upgraded. Any identified issues are shown in a compatibility report that appears during Setup. This report might include guidance on what steps need to be taken to correct these issues.

In-place upgrade	
✓ Advantages: <ul style="list-style-type: none"> • Generally straightforward process which takes less time and planning than a migration strategy • All server roles, features, data and application settings are maintained 	✗ Disadvantages: <ul style="list-style-type: none"> • More difficult to troubleshoot installation failures caused by existing applications or server roles • Existing problems and configuration issues might be brought into the new operating system
Migration	
✓ Advantages: <ul style="list-style-type: none"> • Easier to troubleshoot installation failures • Existing configuration or application issues are not carried forward to the new operating system • You can easily move to updated versions of applications 	✗ Disadvantages: <ul style="list-style-type: none"> • Requires all applications to be re-installed and configured • Requires planning of migration of server roles • Requires migration of data • Requires the purchase of new hardware

The following table lists some advantages and disadvantages of an in-place upgrade.

Advantages	Disadvantages
<ul style="list-style-type: none"> • The process is generally straightforward and takes less time and planning than a migration strategy. • All server roles, features, data and application settings are maintained. 	<ul style="list-style-type: none"> • More difficult to troubleshoot installation failures that are caused by existing applications or server roles. • Existing problems and configuration issues might be brought into the new operating system.

The following table lists some advantages and disadvantages of a migration strategy.

Advantages	Disadvantages
<ul style="list-style-type: none"> • Easier to troubleshoot installation failures. • Any existing configuration or application issues are not carried forward to the new operating system. • Provides the opportunity to easily move to updated versions of applications. 	<ul style="list-style-type: none"> • Requires all applications to be re-installed and configured. • Requires planning of migration of server roles. • Requires migration of data. • Requires the purchase of new hardware.

Hardware Requirements for Windows Server 2012 R2

Hardware requirements define the minimum hardware that is required to run the Windows Server 2012 R2 server. Your actual hardware requirements might be greater, depending on the services that the server is hosting, the load on the server, and the responsiveness of your server.

Each role service and feature places a unique load on network, disk I/O, processor, and memory resources. For example, the File Server role places different stresses on server hardware than the DHCP role.

When considering hardware requirements, remember that Windows Server 2012 R2 can be deployed virtually. Windows Server 2012 R2 is supported on Hyper-V and some other non-Microsoft virtualization platforms. Windows Server 2012 R2 virtualized deployments need to match the same hardware specifications as physical deployments. For example, when creating a virtual machine to host Windows Server 2012 R2, you need to ensure that you configure the virtual machine with enough memory and hard disk space.

Windows Server 2012 R2 has the following minimum hardware requirements:

- Processor architecture: x64
- Processor speed: 1.4 gigahertz (GHz)
- Memory (RAM): 512 megabytes (MB)
- Hard disk drive space: 32 GB, more if the server has more than 16 GB of RAM

The Datacenter edition of Windows Server 2012 R2 supports the following hardware maximums:

- 640 logical processors
- 4 TB of RAM
- 63 failover cluster nodes



Additional Reading: For more information about the Windows Server Virtualization Validation Program, see <http://go.microsoft.com/fwlink/?LinkID=266736>.

Question: Why does a server need more hard disk drive space if it has more than 16 GB of RAM?

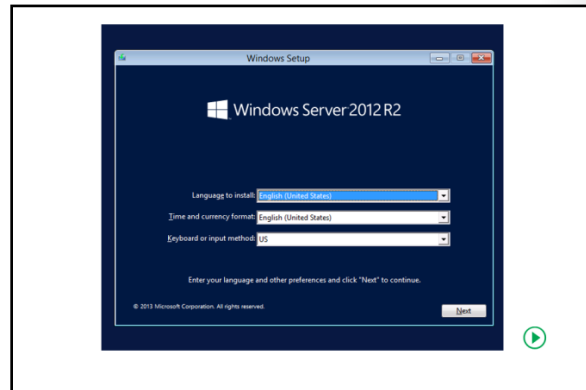
Windows Server 2012 has the following minimum hardware requirements:

- Processor architecture x64
- Processor speed 1.4 GHz
- Memory (RAM) 512 MB
- Hard disk drive space 32 GB
 - More hard disk drive space is needed if the server has more than 16 GB of RAM



Installing Windows Server 2012 R2

The process of deploying a Windows Server operating system is simpler today than it has been historically. The administrator performing the deployment has fewer decisions to make, although the decisions that they do make are critical to the success of the deployment. A typical installation of Windows Server 2012 R2 (if you do not already have an existing answer file), involves performing the following steps:



1. Connect to the installation source. Options for this include:
 - Insert a DVD-ROM containing the Windows Server 2012 R2 installation files, and boot from the DVD-ROM.
 - Connect a specially prepared USB drive that hosts the Windows Server 2012 R2 installation files.
 - Perform a PXE boot, and connect to a Windows Deployment Services 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**. You use this option in the event that an installation has become corrupted, and you are no longer able to boot into Windows Server 2012 R2.
4. In the Windows Setup wizard, on the **Select The Operating System You Want To Install** page, choose from the available operating system installation options. The default option is Server Core Installation.
5. On the **License Terms** page, review the terms of the operating system license. You must choose to accept the license terms before you can proceed with the installation process.
6. On the **Which Type Of Installation Do You Want** page, you have the following options:
 - **Upgrade**. Select this option if you have an existing installation of Windows Server that you want to upgrade to Windows Server 2012 R2. You should launch upgrades from within the previous version of Windows Server rather than 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, choose an available disk on which to install Windows Server 2012 R2. You can also choose to repartition and reformat disks from this page. When you click **Next**, the installation process will copy files and reboot the computer several times.
8. On the **Settings** page, provide a password for the local Administrator account.

Migrating Server Roles

If you choose to perform a clean install of the new server operating system, the first step is to migrate any existing server roles that were running on the server being replaced. Each server role is unique and requires its own migration strategy. Roles and features that rely on specific computer names and IP addresses require thorough planning to be successfully migrated. Some role migrations might require a service outage during the migration. Be sure to schedule these migrations so that they have the least impact possible on the production environment.

Windows Server Migration Tools assist in the migration process

Microsoft provides the following guides to assist in migration of roles and services:

- Migrate Active Directory Federation Services Role Services to Windows Server 2012
- Migrate Health Registration Authority to Windows Server 2012
- Migrate Hyper-VDI to Windows Server 2012
- Migrate IP Configuration to Windows Server 2012
- Migrate Network Policy Server to Windows Server 2012
- Migrate Print and Document Services to Windows Server 2012
- Migrate Remote Access to Windows Server 2012
- Migrate Windows Server Update Services to Windows Server 2012

Using Windows Server Migration Tools

Windows Server Migration Tools are available as a feature in Windows Server 2012 R2. Administrators can use these tools to migrate some server roles, features, operating system settings, shares, and other data from computers that are running certain editions of Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012 to computers that are running Windows Server 2012 R2. Not all role or feature migrations require these migration tools. The guides provide specific instructions on when and how to use the migration tools.

Migration Guides

Microsoft Technet provides the following migration guides to provide instructions for migrating specific roles and features to Server 2012:

- Migrate Active Directory Federation Services Role Services to Windows Server 2012
- Migrate Health Registration Authority to Windows Server 2012
- Migrate Hyper-V to Windows Server 2012
- Migrate IP Configuration to Windows Server 2012
- Migrating Network Policy Server to Windows Server 2012
- Migrate Print and Document Services to Windows Server 2012
- Migrate Remote Access to Windows Server 2012
- Migrate Windows Server Update Services to Windows Server 2012



Reference Links: To view the Windows Server 2012 migration guides, please see <http://go.microsoft.com/fwlink/?LinkID=331068>.

Lesson 3

Post-Installation Configuration of Windows Server 2012 R2

The Windows Server 2012 R2 installation process involves answering a minimal number of questions. Once you have completed installation, you need to perform several post-installation configuration steps before you can deploy it in a production environment. These steps allow you to prepare the server for the role it will perform on your organization's network.

This lesson includes how to perform a range of post-installation configuration tasks, including configuring network addressing information, setting a server's name and joining it to the domain, and understanding product activation options.

Lesson Objectives

After completing this lesson, you will be able to:

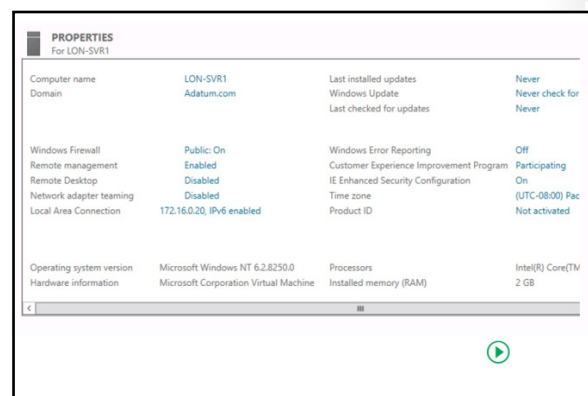
- Describe how to use Server Manager to perform post-installation configuration tasks.
- Describe how to configure server network settings.
- Describe how to join an Active Directory domain.
- Explain how to perform an offline domain join.
- Explain how to activate Windows Server 2012 R2.
- Describe how to configure a Server Core installation.

Overview of Post-Installation Configuration

The Windows Server 2012 R2 installation process minimizes the number of questions that you need to answer during the installation. The only information that you provide during the installation process is the password for the default local Administrator account. The post-installation process involves configuring all of the other settings that the server requires before it can be deployed to a production environment.

You use the Local Server node in the Server Manager console to perform the following tasks:

- 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

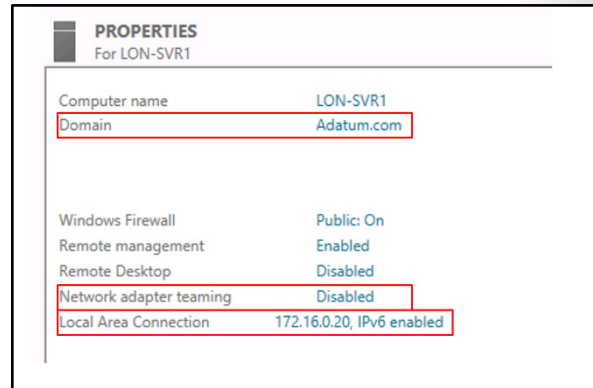



Configuring Server Network Settings

To communicate on the network, a server needs correct IP address information. Once you have completed installation, you need to either set or check the server's IP address configuration. By default, a newly-deployed server attempts 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 IP Addressing (APIPA) range of 169.254.0.1 to 169.254.255.254, then the server

has not yet been configured with an IP address from a DHCP server. This might be because a DHCP server has not yet been configured on the network, or, if there is a DHCP server, because there is a problem with the network infrastructure that blocks the adapter from receiving an address.



 **Note:** If you are using only an IPv6 network, then an IPv4 address in this range is not problematic, and IPv6 address information is still configured automatically.

Configuration Using Server Manager

You can manually configure IP address information for a server by performing the following steps:

1. In the Server Manager console, click the address next to the network adapter that you want to configure.
2. In the Network Connections window, right-click the network adapter for which you want to configure an address, 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** twice:
 - IP address
 - Subnet Mask
 - Default Gateway
 - Preferred DNS server
 - Alternate DNS server

Command-Line IPv4 Address Configuration

You can set IPv4 address information manually from an elevated command prompt by using the netsh.exe command from the interface IPv4 context or by using Windows PowerShell.

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 commands:

```
Netsh interface ipv4 set address "Local Area Connection" static 10.10.10.10 255.255.255.0
New-NetIPAddress -InterfaceIndex 12 -IPAddress 10.10.10.10 -PrefixLength 24
```

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
Set-DNSClientServerAddress -InterfaceIndex 12 -ServerAddresses 10.10.10.5
```

In the Windows PowerShell commands, the InterfaceIndex value identifies which adapter you are configuring. To get a complete list of adapters with corresponding InterfaceIndex values, run the Get-NetIPInterface cmdlet.

Network Interface Card Teaming

With Network Interface Card (NIC) teaming, you can increase the availability of a network resource. When you configure the NIC teaming feature, a computer uses one network address for multiple cards. In the event that one of the cards fails, the computer is able to maintain communication with other hosts on the network that are using that shared address. NIC teaming does not require that the network cards be the same model or use the same driver. To team network cards, perform the following steps:

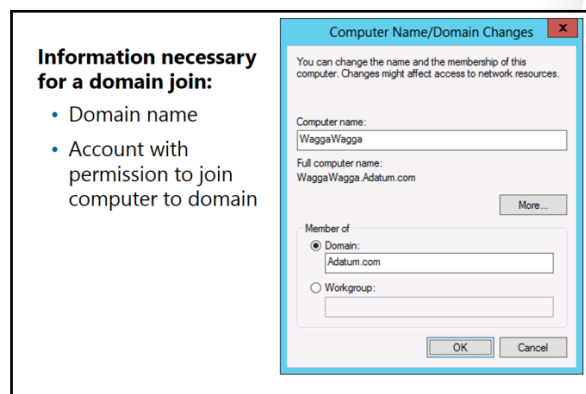
1. Ensure that the server has more than one network adapter.
2. In Server Manager, click the **Local Server** node.
3. Next to Network Adapter Teaming, click **Disabled**. This will launch the **NIC Teaming** dialog box.
4. In the **NIC Teaming** dialog box, hold down the Ctrl key, and then click each 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, provide a name for the team, and then click **OK**.

How to Join a Domain

When you install Windows Server 2012 R2, the computer is assigned a random name. Prior to joining a domain, you should configure the server with the name that you want it to have in the domain. As a best practice, you should use a consistent naming scheme when devising a computer name. Computers should be given names that reflect their function and location, not names with personal ties, such as pet names, or fictional or historical characters. For instance, it is simpler for everyone to determine that a server named MEL-DNS1 is a DNS server in Melbourne, than it is to determine that a server named Copernicus holds the DNS role in the Melbourne office.

You change this name using the Server Manager console by performing the following steps:

1. In Server Manager, click the **Local Server** node.
2. In the Properties window, click the active text next to Computer Name. This launches the **System Properties** dialog box.
3. In the **System Properties** dialog box, in the **Computer Name** tab, click **Change**.



4. In the **Computer Name/Domain Changes** dialog box, enter the new name that you want to assign to the computer.
5. Restart the computer to implement the name change.

Prior to joining the domain, be sure to complete the following steps to verify that the new server is ready to be domain-joined:

1. Ensure that you are able to resolve the IP address of the domain controller and that you can contact that domain controller. Use the PING protocol to ping the domain controller by hostname to accomplish both of these goals.
2. Complete one of the following tasks:
 - Create a computer account in the domain that matches the name of the computer that you want to join to the domain. This is often done when large numbers of computers need to be joined to the domain automatically.
 - Join the computer to the domain using a security account that has the right to perform domain-join operations.
3. Verify that the security account that is used for the domain operation already exists within the domain.

Now that you have renamed your Windows Server 2012 R2 server and have verified that it is ready to be domain-joined, you can join the server to the domain.

To join the domain using Server Manager, perform the following steps:

1. In Server Manager, click the **Local Server** node.
2. In the Properties window, next to Workgroup, click **WORKGROUP**.
3. In the **System Properties** dialog box, on the **Computer Name** tab, click **Change**.
4. In the **Computer Name/Domain Changes** dialog box, in the **Member Of** area, click the **Domain** option. Enter the new domain name, and then click **OK**.
5. In the **Windows Security** dialog box, enter domain credentials that allow you to join the computer to the domain.
6. Restart the computer.

Performing an Offline Domain Join

Offline domain join is a feature that you can use to join a computer to the domain when that computer does not have an active network connection. This feature can be useful in situations where connectivity is intermittent, such as when you are deploying a server to a remote site that is connected via satellite uplink.

Use the **djoin.exe** command-line tool to perform an offline domain join. You can perform an offline domain join by performing the following steps:

1. Sign in to the domain controller with a user account that has the appropriate rights to join other computers to the domain.

Process for performing an offline join:

1. Ensure the user performing the offline domain join is an administrator, or delegate the appropriate rights
2. To provision the computer account object and create the binary file, run **djoin** on the domain controller
3. Transfer the binary file with the domain information to client computer system's hard disk drive
4. To load the binary file, in the destination computer, run **djoin**
5. Restart the client computer



2. Open an elevated command prompt and use the **djoin.exe** command with the **/provision** option. You also need to specify the domain to which you want to join the computer, the name of the computer you will be joining to the domain, and the name of the savefile that you will transfer to the target of the offline domain join.

For example, to join the computer Canberra to the domain adatum.com using the savefile Canberra-join.txt, type the following command:

```
djoin.exe /provision /domain adatum.com /machine canberra /savefile
c:\canberra-join.txt
```

3. Transfer the generated savefile to the new computer, and then run the **djoin.exe** command with the **/requestODJ** option.

For example, to perform the offline domain join, after transferring the savefile Canberra-join.txt to computer Canberra, you would run the following command from an elevated command prompt on Canberra:

```
djoin.exe /requestODJ /loadfile canberra-join.txt /windowspath %systemroot% /localos
```

4. Restart the computer to complete the domain-join operation.

Question: In what situation would you perform an offline domain join rather than a traditional domain join?

Activating Windows Server 2012 R2

To ensure that your organization is correctly licensed and to receive notices for product updates, you must activate every copy of Windows Server 2012 R2 that you install. Windows Server 2012 R2 requires activation after installation. Unlike previous versions of the Windows Server operating system, there is no longer an activation grace period. If you do not perform activation, you cannot perform operating system customization.

To activate Windows Server 2012 R2, you can use one of two general strategies:

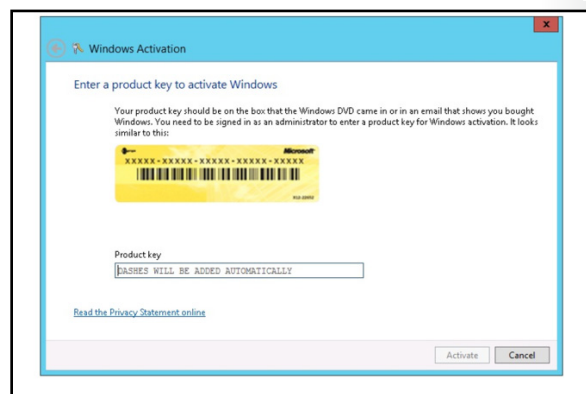
- Manual activation. Suitable when you are deploying a small number of servers.
- Automatic activation. Suitable when you are deploying larger numbers of servers.

Manual Activation

With manual activation, you enter the product key, and the server contacts Microsoft. Alternatively, an administrator performs the activation over the phone or through a special clearinghouse website.

You can perform manual activation from the Server Manager console by performing the following steps:

1. Click the **Local Server** node.
2. In the Properties window, next to Product ID, click **Not Activated**.
3. In the **Windows Activation** dialog box, enter the product key, and then click **Activate**.



4. If a direct connection cannot be established to the Microsoft activation servers, details will display about performing activation using a website from a device that has an Internet connection, or by using a local telephone number.

Because computers running the Server Core installation option do not have the Server Manager console, you can also perform manual activation using the **slmgr.vbs** command. Use the **slmgr.vbs /ipk** command to enter the product key, and **slmgr.vbs /ato** to perform activation once the product key is installed.

You can perform manual activation using either the retail product key, or the multiple activation key. You can use a retail product key to activate only a single computer. However, a multiple activation key has a set number of activations that you can use. Using a multiple activation key, you can activate multiple computers up to a set activation limit.

OEM keys are a special type of activation key that are provided to a manufacturer and allow automatic activation when a computer is first powered on. This type of activation key is typically used with computers that are running client operating systems such as Windows 7 and Windows 8. OEM keys are rarely used with computers that are running server operating systems.

Performing activation manually in large-scale server deployments can be cumbersome. Microsoft provides a method of activating large numbers of computers automatically without having to enter product keys on each system manually.

Automatic Activation

In previous versions of the Windows Server operating system, you could use KMS to perform centralized activation of multiple clients. The Volume Activation Services server role in Windows Server 2012 R2 allows you to manage a KMS server through a new interface. This simplifies the process of installing a KMS key on the KMS server. When you install Volume Activation Services, you can also configure Active Directory–based activation. Active Directory–based activation allows automatic activation of domain-joined computers. When you use Volume Activation Services, each computer activated must periodically contact the KMS server to renew its activation status.

You use the Volume Activation Management Tool (VAMT) 3.0 in conjunction with Volume Activation Services to perform activation of multiple computers on networks that are not connected directly to the Internet. You can use VAMT to generate license reports and manage client and server activation on enterprise networks.

Automatic Virtual Machine Activation

Automatic Virtual Machine Activation (AVMA) is a new feature of Windows Server 2012 R2 that you can use to install virtual machines on a Windows Server 2012 R2 Datacenter Hyper-V host and automatically activate them, provided that the host server is properly activated. The registry on the virtualization host provides real-time tracking data for the guest operations systems. This registry key moves with the virtual machine when it is migrated to a new host.

A special key, the AVMA key, is installed to the virtual machine using the following command from an elevated command prompt:

```
Slmgr /ipk <AVMA_key>
```

The virtual machine automatically activates the license against the virtualization host.

AVMA provides the following benefits:

- You can activate virtual machines in remote locations.
- You do not need an Internet connection to activate virtual machines.

- You can track licenses can be tracked from the virtualization host without requiring access rights on the virtual machines.
- There are no product keys to manage.
- Virtual machines remain activated when migrated across virtualization hosts.

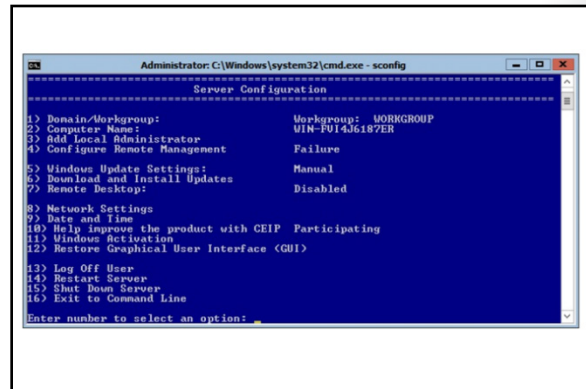
The following guest virtual machine operating systems are supported:

- Windows Server 2012 R2 Datacenter edition
- Windows Server 2012 R2 Standard edition
- Windows Server 2012 R2 Essentials edition

Configuring a Server Core Installation

Performing post installation on a computer running the Server Core operating system option can be daunting to administrators that have not performed the task before. Instead of having GUI-based tools that simplify the post-installation configuration process, IT professionals are faced with performing complex configuration tasks from a command-line interface.

The good news is that you can perform the majority of post-installation configuration tasks using the `sconfig.cmd` command-line tool. Using this tool minimizes the possibility of making syntax errors when using more complicated command-line tools.



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 WinRM
- Enable Windows Update
- Download and install updates
- Enable Remote Desktop
- Configure Network Address information
- Set the date and time
- Perform Windows Activation
- Sign out
- Restart the server
- Shut down the server

Configure IP Address Information

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

1. From a command-line command, run **sconfig.cmd**.
2. Choose **option 8** to configure Network Settings.
3. Choose the index number of the network adapter to which you want to assign an IP address.
4. In the Network Adapter Settings area, choose 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 a server's 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's name using **sconfig.cmd** by performing the following steps:

1. From a command-line command, run **sconfig.cmd**.
2. Choose **option 2** to configure the new computer name.
3. Type the new computer name, and then press Enter.

You must restart the 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, type the following command:

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



Note: Prior to joining the domain, verify that you are able to ping the DNS server by hostname.

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

1. From a command-line command, run **sconfig.cmd**.
2. Choose **option 1** to configure Domain/Workgroup.
3. To choose the Domain option, type **D**, and then press Enter.
4. Type the name of the domain to which you want to join the computer.
5. Provide the details in *domain\username* format, of an account that is authorized to join the domain.
6. Type the password associated with that account.

To complete a domain join operation, it is necessary to restart the computer.

Adding Roles and Features

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

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

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

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

```
Install-WindowsFeature NLB
```

Not all features are available directly 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 directly 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 R2. You can mount an installation image using the **DISM.exe** command-line tool. If you do not specify a source path when installing a component that is not available and the server has Internet connectivity, **Install-WindowsFeature** will attempt to retrieve source files from Windows Update.

Add the GUI

You can configure a Server Core computer with the GUI using the **sconfig.cmd** command-line tool. To do this, choose **option 12** from within the sconfig.cmd Server Configuration menu.



Note: You can add or remove the graphical component of the Windows Server 2012 R2 operating system by using the **Install-WindowsFeature** cmdlet.

You can also use the **dism.exe** command-line tool to add and remove Windows roles and features from a Server Core deployment, even though this tool is used primarily for managing image files.

Demonstration: Using DISM to Add Windows Features

Deployment Image Servicing and Management (DISM) is a command-line tool that you can use to service offline images or running operating systems. Use it to install, uninstall, configure, and update Windows features, packages, drivers and international settings.

After an image has been mounted to the file system, use DISM to service that image by specifying the path to the image and the servicing options in the command line. DISM can also service running systems when you specify the **/online** parameter and the servicing options in the command line.

In this demonstration, you will see how to use DISM to enable the Windows Server Backup feature for a running system. For example, if you were servicing an offline image, you would first use the DISM **/mount-image** parameter to mount the image to the file system. Then you would use the DISM **/image:<path to imagefile>** parameter and pass servicing commands to the image.

To service a VHD file you need to attach the virtual disk using the DiskPart.exe command line tool. For example, to attach a VHD file at C:\Server.VHD to the drive letter V, you can perform the following steps.

At an elevated command prompt, run the following commands:

```
DiskPart
Select vdisk file C:\Server.vhd
Attach vdisk
Assign letter=V
Exit
```

After you have finished servicing the VHD file using DISM, you can detach the VHD file using the following commands:

```
DiskPart
Select vdisk file C:\Server.vhd
Detach vdisk
Exit
```



Reference Links: For more information about using DISM, see the article Enable or Disable Windows Features at <http://go.microsoft.com/fwlink/?LinkID=331067>.



Additional Reading: For more information about using DISM to service VHD files, see the article Walkthrough: Service a Virtual Hard Disk Image Offline at <http://go.microsoft.com/fwlink/?LinkID=331066>.

Demonstration Steps

View a list of all Windows features and their current state

1. Use Server Manager to launch the Windows Server Backup MMC.
Notice that Windows Server Backup is not installed on the computer.
2. Close the wbadmin-[Windows Server Backup(Local)] window.

Gather information about the Windows Server Backup feature

1. Launch **Windows PowerShell** as Administrator.
2. Execute the following command:

```
DISM /online /get-features
```

3. Execute the following command:

```
DISM /online /get-featureinfo /featurename:WindowsServerBackup
```

Enable the Windows Server Backup feature

1. Execute the following command:

```
DISM /online /enable-feature /featurename:WindowsServerBackup
```



Note: The feature name is case-sensitive.

2. Use Server Manager to launch the Windows Server Backup MMC.
Notice that Windows Server Backup is now available.
3. Close all open windows.

Lesson 4

Overview of Windows Server 2012 R2 Management

Configuring a server correctly can save you from substantial problems later. Windows Server 2012 R2 provides multiple tools to perform specific administrative tasks, each of which is appropriate for a given set of circumstances. The Windows Server 2012 R2 management interface also enhances your ability to perform administrative tasks on more than one server simultaneously.

This lesson covers the different management tools that you can use to perform administrative tasks on computers that are running the Windows Server 2012 R2 operating system.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe Server Manager.
- Describe how to use administrative tools and Remote Server Administration Tools.
- Describe how to use Server Manager to perform administrative tasks.
- Describe how to configure services.
- Describe how to configure Windows Remote Management.

What Is Server Manager?

Server Manager is the primary graphical tool used to manage computers running Windows Server 2012 R2. You can use the Server Manager console to manage both the local server and remote servers. You can also manage servers as groups. By managing servers as groups, you can perform the same administrative tasks quickly across multiple servers that either perform the same role, or are members of the same group.

You can use the server manager console to perform the following tasks on both local servers and remote servers:

- Add roles and features
- Launch Windows PowerShell sessions
- View events
- Perform server configuration tasks

Server Manager can be used to manage the following down-level systems:

- Windows Server 2008 R2 SP1 (both full server or Server Core)
- Windows Server 2008 SP2 (full server only)

To manage these systems, you must install Windows Management Framework 3.0 (WMF 3.0) on the managed systems.

You can use Server Manager to:

- Manage multiple servers on a network from one console
- Add roles and features
- Launch Windows PowerShell sessions
- View events
- Perform server configuration tasks
- Manage down-level servers



You can use Best Practices Analyzer to:

- Determine whether roles on your network are functioning efficiently
- Query event logs for warning and error events
- Diagnose health issues with specific roles

Best Practices Analyzer

Server Manager includes a Best Practices Analyzer tool for all Windows Server 2012 R2 roles. With Best Practices Analyzer, you can determine whether roles on your network are functioning efficiently or if there are problems that you need to remediate. Best Practices Analyzer examines how a role functions—including querying associated event logs for warning and error events—so you can be aware of health issues associated with specific roles before those health issues cause a failure that impacts the server functionality.

Administrative Tools and Remote Server Administration Tools

When you use Server Manager to perform a specific role-related or feature-related administrative task, the console launches the appropriate administrative tool. When you install a role or feature using Server Manager locally or remotely, you are prompted to install the appropriate administrative tool. For example, when you use Server Manager to install the DHCP role on another server, you are prompted to install the DHCP console on the local server.

Administrative tools:

- Active Directory Administrative Center
- Active Directory Users and Computers
- DNS console
- Event Viewer
- Group Policy Management Console
- IIS Manager
- Performance Monitor
- Resource Monitor
- Task Scheduler
- Can be used to manage non-domain joined systems

Remote Server Administration Tools

You can install the complete set of administrative tools for Windows Server 2012 R2 by installing the Remote Server Administration Tools (RSAT) feature. When you install RSAT, you can choose to install all of the tools, or only the tools to manage specific roles and features. You can also install RSAT on computers running the Windows 8 operating system. This allows administrators to manage servers remotely, without having to sign in directly to each server.

It is a general best practice to run a Windows Server 2012 R2 server as a Server Core installation and manage it remotely via RSAT for Windows 8, or with one of the many other remote management methods.

In addition to Windows PowerShell, the tools that administrators most commonly use include the following:

- **Active Directory Administrative Center.** With this console, you can perform Active Directory administrative tasks such as raising domain and forest functional levels, and enabling the Active Directory Recycle Bin. You also use this console to manage Dynamic Access Control.
- **Active Directory Users and Computers.** With this tool, you can create and manage Active Directory users, computers, and groups. You can also use this tool to create Organizational Units (OUs).
- **DNS console.** With the DNS console, you can configure and manage the DNS Server role. This includes creating forward and reverse lookup zones, and managing DNS records.
- **Event Viewer.** You can use the Event Viewer to view events recorded in the Windows Server 2012 R2 event logs.
- **Group Policy Management Console.** With this tool, you can edit Group Policy Objects (GPOs) and manage their application in AD DS.
- **IIS Manager Tool.** You can use this tool to manage websites.
- **Performance Monitor.** You can use this console to view record performance data by selecting counters associated with specific resources that you want to monitor.

- Resource Monitor. You can use this console to view real-time information on CPU, memory, and disk and network utilization.
- Task Scheduler. You can use this console to manage the execution of scheduled tasks.

You can access each of these tools in Server Manager by accessing the Tools menu.



Note: You can also pin frequently used tools to the Windows Server 2012 R2 taskbar, or to the Start screen.

Managing non-domain joined Windows Server 2012 R2 with RSAT and Server Manager

Under normal circumstances, you cannot manage non-domain joined computers with RSAT from Windows 8 or Windows Server 2012 R2 systems. In Windows Server 2012 R2, however, there are Windows PowerShell commands that allow this configuration.

First, ensure you can resolve the name of the computer you will manage by its host name, because you cannot use its IP address. To do this, either manually add the server to your DNS or add an entry in the local Hosts file on the Windows 8 computer.

Next, on the Windows 8 computer running the RSAT, start Windows PowerShell as an Administrator and run the following command:

```
Set-Item WSMan:\localhost\Client\TrustedHosts -Value <YourtargetServernameHere> -Force
```

If the computer has a DNS suffix configured, then include both the host name and the fully qualified domain name (FQDN) of the server in quotes.

For example, to manage a server named SVR1 that has been manually configured to have a DNS suffix of contoso.com, the command would be:

```
Set-Item WSMan:\localhost\Client\TrustedHosts -Value "SVR1,SVR1.Contoso.com" -Force
```

In RSAT Server Manager, click **Manage** on the menu bar, click the DNS tab, enter the server name, and then add the server to the Selected column.

The server should appear in the Servers pane, and displaying a Kerberos error. Right-click the entry and select **Manage As**, and then enter the local Administrator credentials of the server.

Demonstration: Using Server Manager

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

- Sign in to Windows Server 2012 R2 and view the Windows Server 2012 R2 desktop.
- Add a feature by using the Add Roles and Features Wizard.
- View role-related events.
- Run the Best Practice Analyzer for a role.
- List the tools available from Server Manager.
- Restart Windows Server 2012 R2.

Demonstration Steps

Add a feature by using the Add Roles and Features Wizard

1. In Server Manager, start the Add Roles and Features Wizard.
2. Select the **Role-based or featured-based installation** check box.
3. Click **Select a server from the server pool**, verify that **LON-DC1.Adatum.com** is selected, and then click **Next**.
4. On the **Select server roles** page, select **Fax Server**.
5. In the **Add Roles and Features Wizard** dialog box, click **Add Features**.
6. On the **Select features** page, click **BranchCache**.
7. On the **Print and Document Services** page, click **Next twice**.
8. On the **Fax Server** page, click **Next**.
9. On the **Confirmation** page, select the **Restart the destination server automatically if required** check box, click **Yes**, click **Install**, and then click **Close**.
10. Click the flag icon next to **Server Manager Dashboard**, and review the messages.



Note: You can close this console without terminating the task.

View role-related events

1. Click the **Dashboard** node.
2. In the **Roles and Server Groups** pane, under **DNS**, click **Events**.
3. On the **DNS - Events Detail View**, change the time period to **12 hours** and the **Event Sources** to **All**.

Run the Best Practice Analyzer for a role

1. Under DNS, click **BPA** results.
2. Select **All** on the **Severity Levels** drop-down menu, and then click **OK**.

List the tools available from Server Manager

- Click the **Tools** menu, and review the tools that are installed on **LON-DC1**.

Sign out the currently signed-in user

1. Sign out from LON-DC1.
2. Sign back in to **LON-DC1** using the **Adatum\Administrator** account and the password **Pa\$\$w0rd**.

Restart Windows Server 2012 R2

- In a Windows PowerShell window, type the following command, and then press Enter:

```
Shutdown /r /t 5
```

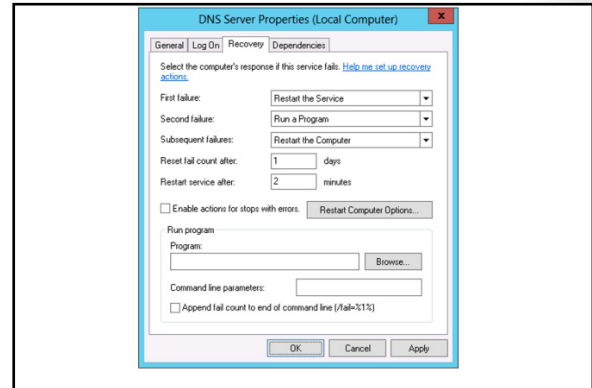

Configuring Services

Services are programs that run in the background and provide services to clients and to the host server. You can manage services through the Services console, which is available in Server Manager from the Tools menu. When securing a computer, you should disable all services except those that are required by the roles, features, and applications that are installed on the server.

Startup Types

Services use one of the following startup types:

- Automatic. The service starts automatically when the server boots.
- Automatic (Delayed Start). The service starts automatically after the server has booted.
- Manual. The service must be started manually, either by a program or by an administrator.
- Disabled. The service is disabled and cannot be started.



Note: If a server is behaving problematically, open the Services console, sort by startup type, and then locate those services that are configured to start automatically and which are not in a running state.

Service Recovery

Recovery options determine what a service does in the event that it fails. You access the Recovery tab from the DNS Server Properties window. On the Recovery tab, you have the following recovery options:

- Take no action. The service remains in a failed state until attended to by an administrator.
- Restart the Service. The service restarts automatically.
- Run a Program. Allows you to run a program or a script.
- Restart the Computer. The computer restarts after a preconfigured number of minutes.

You can configure different recovery options for the first failure, the second failure, and subsequent failures. You can also configure a period of time after which the service failure clock resets.

Managed Service Accounts

Managed service accounts are special domain-based accounts that you can use with services. The advantage of a managed service account is that the account password is rotated automatically according to a schedule. These password changes are automatic, and do not require administrator intervention, which minimizes the chance that the service account password will be compromised. This happens typically because administrators traditionally assign simple passwords to service accounts with the same service across a large number of servers, and never bother to update those passwords. Virtual accounts are service-specific accounts that are local rather than domain-based. Windows Server 2012 R2 rotates and manages the password for virtual accounts.

Question: What is the advantage of a managed service account compared to a traditional domain-based service account?

Configuring Windows Remote Management

Most administrators no longer perform systems administration tasks solely from the server room. Almost all tasks that they perform on a daily basis are now performed using remote management technologies.

With Windows Remote Management (WinRM), you can use Remote Shell, remote Windows PowerShell, and other remote management tools to manage a computer remotely.

You can enable WinRM from Server Manager by performing the following steps:

1. In the Server Manager console, click the **Local Server** node.
2. In the **Properties** dialog box for the local server, next to Remote Management, click **Disabled**. This opens the **Configure Remote Management** dialog box.
3. In the **Configure Remote Management** dialog box, select the **Enable remote management of this server from other computers** check box, and then click **OK**.

When deciding to use Remote Management, consider the following:

- You are more likely to manage a server remotely than by locally signing on
- With WinRM, you can use consoles, command-line utilities, or Windows PowerShell to perform remote management tasks
- With Remote Desktop, you can sign in to a server locally or from across the network

You also can enable WinRM from a command line by running the command **WinRM qc**. You disable WinRM by using the same method that you use to enable it. You can disable WinRM on a computer running the Server Core installation option using the `sconfig.cmd` tool.

Remote Desktop

Remote Desktop is the traditional method by which systems administrators connect remotely to the servers that they manage. You can configure Remote Desktop on a computer that is running the full version of Windows Server 2012 R2 by performing the following steps:

1. In the Server Manager console, click the **Local Server** node.
2. Next to Remote Desktop, click **Disabled**.
3. In the **System Properties** dialog box, on the **Remote** tab, select 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**. Allows connections from Remote Desktop clients that do not support Network Level Authentication
 - **Allow Connections only from Computers running Remote Desktop with Network Level Authentication**. Allows secure connections from computers running Remote Desktop clients that support network-level authentication.

You can enable and disable Remote Desktop on computers that are running the Server Core installation option by using the `sconfig.cmd` command-line tool.

Managing Non-Domain Joined Computers

Many users want to use their own mobile devices, such as smart phones and tablets, to access company apps and data. Managing these devices is difficult when they do not belong to the domain. The domain administrator has little or no control over how the device is used off-premise.

There might also be pockets of workgroup computers in your environment. In this case, you can use local group policies to manage these systems, and use Remote Desktop technologies to manage them.

Demonstration: Performing Remote Management

In this demonstration, you will see how to use Server Manager to manage a remote server, how to add the DNS Server role on a remote server, and how to connect to and configure the remote server by using RDP.

Demonstration Steps

Use Server Manager to manage a remote server

1. Sign in to LON-DC1 as **Adatum\Administrator** with a password of **Pa\$\$w0rd**.
2. In Server Manager, click **Add other servers to manage**.
3. Add LON-SVR1.

Add the DNS Server role on a remote server

- Use the Add Roles and Features wizard to add the **DNS Server** role to LON-SVR1

Connect to and configure a remote server by using RDP

1. Sign in to LON-SVR1 as **Adatum\Administrator** with a password of **Pa\$\$w0rd**.
2. In Server Manager, open the properties for the Local Server, and then enable Remote Desktop.
3. On LON-DC1, go to the **Start** screen.
4. Type **Remote**, and then open **Remote Desktop Connection**.
5. Connect to LON-SVR1 as **Adatum\Administrator**.
6. After successfully connecting, sign out from LON-SVR1.

Lesson 5

Introduction to Windows PowerShell

Windows PowerShell is a command-line interface and task-based scripting technology that is built into the Windows Server 2012 R2 operating system. Windows PowerShell simplifies the automation of common systems administration tasks. With Windows PowerShell, you can automate tasks, leaving you more time for more difficult systems administration tasks.

In this lesson, you will learn about Windows PowerShell, and why Windows PowerShell is a critical piece of a server administrator's toolkit.

This lesson describes how to use Windows PowerShell's built-in discoverability features to learn how to use specific cmdlets and to find related cmdlets. This lesson also discusses how to leverage the Windows PowerShell Integrated Scripting Environment (ISE) to assist you in creating effective Windows PowerShell scripts.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the purpose of Windows PowerShell.
- Describe Windows PowerShell cmdlet syntax, and explain how to determine commands associated with a particular cmdlet.
- Describe common Windows PowerShell cmdlets used to manage services, processes, roles and features.
- Describe the functionality of Windows PowerShell ISE.
- Explain how to use Windows PowerShell.
- Explain how to use Windows PowerShell ISE.

What Is Windows PowerShell?

Windows PowerShell is a scripting language and command-line interface that is designed to assist you in performing day-to-day administrative tasks. Windows PowerShell is made up of cmdlets that you execute at a Windows PowerShell command prompt, or combine into Windows PowerShell scripts. Unlike other scripting languages that were designed initially for another purpose, but have been adapted for system administration tasks, Windows PowerShell is designed with system administration tasks in mind.

name	ModuleName
----	-----
Set-BCAuthentication	BranchCache
Set-BCCache	BranchCache
Set-BCDataCacheEntryMaxAge	BranchCache
Set-BCInhibitStaleness	BranchCache
Set-BCSecurity	BranchCache
Set-ClientDownloadLocalTask	ScheduleTasks
Set-DACLientExperienceConfiguration	DirectAccessClientComponents
Set-GroupPolicyTableItem	DirectAccessClientComponents
Set-Disk	Storage
Set-DnsClient	DnsClient
Set-DnsClientGlobalSetting	DnsClient
Set-DnsClientInterglobal	DnsClient
Set-DnsClientIntersite	DnsClient
Set-DnsClientServerAddress	DnsClient
Set-DnsAdvancedSetting	MsiDc
Set-DnsAdvancedSetting	MsiDc
Set-DnsClusterDefault	MsiDc
Set-DnsClusterMapping	MsiDc
Set-DnsDefault	MsiDc
Set-DnsLog	MsiDc
Set-DnsNetworkSetting	MsiDc
Set-DnsTransaction	MsiDc
Set-DnsTransactionsTraceSession	MsiDc
Set-DnsTransactionsTraceSetting	MsiDc
Set-InitiatorPort	Storage
Set-IPSecChipsetcert	IPSec
Set-LogProperties	PSDiagnostics
Set-MgmtNet	MgmtNet
Set-NetIPolicyConfiguration	NetworkConnectivityStatus
Set-NetKerberosConfiguration	NetworkTransition
Set-NetKerberos	NetAdapter
Set-NetAdapterAdvancedProperty	NetAdapter
Set-NetAdapterBundling	NetAdapter
Set-NetAdvancedCacheDumpOffload	NetAdapter

An increasing number of Microsoft products—such as Exchange Server 2010—have graphical interfaces that build Windows PowerShell commands. These products allow you to view the generated Windows PowerShell script so you can execute the task at a later time without having to complete all of the steps in the GUI. Being able to automate complex tasks simplifies a server administrator's job, and saves time.

You can extend Windows PowerShell functionality by adding modules. For example, the Active Directory module includes Windows PowerShell cmdlets that are specifically useful for performing Active Directory–related management tasks. The DNS Server module includes Windows PowerShell cmdlets that are specifically useful for performing DNS server-related management tasks. Windows PowerShell includes features such as tab completion, which allows administrators to complete commands by pressing the tab key rather than having to type the complete command. You can learn about the functionality of any Windows PowerShell cmdlet by using the **Get-Help** cmdlet.



Note: Windows PowerShell Version 4.0 ships with Windows Server 2012 R2 and is backward compatible with previous versions of Windows PowerShell.

Windows PowerShell Cmdlet Syntax

Windows PowerShell cmdlets use a verb-noun syntax. Each noun has a collection of associated verbs. The available verbs differ with each cmdlet's noun.

Common Windows PowerShell cmdlet verbs include:

- Get
- New
- Set
- Restart
- Resume
- Stop
- Suspend
- Clear
- Limit
- Remove
- Add
- Show
- Write

You can view the available verbs for a particular Windows PowerShell noun by executing the following command:

```
Get-Command -Noun NounName
```

You can view the available Windows PowerShell nouns for a specific verb by executing the following command:

```
Get-Command -Verb VerbName
```

Windows PowerShell parameters start with a dash. Each Windows PowerShell cmdlet has its own associated set of parameters. You can learn what the parameters are for a particular Windows PowerShell cmdlet by executing the following command:

```
Get-Help CmdletName
```

Windows PowerShell Cmdlet

Syntax:

- Get-Help -Noun
NounName
- Get-Help -Verb
VerbName
- Help *CmdletName*
- Get-Command

Capability	Name
----	----
CIM	Set-BCAuthentication
CIM	Set-BCCache
CIM	Set-BCDataCacheEntryMaxAge
CIM	Set-BCWinSMBLatency
CIM	Set-BCSecretKey
CIM	Set-ClusteredScheduledTask
CIM	Set-DAClientExperienceConfiguration
CIM	Set-DAEntryPointTableItem
CIM	Set-Disk
CIM	Set-DnsClient
CIM	Set-DnsClientGlobalSetting
CIM	Set-DnsClientNrptGlobal
CIM	Set-DnsClientNrptRule
CIM	Set-DnsClientServerAddress
CIM	Set-DtcAdvancedHostSetting
CIM	Set-DtcAdvancedSetting
CIM	Set-DtcClusterDefault
CIM	Set-DtcClusterTMMapping
CIM	Set-DtcDefault
CIM	Set-DtcLog

You can determine which Windows PowerShell cmdlets are available by executing the **Get-Command** cmdlet. The Windows PowerShell cmdlets that are available depend on which modules are loaded. You can load a module using the **Import-Module** cmdlet.

Common Cmdlets for Server Administration

As a server administrator, there are certain cmdlets that you are more likely to use. These cmdlets relate primarily to services, event logs, processes, and the ServerManager module running on the server.

Service Cmdlets

You can use the following Windows PowerShell cmdlets to manage services on a computer that is running Windows Server 2012 R2:

- **Get-Service.** View the properties of a service.
- **New-Service.** Creates a new service.
- **Restart-Service.** Restarts an existing service.
- **Resume-Service.** Resumes a suspended service.
- **Set-Service.** Configures the properties of a service.
- **Start-Service.** Starts a stopped service.
- **Stop-Service.** Stops a running service.
- **Suspend-Service.** Suspends a service.

System Administration cmdlets	Details
Service Cmdlets	Use the Service noun
Event Log Cmdlets	Use the Eventlog noun
Process Cmdlets	Use the Process noun
ServerManager module	Allows the WindowsFeature noun
Windows PowerShell Remote Management	Allows cmdlets or scripts to be run on remote computers

Event Log Cmdlets

You can use the following Windows PowerShell cmdlets to manage event logs on a computer that is running Windows Server 2012 R2:

- **Get-EventLog.** Displays events in the specified event log.
- **Clear-EventLog.** Deletes all entries from the specified event log.
- **Limit-EventLog.** Sets event log age and size limits.
- **New-EventLog.** Creates a new event log and a new event source on a computer running Windows Server 2012 R2.
- **Remove-EventLog.** Removes a custom event log and unregisters all event sources for the log
- **Show-EventLog.** Shows the event logs of a computer.
- **Write-EventLog.** Allows you to write events to an event log.

Process Cmdlets

You can use the following Windows PowerShell cmdlets to manage processes on a computer that is running Windows Server 2012 R2:

- **Get-Process.** Provides information on a process.
- **Start-Process.** Starts a process.
- **Stop-Process.** Stops a process.

- **Wait-Process.** Waits for the process to stop before accepting input.
- **Debug-Process.** Attaches a debugger to one or more running processes.

ServerManager Module

The ServerManager module allows you to add one of three cmdlets that are useful for managing features and roles. These cmdlets are:

- **Get-WindowsFeature.** View a list of available roles and features. Also displays whether the feature is installed, and whether the feature is available. You can only install an unavailable feature if you have access to an installation source.
- **Install-WindowsFeature.** Installs a particular Windows Server role or feature. The **Add-WindowsFeature** cmdlet is aliased to this command and is available in previous versions of Windows operating systems.
- **Remove-WindowsFeature.** Removes a particular Windows Server role or feature.

Windows PowerShell Remote Management

You can use Windows PowerShell to remotely run cmdlets on other Windows systems, called *remoting*. Windows PowerShell remoting depends on the WinRM service running on the target systems. This service can be enabled manually or by running the Enable-PSRemoting cmdlet on the target.

The simplest way to use remoting is one-to-one remoting, which allows you to bring up an interactive Windows PowerShell session on the remote system. Once connected, the Windows PowerShell prompt displays the name of the remote computer.

To start or end a remote session use the following cmdlet:

```
Enter-PSSession -computername <name of computer to connect to>  
Exit-PSSession
```

You can also use the Invoke-command cmdlet to run commands on multiple remote computers. Specify other computers by using the *ComputerName* parameter, and specify the commands to run by setting the *ScriptBlock* parameter.

To use the invoke-command cmdlets to get the service status from multiple computers, first enable WinRM on the targets, and then use the following cmdlet:

```
Invoke-command -computername Lon-Srv1, Lon-Srv2 -scriptblock {get-service}
```

Demonstration: Using Windows PowerShell

In this demonstration, you will see how to use Windows PowerShell to display the running services and processes on a server and a remote server, to connect to a remote computer and display all services, and to invoke commands to multiple computers to display running processes.

Demonstration Steps

Use Windows PowerShell to display the running services and processes on a server

1. On LON-DC1, open a Windows PowerShell session.
2. Execute the following commands, and then press Enter:

```
Get-Service | where-object {$_.status -eq "Running"}  
Get-Command -Noun Service  
Get-Process  
Get-Help Process  
Get-Help -Full Start-Process
```

3. Close the Windows PowerShell window.
4. On the taskbar, right-click the **Windows PowerShell** icon, and then click **Run as Administrator**.

Use Windows PowerShell to connect to a remote computer and display all services and their current status

1. On LON-SVR1, open a Windows PowerShell session.
2. Execute the following command:

```
Enable-PSRemoting
```

3. Accept all the default prompts.
4. On LON-DC1, open a Windows PowerShell session.
5. Execute the following commands:

```
Enter-PSSession -Computername LON-SVR1  
Get-Service  
Exit-PSSession
```

Use Windows PowerShell to invoke commands to multiple computers and display running processes

1. On LON-DC1, execute the following command:

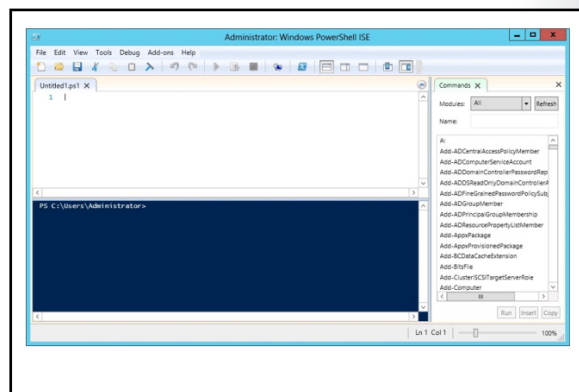
```
Invoke-Command -computername LON-DC1, LON-SVR1 -Scriptblock {Get-Process}
```

2. Examine the output, and then close the **Windows PowerShell** window.

Windows PowerShell ISE is an integrated scripting environment that provides you with assistance when using Windows PowerShell. It provides command completion functionality, and enables you to see all available commands and the parameters that you can use with those commands.

The ability to view cmdlet parameters ensures that you are aware of the full functionality of each cmdlet, and can create syntactically-correct Windows PowerShell commands.

You can use the Windows PowerShell ISE environment to view available cmdlets by module. You can then determine which Windows PowerShell module you need to load to access a particular cmdlet.



In this demonstration, you will see how to complete the following tasks:

- Use Windows PowerShell ISE to import the ServerManager module.
- View the cmdlets made available in the ServerManager module.
- Use the **Get-WindowsFeature** cmdlet from Windows PowerShell ISE.
- View and run a Windows PowerShell script

Use Windows PowerShell ISE to import the ServerManager module

1. Ensure that you are signed in to LON-DC1 as **Administrator**.
2. In Server Manager, click **Tools**, and then click **Windows PowerShell ISE**.
3. At the command prompt, type **Import-Module ServerManager**.

- In the Commands pane, use the **Modules** drop-down menu to select the **ServerManager** module.

1. Click **Get-WindowsFeature**, and then click **Show Details**.
2. In the **ComputerName** field, type **LON-DC1**, and then click **Run**.

Run a Windows PowerShell script from the scripting pane to create a universal group named Helpdesk and add members

1. In Server Manager, click **Tools**, and then open **Active Directory Users and Computers**.
2. Open the IT organizational unit (OU). Note that there is no group named Helpdesk.
3. Use File Explorer to navigate to **E:\Labfiles\Mod01** and edit the **CreateAndPopulateHelpdesk.ps1** script.
4. View the script and then click the green arrow on the toolbar to run the script.
5. Switch back to Active Directory Users and Computers and refresh the view. You should now see there is a group named Helpdesk.
6. Open the properties of the Helpdesk group and see that the group is populated by the members of the IT department.
7. Close all open windows.

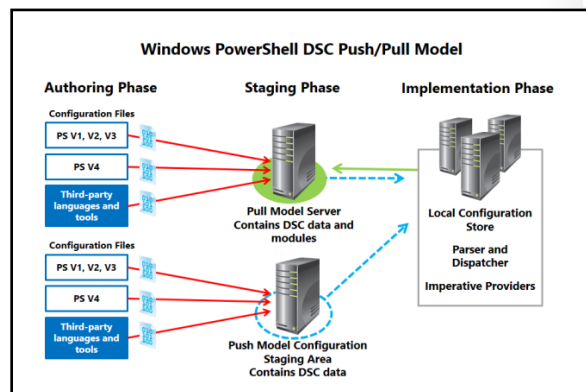
Windows PowerShell Desired State Configuration

Windows PowerShell Desired State Configuration (DSC) is a new management platform in Windows PowerShell that enables deploying and managing configurations for software services and managing the environment in which these services run. These Windows PowerShell extensions are available only in Windows Server 2012 R2 and Windows 8.1.

DSC uses WMI providers to implement configurations. DSC comes with 12 WMI providers that enable the configuration of roles and features, managing services and more. You can also create custom providers. This topic provides an overview of DSC.

DSC can be implemented in two different models; either Push or Pull. There are three main phases to implementing DSC:

1. The Authoring phase
2. The Staging phase (if implementing a Pull model)
3. The Implementation phase



These three main phases are described in the following table.

Phase	Description
Authoring phase	The DSC is created by using Windows PowerShell or by using third-party tools. Windows PowerShell commands are used to create one or more Management Object Format (MOF) files that describe the configuration settings.
Staging phase	<p>In a Pull model, DSC data and any custom providers are kept on the Pull server, which is an IIS server. The target system contacts the Pull server by passing a Uniform Resource Identifier (URI) along with a unique identifier to pull its DSC configuration and verify if any required providers are available. If they are not, those providers are downloaded to the target computer.</p> <p>In the case of a Push model, you need to ensure that any required providers are already in place on the target computer because only the configuration settings are pushed.</p>
Implementation phase	The final phase is the application of the configuration. Once DSC data is either pushed or pulled to the target server's Local Configuration Store, the configuration is then parsed and the appropriate WMI provider implements the settings.

DSC can be used to perform many different functions, including the following:

- Install or remove server roles and features
- Manage registry settings
- Manage files and directories
- Start, stop, and manage processes and services
- Manage local groups and user accounts
- Install and manage packages such as .msi and .exe
- Manage environment variables
- Run Windows PowerShell scripts
- Fix a configuration that has drifted away from the desired state
- Discover the actual configuration state on a given node

Lab: Deploying and Managing Windows Server 2012 R2

Scenario

A. Datum Corporation is a global engineering and manufacturing company with a head office based in London, England. A. Datum has recently deployed a Windows Server 2012 R2 infrastructure with Windows 8 clients.

You have been working for A. Datum for several years as a desktop support specialist and have recently accepted a promotion to the server support team.

The marketing department has purchased a new web-based application. You need to install and configure the servers in the data center for this application. One server has a GUI interface, and the other server is configured as Server Core.

Objectives

After completing this lab, you should be able to:

- Deploy Windows Server 2012 R2.
- Configure Windows Server 2012 R2 Server Core.
- Manage servers by using Server Manager.
- Manage servers with Windows PowerShell.

Lab Setup

Estimated Time: 60 minutes

Virtual machines	20410C-LON-DC1 20410C-LON-SVR3 20410C-LON-CORE
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 **20410C-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in using the following credentials:
 - a. User name: **Adatum\Administrator**
 - b. Password: **Pa\$\$w0rd**
5. Repeat steps 1 through 3 for 20410C-LON-CORE. Do not sign in until directed to do so.

Exercise 1: Deploying Windows Server 2012 R2

Scenario

The first Windows Server 2012 R2 server that you are installing for the Marketing department will host a SQL Server 2012 database engine instance. You want to configure the server so that it will have the full GUI, as this will allow the application vendor to run support tools directly on the server, rather than requiring a remote connection.

The first server you are installing for the new marketing app is for a SQL Server 2012 database. This server will have the full GUI to allow the application vendor to run support tools directly on the server.

The main tasks for this exercise are as follows:

1. Install the Windows Server 2012 server.
2. Change the server name.
3. Change the date and time.
4. Configure the network.
5. Add the server to the domain.

► Task 1: Install the Windows Server 2012 server

1. In the Hyper-V Manager console, open the settings for 20410C-LON-SVR3.
2. Configure the DVD drive to use the Windows Server 2012 R2 image file named **Windows2012R2RTM.iso**. This file is located at **D:\Program Files\Microsoft Learning\20410\Drives**.
3. Start **20410C-LON-SVR3**. In the Windows Setup Wizard, on the **Windows Server 2012 R2** page, 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. Click to install the **Windows Server Windows Server 2012 R2 Datacenter Evaluation (Server with a GUI)** operating system.
5. Accept the license terms, and then click **Custom: Install Windows only (advanced)**.
6. Install Windows Server 2012 R2 on **Drive 0**.



Note: Depending on the speed of the equipment, the installation takes approximately 20 minutes. The virtual machine will restart several times during this process.

7. Enter the password **Pa\$\$w0rd** in both the **Password** and **Reenter password** boxes, and then click **Finish** to complete the installation.

► Task 2: Change the server name

1. Sign in to LON-SVR3 as **Administrator** with the password **Pa\$\$w0rd**.
2. In Server Manager, on the **Local Server** node, click the randomly-generated name next to **Computer name**.
3. In the **System Properties** dialog box, on the **Computer Name** tab, click **Change**.
4. In the **Computer name** box, type **LON-SVR3**, and then click **OK**.

5. Click **OK** again, and then click **Close**.
6. Restart the computer.

► **Task 3: Change the date and time**

1. Sign in to server LON-SVR3 as **Administrator** with the password **Pa\$\$w0rd**.
2. On the taskbar, click the time display, and then click **Change date and time settings**.
3. Click **Change Time Zone**, and set the time zone to your current time zone.
4. Click **Change Date and Time**, and verify that the date and time that display in the **Date and Time Settings** dialog box match those in your classroom.
5. Close the **Date and Time** dialog box.

► **Task 4: Configure the network**

1. On LON-SVR3, click **Local Server**.
2. Next to Ethernet, click **IPv4 Address Assigned by DHCP, IPv6 Enabled**.
3. In the **Network Connections** dialog box, right-click **Ethernet**, and then click **Properties**.
4. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
5. Enter the following IP address information, then and click **OK**:
 - IP address: **172.16.0.101**
 - Subnet Mask: **255.255.0.0**
 - Default Gateway: **172.16.0.1**
 - Preferred DNS server: **172.16.0.10**
6. Close all dialog boxes.

► **Task 5: Add the server to the domain**

1. On LON-SVR3, in the Server Manager console, click **Local Server**.
2. Next to Workgroup, click **WORKGROUP**.
3. On the **Computer Name** tab, click **Change**.
4. Click the **Domain** option, and in the **Domain** box, enter **adatum.com**.
5. Enter the following account details:
 - Username: **Administrator**
 - Password: **Pa\$\$w0rd**
6. In the **Computer Name/Domain Changes** dialog box, click **OK**.
7. Restart the computer to apply changes.
8. In the **System Properties** dialog box, click **Close**.
9. After LON-SVR3 restarts, sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.

Results: After completing this exercise, you should have deployed Windows Server 2012 R2 on LON-SVR3. You also should have configured LON-SVR3, including name change, date and time, networking.

Exercise 2: Configuring Windows Server 2012 R2 Server Core

Scenario

The web-based tier of the marketing application is a .NET application. To minimize the operating system footprint and reduce the need to apply software updates, you have chosen to host the IIS component on a computer that is running the Server Core installation option of the Windows Server 2012 R2 operating system.

To enable this, you need to configure a computer that is running Windows Server 2012 R2 with the Server Core installation option.

The main tasks for this exercise are as follows:

1. Set computer name.
2. Change the computer's date and time.
3. Configure the network.
4. Add the server to the domain.

► Task 1: Set computer name

1. Sign in to LON-CORE as **Administrator** with the password **Pa\$\$w0rd**.
2. On LON-CORE, type **sconfig.cmd**.
3. Click option **2** to select **Computer Name**.
4. Set the computer name as **LON-CORE**.
5. In the **Restart** dialog box, click **Yes** to restart the computer.
6. After the computer restarts, sign in to server LON-CORE using the **Administrator** account with the password **Pa\$\$w0rd**.
7. At the command prompt, type **hostname**, and then press Enter to verify the computer's name.

► Task 2: Change the computer's date and time

1. Ensure you are signed in to server LON-CORE as **Administrator** with the password **Pa\$\$w0rd**.
2. At the command prompt, type **sconfig.cmd**.
3. To select **Date and Time**, type **9**.
4. Click **Change time zone**, and then set the time zone to the same time zone that your classroom uses.
5. In the **Date and Time** dialog box, click **Change Date and Time**, and verify that the date and time match those in your location.
6. Exit **sconfig.cmd**.

► Task 3: Configure the network

1. Ensure that you are signed in to server LON-CORE using the account **Administrator** and password **Pa\$\$w0rd**.
2. At the command prompt, type **sconfig.cmd**, and then press Enter.
3. To configure Network Settings, type **8**.
4. Type the number of the network adapter that you want to configure.
5. Type **1** to set the Network Adapter Address.
6. Click **static IP address configuration**, and then enter the address **172.16.0.111**.

7. At the Enter subnet mask prompt, type **255.255.0.0**.
8. At the Enter default gateway prompt, type **172.16.0.1**.
9. Type **2** to configure the DNS server address.
10. Set the preferred DNS server to **172.16.0.10**.
11. Do not configure an alternate DNS server address.
12. Exit `sconfig.cmd`.
13. Verify network connectivity to `lon-dc1.adatum.com` using the PING tool.

► Task 4: Add the server to the domain

1. Ensure that you are signed in to server LON-CORE using the account **Administrator** with password **Pa\$\$w0rd**.
2. At the command prompt, type **sconfig.cmd**, and then press Enter.
3. Type **1** to switch to **configure Domain/Workgroup**.
4. Type **D** to join a domain.
5. At the Name of domain to join prompt, type **adatum.com**.
6. At the Specify an authorized domain\user prompt, type **Adatum\Administrator**.
7. At the Type the password associated with the domain user prompt, type **Pa\$\$w0rd**.
8. At the prompt, click **No**.
9. Restart the server.
10. Sign in to server LON-CORE with the **Adatum\Administrator** account using the password **Pa\$\$w0rd**.

Results: After completing this exercise, you should have configured a Windows Server 2012 R2 Server Core deployment and verified the server's name.

Exercise 3: Managing Servers

Scenario

After deploying the servers LON-SVR3 and LON-CORE for hosting the Marketing application, you need to install appropriate server roles and features to support the application. With this in mind, you will install the Windows Server Backup feature on both LON-SVR3 and LON-CORE. You will install the Web Server role on LON-CORE.

You also need to configure the World Wide Web Publishing service on LON-CORE.

The main tasks for this exercise are as follows:

1. Create a server group.
2. Deploy features and roles to both servers.
3. Review services and change a service setting.

► **Task 1: Create a server group**

1. Sign in to LON-DC1 with the **Administrator** account and the password **Pa\$\$w0rd**.
2. In the Server Manager console, click **Dashboard**, and then click **Create a server group**.
3. Click the **Active Directory** tab, and then click **Find Now**.
4. In the **Server group** name box, type **LAB-1**.
5. Add **LON-CORE** and **LON-SVR3** to the server group.
6. Click **LAB-1**. Select both **LON-CORE** and **LON-SVR3**.
7. Scroll down, and under the **Performance** section, select both **LON-CORE** and **LON-SVR3**.
8. Right-click **LON-CORE**, and then click **Start Performance Counters**.

► **Task 2: Deploy features and roles to both servers**

1. In Server Manager on LON-DC1, click the **LAB-1** server group, right-click **LON-CORE**, and then click **Add Roles and Features**.
2. In the Add Roles and Features Wizard, click **Next**, click **Role-based or feature-based installation**, and then click **Next**.
3. Verify that **LON-CORE.Adatum.com** is selected, and then click **Next**.
4. Select the **Web Server (IIS)** Server role.
5. Select the **Windows Server Backup** feature.
6. Add the **Windows Authentication** role service, and then click **Next**.
7. Select the **Restart the destination server automatically if required** check box, and then click **Install**.
8. Click **Close**.
9. Right-click **LON-SVR3**, click **Add Roles and Features**, and then click **Next**.
10. In the Add Roles and Features Wizard, click **Role-based or feature-based installation**, and then click **Next**.
11. Verify that **LON-SVR3.Adatum.com** is selected, and then click **Next** twice.
12. Click **Windows Server Backup**, and then click **Next**.
13. Select the **Restart the destination server automatically if required** check box, click **Install**, and then click **Close**.
14. In Server Manager, click the **IIS** node, and verify that LON-CORE is listed.

► **Task 3: Review services and change a service setting**

1. Sign in to LON-CORE with the **Adatum\Administrator** account and the password **Pa\$\$w0rd**.
2. In the Command Prompt window, execute the following command:

```
netsh.exe advfirewall firewall set rule group="remote desktop" new enable=yes
```

3. Sign in to LON-DC1 with the **Adatum\Administrator** account.
4. In Server Manager, click **LAB-1**, right-click **LON-CORE**, and then click **Computer Management**.
5. Expand **Services and Applications**, and then click **Services**.
6. Verify that the **Startup type** of the **World Wide Web Publishing** service is set to **Automatic**.

7. Verify that the service is configured to use the **Local System account**.
8. Configure the following service recovery settings:
 - First failure: **Restart the Service**
 - Second failure: **Restart the Service**
 - Subsequent failures: **Restart the Computer**.
 - Reset fail count after: **1** days
 - Reset service after: **1** minute
9. Configure the Restart Computer option to **2 minutes**, and then close the **World Wide Web Publishing Services Properties** dialog box.
10. Close the Computer Management console.

Results: After completing this exercise, you should have created a server group, deployed roles and features, and configured the properties of a service.

Exercise 4: Using Windows PowerShell to Manage Servers

Scenario

The marketing application vendor has indicated that they can provide some Windows PowerShell scripts to configure the web server that is hosting the application. You need to verify that remote administration is functional before running the scripts.

The main tasks for this exercise are as follows:

1. Use Windows PowerShell to connect remotely to servers and view information.
2. Use Windows PowerShell to remotely install new features.

► Task 1: Use Windows PowerShell to connect remotely to servers and view information

1. Sign in to LON-DC1 with the **Adatum\Administrator** account and the password **Pa\$\$w0rd**.
2. On LON-DC1, in Server Manager, click the **LAB-1** server group.
3. Right-click **LON-CORE**, and then click **Windows PowerShell**.
4. Type **Import-Module ServerManager**.
5. Type **Get-WindowsFeature**, and review roles and features.
6. Use the following command to review the running services on LON-CORE:

```
Get-service | where-object {$_.status -eq "Running"}
```

7. Type **get-process** to view a list of processes on LON-CORE.
8. Review the IP addresses assigned to the server by typing the following command:

```
Get-NetIPAddress | Format-table
```

9. Review the most recent 10 items in the security log by typing the following command:

```
Get-EventLog Security -Newest 10
```

10. Close Windows PowerShell.

► Task 2: Use Windows PowerShell to remotely install new features

1. On LON-DC1, on the taskbar, click the **Windows PowerShell** icon.
2. Type the following command to verify that the **XPS Viewer** feature has not been installed on LON-SVR3

```
Get-WindowsFeature -ComputerName LON-SVR3
```

3. To deploy the XPS Viewer feature on LON-SVR3, type the following command, and then press Enter:

```
Install-WindowsFeature XPS-Viewer -ComputerName LON-SVR3
```

4. Type the following command to verify that the XPS Viewer feature has now been deployed on LON-SVR3:

```
Get-WindowsFeature -ComputerName LON-SVR3
```

5. In the Server Manager console, in the **Tools** drop-down menu, click **Windows PowerShell ISE**.
6. In the Untitled1.ps1 script pane, type the following:

```
Import-Module ServerManager  
Install-WindowsFeature WINS -ComputerName LON-SVR3  
Install-WindowsFeature WINS -ComputerName LON-CORE
```

7. Save the script as **InstallWins.ps1** in a new folder named **Scripts**.
8. Press the F5 key to execute InstallWins.ps1.

Results: After completing this exercise, you should have used Windows PowerShell to perform a remote installation of features on multiple servers.

► Prepare for the next module

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

1. On the host computer, switch to the Hyper-V Manager console.
2. In the **Virtual Machines** list, right click **20410C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for 20410C-LON-CORE and 20410C-LON-SVR3.

Lab Review Questions

Question: What IP address range is used by the computers in the lab?

Question: Why must you set the DNS server address prior to joining the domain?

Question: Besides **sconfig.cmd**, what other tool can you use to rename a computer running the Server Core operating system?

Module Review and Takeaways

Module Review Questions

Question: What is the benefit of using Windows PowerShell to automate common tasks?

Question: What are the advantages to performing a Server Core deployment compared to the full GUI deployment?

Question: What tool can you use to determine which cmdlets are contained in a Windows PowerShell module?

Question: Which role can you use to manage KMS?

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
WinRM connections fail.	
Windows PowerShell cmdlets not available.	
Cannot install the GUI features on Server Core deployments.	
Unable to restart a computer running Server Core.	
Unable to join the domain.	

Tools

Tool	Use	Where to find it
Windows PowerShell	Performing multiple administrative and configuration tasks	taskbar
DISM.exe	Image servicing and management	From Search on the Start screen
Djoin.exe	Performing offline domain joins of computers	From Search on the Start screen
RSAT	Managing servers remotely from a Windows 8.1 system	Download from the Microsoft download center and install
Windows Server Migration Tools	Assisting with migrating to different versions of Windows Server	Download from the Microsoft download center and install
DiskPart	Managing physical disks	From the command prompt using the DiskPart command

Module 2

Introduction to Active Directory Domain Services

Contents:

Module Overview	2-1
Lesson 1: Overview of AD DS	2-2
Lesson 2: Overview of Domain Controllers	2-12
Lesson 3: Installing a Domain Controller	2-18
Lab: Installing Domain Controllers	2-26
Module Review and Takeaways	2-30

Module Overview

Active Directory® Domain Services (AD DS) and its related services form the foundation for enterprise networks that run Windows® operating systems. The AD DS database is the central store of all the domain objects, such as user accounts, computer accounts, and groups. AD DS provides a searchable hierarchical directory, and provides a method for applying configuration and security settings for objects in the enterprise. This module covers the structure of AD DS and its various components, such as forest, domain, and organizational units (OUs).

The process of installing AD DS on a server has been refined and improved with Windows Server® 2012 compared to the process of installing AD DS with earlier Windows server operating systems. This module examines some of the choices that are available with Windows Server 2012 for installing AD DS on a server.

Objectives

After completing this module, you will be able to:

- Describe the structure of AD DS.
- Describe the purpose of domain controllers.
- Install a domain controller.

Lesson 1

Overview of AD DS

The AD DS database stores information on user identity, computers, groups, services and resources. AD DS domain controllers also host the service that authenticates user and computer accounts when they sign in to the domain. Because AD DS stores information about all of the objects in the domain, and all users and computers must connect to AD DS domain controllers when they are signing into the network, AD DS is the primary means by which you can configure and manage user and computer accounts on your network.

This lesson covers the core logical components and physical components that make up an AD DS deployment.

Lesson Objectives

After completing this lesson you will be able to:

- Describe the components of AD DS.
- Describe AD DS domains.
- Describe OUs and their purpose.
- Describe AD DS forests and trees, and explain how you can deploy them in a network.
- Explain how an AD DS schema provides a set of rules that manage the objects and attributes that are stored in the AD DS domain database.
- Describe what is new for Active Directory in Windows Server 2012 R2.

Overview of AD DS

AD DS is composed of both logical and physical components. You need to understand the way the components of AD DS work together so that you can manage your infrastructure efficiently. In addition, you can use many other AD DS options to perform actions such as installing, configuring, and updating apps, managing the security infrastructure, enabling Remote Access and DirectAccess, and issuing and managing digital certificates.

One of the most used AD DS features is Group Policy, which enables you to configure centralized policies that you can use to manage most objects in AD DS. Understanding the various AD DS components is important to successfully using Group Policy.

Active Directory Domain Services (AD DS) is composed of both logical and physical components

Logical components	Physical components
<ul style="list-style-type: none"> • Partitions • Schema • Domains • Domain trees • Forests • Sites • Organizational units (OUs) • Containers 	<ul style="list-style-type: none"> • Domain controllers • Data stores • Global catalog servers • Read-only domain controllers (RODC)

Logical Components

AD DS logical components are structures that you use to implement an Active Directory design that is appropriate for an organization. The following table describes the types of logical structures that an Active Directory database contains.

Logical component	Description
Partition	A section of the AD DS database. Although the database is one file named Ntds.dit, it is viewed, managed, and replicated as if it consisted of distinct sections or instances. These are called <i>partitions</i> , which are also referred to as <i>naming contexts</i> .
Schema	The set of definitions of the object types and attributes that are used to create objects in AD DS.
Domain	A logical, administrative container for users and computers.
Domain tree	A collection of domains that share a common root domain and a contiguous Domain Name System (DNS) namespace.
Forest	A collection of domains that share a common AD DS.
Site	A collection of users, groups, and computers as defined by their physical location. Sites are useful in planning administrative tasks such as replication of changes to the AD DS database.
Organizational unit (OU)	An organizational unit is a container object that provides a framework for delegating administrative rights and for linking Group Policy Objects (GPOs).
Container	A container is an object that provides an organizational framework for use in AD DS. Containers cannot have GPOs linked to them.

Physical Components

The following table lists some of the physical components of AD DS and gives a brief description of each.

Physical component	Description
Domain controller	Contains a copy of the AD DS database. For most operations, each domain controller can process changes and replicate the changes to all the other domain controllers in the domain.
Data store	The files on each domain controller that holds the AD DS database. The Ntds.dit file, and associated log files, is a Microsoft JET database, which is stored in the C:\Windows\NTDS folder by default.
Global catalog server	A domain controller that hosts the <i>global catalog</i> , which is a partial, read-only copy of all the objects in the forest. A global catalog speeds up searches for objects that might be stored on domain controllers in a different domain in the forest.
Read-only domain controller (RODC)	A special read-only installation of AD DS. These are often used in branch offices where security and IT support are less advanced than in the main corporate centers.



Additional Reading: For more information about domains and forests, see Active Directory Domain Services Overview at <http://go.microsoft.com/fwlink/?LinkID=331086>.

What Are AD DS Domains?


The AD DS Domain Contains User, Computers, Groups

An AD DS domain is a logical container used to manage user, computer, group, and other objects. All of the domain objects are stored in the AD DS database, a copy of which is stored on each domain controller.

There are many types of objects in the AD DS database, including user accounts, computer accounts, and groups. The following list gives a brief description of these three object types:

- User accounts. User accounts contain the required information to authenticate a user during the sign-in process and to build the users access token.
- Computer accounts. Each domain-joined computer has an account in AD DS. Computer accounts are used in the same matter as user accounts, only for the domain-joined computers.
- Groups. Groups are used to organize users or computers to make it easier to manage permissions and group policy in the domain.

- AD DS requires one or more domain controllers
- All domain controllers hold a copy of the domain database which is continually synchronized
- The domain is the context within which user accounts, computer accounts, and groups are created
- The domain is a replication boundary
- The domain is an administrative center for configuring and managing objects
- Any domain controller can authenticate any sign-in anywhere in the domain
- The domain provides authorization



The AD DS Domain is a Replication Boundary

When changes are made to any object in the domain, the domain controller where the change occurred replicates that change to all the other domain controllers in the domain. If there are multiple domains in the forest, only subsets of the changes are replicated to other domains. AD DS uses a multi-master replication model that allows every domain controller to make changes to objects in the domain. Changes to relative ID (RID) management in the Windows Server 2012 version of Active Directory Domain Services (Windows Server 2012 Active Directory now allow a single domain to contain nearly 2 billion objects. With this capacity, most organizations could deploy only a single domain although organizations that have decentralized administrative structures, or that are distributed across multiple locations, might consider implementing multiple domains in the same forest.

The AD DS Domain is an Administrative Center

It contains an Administrator account and a Domain Admins group. By default the Administrator account is a member of the Domain Admins group, and the Domain Admins group is a member of every local Administrators group of domain-joined computers. Also, by default, the members of the Domain Admins group have full control over every object in the domain. The Administrator account in the forest root domain has additional rights as detailed in the "What Is an AD DS Forest?" topic.

The AD DS Domain Provides Authentication

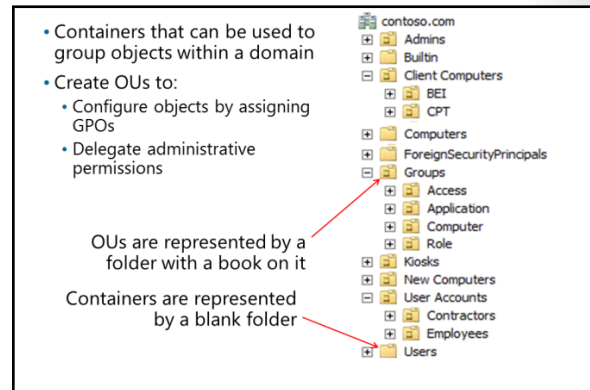
Whenever a domain-joined computer starts up, or a user signs in to a domain-joined computer, AD DS authenticates them. Authentication verifies the computer or user has the proper credentials for an AD DS account.

The AD DS Domain Provides Authorization

Windows operating systems use authorization and access control technologies to allow authenticated users to access resources. Typically the authorization process is performed locally at the resource. Windows Server 2012 introduced domain-based Dynamic Access Control to allow for central access rules to control access to resources. Central access rules do not replace the current access control technology; it just provides an additional level of control.

What Are OUs?

An *organizational unit* (OU) is a container object within a domain that you can use to consolidate users, computers, groups, and other objects. These should not be confused with the generic container objects in AD DS. The primary difference between OUs and containers are the management capabilities. Containers have limited management capabilities, for instance you cannot directly apply a GPO to a container. Typically containers are used for system objects and as the default locations for new objects. With OUs you have more management options; you can directly link GPOs, assign an OU manager and associate a COM+ partition with an OU.



While there is not a menu option for creating new containers in Active Directory Users and Computers, you can create new OUs in AD DS at any time. There are two reasons to create OUs:

- To configure objects in the OU. You can assign GPOs to the OU, and the settings apply to all objects within the OU. GPOs are policies that administrators create to manage and configure settings for computers and/or users. The GPOs are deployed by linking them to OUs, Domains or Sites.
- To delegate administrative control of objects within the OU. You can assign management permissions on an OU, thereby delegating control of that OU to a user or group within AD DS in addition to the administrators group.

You can use OUs to represent the hierarchical, logical structures within your organization. For example, you can create OUs that represent the departments within your organization, the geographic regions within your organization, or a combination of both departmental and geographic regions. You can use OUs to manage the configuration and use of user, group, and computer accounts based on your organizational model.

Every AD DS domain contains a standard set of containers and OUs that are created when you install AD DS, some of the default objects are used primarily by AD DS and are hidden from view by default. The follow list includes the objects that are visible by default:

- Domain. Serves as the top level of the domain organizational hierarchy.
- Builtin container. Stores a number of default groups.
- Computers container. The default location for new computer accounts that you create in the domain.
- Domain Controllers OU. The default location for the computer accounts for domain controller computer accounts. This is the only OU that is present in a new installation of AD DS.
- ForeignSecurityPrincipals. The default location for trusted objects from domains outside the AD DS forest. These are typically created when an objected from an external domain is added to a group in the AD DS domain.
- Managed Service Accounts. The default location for managed service accounts. A managed service account is domain account where AD DS provides automatic password management.
- Users container. The default location for new user accounts and groups that you create in the domain. The users container also holds the administrator and guest accounts for the domain, and some default groups.

There are several containers that can only be seen when the Advanced Features is selected from the View menu. The following list includes the objects that are hidden by default:

- **LostAndFound.** This container holds objects that have been orphaned.
- **Program Data.** This container holds Active Directory data for Microsoft applications, such as Active Directory Federation Services (AD FS).
- **System.** This container holds the built-in system settings.
- **NTDS Quotas.** This container holds directory service quota data.
- **TPM Devices.** This container is new with Windows Server 2012. It stores the recovery information for Trusted Platform Module (TPM) devices.



Note: Containers in an AD DS domain cannot have GPOs linked to them. To link GPOs to apply configurations and restrictions, create a hierarchy of OUs, and then link GPOs to them.

Hierarchy Design

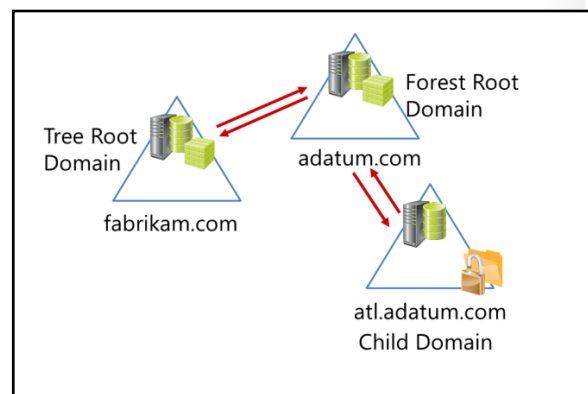
The design of an OU hierarchy is dictated by the administrative needs of the organization. The design could be based on geographic, functional, resource, or user classifications. Whatever the order, the hierarchy should make it possible to administer AD DS resources as effectively and with as much flexibility as possible. For example, if all computers that IT administrators use must be configured in a certain way, you can group all the computers in an OU, and then assign a GPO to manage those computers.

You also can create OUs within other OUs. For example, your organization might have multiple offices, and each office might have a set of IT administrators who are responsible for managing user and computer accounts in their office. In addition, each office might have different departments with different computer configuration requirements. In this situation, you could create an OU for each office, and then within each of those OUs, create an OU for the IT administrators and OUs for each of the other departments.

Although there is no technical limit to the number of levels in your OU structure, for the purpose of manageability, limit your OU structure to a depth of no more than 10 levels. Most organizations use five levels or fewer to simplify administration. Note that Active Directory–enabled applications can impose restrictions on the OU depth within the hierarchy for the parts of the hierarchy they use.

What Is an AD DS Forest?

A domain *tree* is a collection of one or more domains that share a contiguous name space. A *forest* is a collection of one or more domain trees that share a common directory schema and global catalog. The first domain that is created in the forest is called the *forest root domain*. The *forest root domain* contains a few objects that do not exist in other domains in the forest. Since these objects are always created on the very first domain controller created, a forest can consist of as little as one domain with a single domain controller, or



it can consist of hundreds of domains across multiple trees. These objects that only exist in the forest root domain are as follows:

- The schema master role. This is a special forest-wide domain controller role. There is only one schema master in any forest. The schema can be changed only on the domain controller which holds the schema master.
- The domain naming master role. This is also special forest-wide domain controller role. There is only one domain naming master in any forest. New domain names can be added to the directory only by the domain naming master.
- The Enterprise Admins group. By default, the Enterprise Admins group has the Administrator account for the forest root domain as a member. The Enterprise Admins group is a member of the local administrators group in every domain in the forest. This allows members of the Enterprise Admins group to have full control administrative rights to every domain throughout the forest.
- The Schema Admins group. By default, the Schema Admins group has no members. Only members of the Enterprise Admins group can add members to the Schema Admins group. Members of the Schema Admins group are the only administrators that can make changes to the Schema.

Security Boundary

An AD DS forest is a security boundary. By default, no users from outside the forest can access any resources inside the forest. Typically an organization will create only one forest, although you can create multiple forests to isolate administrative permissions between different parts of the organization.

By default, all the domains in a forest automatically trust the other domains in the forest. This makes it easy to enable access to resources such as file shares and websites for all users in a forest, regardless of the domain in which the user account is located.

Replication Boundary

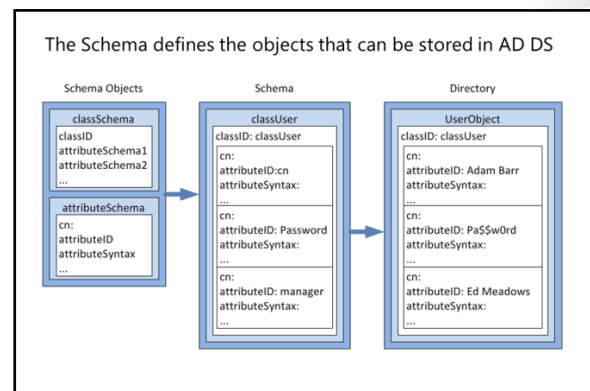
An AD DS forest is the replication boundary for the configuration and schema partitions in the AD DS database. This means that all domain controllers in the forest must share the same schema. Because of this organizations that want to deploy applications with incompatible schemas will need to deploy additional forests.

The AD DS forest is also the replication boundary for the global catalog. The global catalog provides the ability to find objects from any domain in the forest. The global catalog is used whenever universal principal name (UPN) sign-in credentials are used or when looking up users in Microsoft Exchange Server address books.

What Is the AD DS Schema?

The *AD DS schema* is the component that defines all object classes and attributes that AD DS uses to store data. It is sometimes referred to as the blueprint for AD DS. The schema is replicated among all domain controllers in the forest. Any change that is made to the schema is replicated to every domain controller in the forest from the schema master holder, typically the first domain controller in the forest.

AD DS stores and retrieves information from a wide variety of applications and services; it does this, in part, by standardizing how data is stored in



the AD DS directory. By standardizing how data is stored, AD DS can retrieve, update, and replicate data, while ensuring that the integrity of the data is maintained.

Objects

AD DS uses objects as units of storage. All object types are defined in the schema. Each time that the directory handles data, the directory queries the schema for an appropriate object definition. Based on the object definition in the schema, the directory creates the object and stores the data.

Object definitions specify both the types of data that the objects can store, and the syntax of the data. You can only create objects in AD DS that are defined by the schema. Because the data is stored in a rigidly defined format, AD DS can store, retrieve, and validate the data that it manages, regardless of which application the data came from.

Relationships between Objects, Rules, Attributes, and Classes

In AD DS, the schema defines the following:

- Objects that are used to store data in the directory
- Rules that define the structure of the objects
- Structure and content of the directory itself

AD DS schema objects consist of attributes, which are grouped together into classes. Each class has rules that define which attributes are required and which are optional. For example, the user class consists of over 400 possible attributes, including **cn** (the common name attribute), **givenName**, **displayName**, **objectSID** and **manager**. Of these attributes the **cn** and **objectSID** attributes are mandatory, and the **cn** attribute is defined as a single value Unicode String from 1 to 64 characters long and is replicated to the global catalog.

Making Changes to the Schema

Only members of the Schema Administrators can modify the AD DS schema. You cannot remove anything from the AD DS schema; you can only extend the AD DS schema using AD DS schema extensions, or modify the attributes of existing objects. For example, when you are preparing to install Exchange Server 2013 you must apply the Exchange Server 2013 Schema Extensions. This extension adds or modifies over 200 classes and over 100 different attributes.

Because the schema dictates how information is stored, and because any changes that are made to the schema affect every domain controller, changes to the schema should be made only when necessary. Before making any changes, you should review the changes through a tightly controlled process, and then implement them only after you have performed testing to ensure that the changes will not adversely affect the rest of the forest and any applications that use AD DS.

The schema master is one of the flexible single master operations (FSMO) roles that are hosted on a single domain controller in AD DS. Because it is a single master, you must make changes to the schema by targeting the domain controller that holds the schema master.

What Is New for Windows Server 2012 Active Directory?

Windows Server 2012 Active Directory (Windows Server 2012 AD) includes many enhancements and improvements from previous versions. Some of these enhancements have been added to help manage the proliferation of consumer devices in the workplace. For example, to provide users with an easier way to integrate their consumer devices into the workplace, the new features Workplace Join and Web Application Proxy were implemented. Also, the security around the use of consumer devices in the workplace has been improved; multi-factor access control and multi-factor authentication were implemented to manage the risk associated with allowing consumer devices to participate in the domain.

Improvements for using consumer devices in the enterprise:

Workplace Join

- Allows consumer devices to be joined to the domain

Web Application Proxy

- Allows applications to be published to the Internet

Multi-Factor Access Control

- Allows claims using different factors

Multi-Factor Authentication

- Allows you to specify the use of multiple factors for authentication

Workplace Join

Windows Server 2012 R2 allows users to join their personal devices to the domain. Both Windows-based devices and iOS®-based devices can join a Windows Server 2012 R2-based domain. Using a feature from Active Directory Federation Services (AD FS) called Device Registration Service (DRS) a user's personal device can be registered in AD DS. The DRS service will create an AD DS object for the device and issue a certificate to the device that is used to authenticate the device. If both DRS and the Web Application Proxy are used, any device with a working internet connection can be workplace joined.

When a personal device is workplace joined administrators can:

- Use the information about the device that is stored in AD DS and configure conditional access.
- Provide a seamless experience to users who are accessing company resources from workplace joined devices.
- Provide a single sign on (SSO) experience for accessing resources.

Web Application Proxy

Web Application Proxy is a new Remote Access role service which you can use to provide external users access to apps and running on internal servers from anywhere at any time. Web Application Proxy can be used with personal devices that have been workplace joined as well as company owned laptops, smartphones and other devices. Web Application Proxy gives administrators more detailed control than a traditional virtual private network (VPN), because users can only access apps that have been published to them through the Web Application Proxy. You can apply additional security through Web Application Proxy to further control the apps the user can access. Web Application Proxy requires the use of AD FS and uses AD FS features such as SSO.

Multi-Factor Access Control

Workplace Join and Web Application Proxy both use AD FS; the primary function of AD FS is to issue access tokens that contain claims. Claims-based authentication is used extensively in cloud-based apps and services. Claims-based authentication is similar to the traditional authentication process used in a Windows domain; the primary functional difference is that the claims-based security token only includes who the user is and does not define anything they can do. AD FS evaluates claims requests that can be

based on one or more factors. An app that uses claims-based authentication is also known as *relying party application*. There are over 50 factors that AD FS can use to authenticate a claim request, some of those factors are:

Claim type	Description
Email Address	The user's email address.
Name	The user's name, which must be unique.
Role	A role that the user has been assigned.
Primary group security identifier (SID)	The primary group SID of the user.
Issuer	The name of the certificate authority that issued the X.509 certificate.

The relying party application defines what the user will be able to do based on the information in the claim.

There are several benefits of using multi-factor access control in AD FS, including:

- You can permit or deny access based on the user, device, location, authentication state or other factors by using the flexible and granular per-application authorization policies.
- You can create different rules for each application by using the individual issuance authorization rules for relying party applications.
- You can deliver a rich UI experience for the common multi-factor scenarios to users by using AD FS's web-based authentication with customizable forms for some common scenarios.
- For more complex scenarios you can use Windows PowerShell® to develop your rules by using the rich claims language and Windows PowerShell support.
- You can tell the users why their request was denied and not just display a generic access denied message by using individual, customized access denied messages for relying party applications.



Additional Reading: For more information about how to manage risk with multi-factor access control, see the article " Overview: Manage Risk with Multi-Factor Access Control" at <http://go.microsoft.com/fwlink/?LinkID=331088>.

Multi-Factor Authentication

AD FS has multiple authentication methods which you can use to create flexible authentication scenarios; these flexible authentication scenarios allow your users to access company resources in more ways. You can create a global authentication policy that applies to all access attempts or you can create custom authentication rules for individual, AD FS-secured resources. Custom, per-relying party application authentication rules do not override global authentication rules. You can configure the authentication rules to require only a primary authentication method, or to use multi-factor authentication. When creating both global and per-relying party application authentication rules, if either the global rule or application specific rule requires the use of multi-factor authentication, then the user will be required to use multi-factor authentication.

When you create a global authentication policy you can configure the following settings:

- Primary authentication method. By default, external connections use Forms Authentication and internal connections use Windows Authentication.
- Settings and methods for multi-factor authentication. You can configure the conditions under which multi-factor authentication is used, and what additional authentication method is used. You can use Certificate Authentication, such as with a smart card, or other third-party authentication methods.
- Whether device authentication is enabled. This option is used with Workplace Join and allows you to configure the device as a secondary authentication factor.

When configuring per-relying party application authentication rules you can configure the following settings:

- Whether the users need to provide credentials each time they sign in
- Multi-factor authentication settings for the relying party application

The parameters on which multi-factor authentication rules can be based include the following:

- Users or groups in the AD DS directory
- The workplace joined status of devices
- The connection being from the intranet or the Internet

A Work from Home Scenario

These features work together to greatly enhance a work from home scenario. In a typical work from home scenario that does not use the new features described above, users are connected through a VPN to the corporate network. They may be required to use two factor authentication, such as a smart card, but that only protects the connection. Once they are connected they can access anything they could access from within the corporate network.

Security departments generally have several concerns about work from home programs:

- Are the users accessing the corporate network from a secure system or are they on a less secure system that can be found, for instance, in public libraries?
- Do the users have access to sensitive files and are they downloading them to their local system?
- How are users accessing line of business (LOB) applications?
- What are their screen saver settings, if they walk away from their computer how easy is it for someone else to walk up and use it?

When you use the new Windows Server 2012 R2 features you can allow the users to work from home while you still maintain secure control over what they can access. The users use Workplace Join to add their personal systems to the domain; then security settings can be configured for several different scenarios. For example, certain files might be configured to only be accessible from the users workplace joined computer, or sensitive files might be configured to only be accessible from domain joined systems.

Similarly, LOB apps can be published through the Web Application Proxy, using the claims defined through multi-factor access control to specify what the users are allowed to do in the apps. Additionally multi-factor authentication can be specified for certain apps to help ensure the appropriate user is running the apps.

Lesson 2

Overview of Domain Controllers

Because domain controllers authenticate all users and computers in the domain, domain controller deployment is critical to the correct functioning of the network.

This lesson examines domain controllers, the sign-in process, and the importance of DNS in that process. In addition, this lesson discusses the purpose of the global catalog.

All domain controllers are essentially the same, with two exceptions. RODCs contain a read-only copy of the AD DS database, while other domain controllers have a read/write copy. There are also certain operations that can only be performed on specific domain controllers called *operations masters*, which are discussed at the end of this lesson.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the purpose of domain controllers.
- Describe the purpose of the global catalog.
- Describe the AD DS sign-in process, and the importance of DNS and SRV records in the sign-in process.
- Describe the functionality of SRV records.
- Explain the functions of operations masters.

What Is a Domain Controller?

A *domain controller* is a server that is configured to store a copy of the AD DS directory database (Ntds.dit) and a copy of the SYSVOL folder. All domain controllers except RODCs store a read/write copy of both Ntds.dit and the SYSVOL folder. Ntds.dit is the database itself, and the SYSVOL folder contains all the template settings and files for GPOs.

Domain controllers use a multi-master replication process; for most operations, data can be modified on any domain controller, except on RODCs. The AD DS replication service then synchronizes the changes that have been made to the AD DS database to all other domain controllers in the domain. The SYSVOL folders are replicated either by the File Replication Service (FRS), or by the newer Distributed File System (DFS) Replication.

Domain controllers host several other Active Directory–related services, including the Kerberos authentication service, which is used by User and Computer accounts for sign-in authentication, and the Key Distribution Center (KDC), which issues the ticket-granting tickets (TGTs) to an account that signs in to the AD DS domain. You can optionally configure domain controllers to host a copy of the global catalog.

All users in an AD DS domain exist in the AD DS database, if the database is unavailable for any reason all operations depending on domain-based authentication will fail. As a best practice, an AD DS domain should have a minimum of two domain controllers. This supports high availability of the AD DS database and spreads the authentication load during peak sign-in times.

Domain Controllers

Servers that host the AD DS database (Ntds.dit) and SYSVOL

Kerberos authentication service and Key Distribution Center (KDC) services perform authentication

Best practices:

- Availability:
At least two domain controllers in a domain
- Security:
RODC and BitLocker



Note: Two domain controllers should be considered an absolute minimum.

When you deploy a domain controller in a branch office where physical security is less than optimal, there are some additional measures that you can use to reduce the impact of a breach of security. One option is to deploy an RODC.

The RODC contains a read-only copy of the AD DS database, and by default, it does not cache any user passwords. You can configure the RODC to cache the passwords for users in the branch office. If an RODC is compromised, the potential loss of information is much lower than with a full read/write domain controller. Another option is to use Windows BitLocker® Drive Encryption to encrypt the domain controller hard drive. If the hard drive is stolen, BitLocker encryption ensures that there is a very low chance of a malicious user getting any useful information from it.



Note: BitLocker is a drive encryption system that is available for Windows Server® operating systems, and for certain Windows client operating system versions. BitLocker securely encrypts the entire operating system so that the computer cannot start without being supplied a private key and (optionally) passing an integrity check. A disk remains encrypted even if you transfer it to another computer.

What Is the Global Catalog?

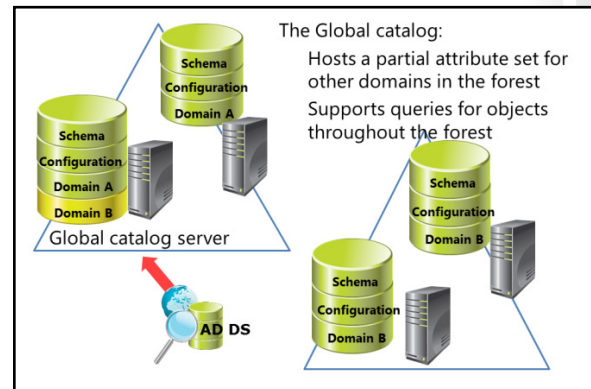
The *global catalog* is a partial, read-only, searchable copy of all the objects in the forest. It speeds up searches for objects that might be stored on domain controllers in a different domain in the forest.

Within a single domain, the AD DS database on each domain controller contains all the information about every object in that domain, but only a subset of this information is replicated to global catalog servers in other domains in the forest. Within a given domain, a query for an object is directed to one of the domain controllers

in that domain, but that query cannot provide any results for objects in other domains in the forest. For a query to include results from other domains in the forest you must query a domain controller that is a global catalog server. By default, the only global catalog server that is hosted is the first domain controller in the forest root domain. To enhance searching across domains in a forest, you should configure additional domain controllers to store a copy of the global catalog.

The global catalog does not contain all attributes for each object. Instead, the global catalog maintains the subset of attributes that are most likely to be useful in cross-domain searches. These attributes include **givenName**, **displayName**, and **mail**.

There are a variety of reasons why you might perform a search against a global catalog rather than a domain controller that is not a global catalog. For example, when a server that is running Exchange Server receives an incoming email, it needs to search for the recipient's account so that it can decide how to route the message. By automatically querying a global catalog, the server that is running Exchange Server is able to locate the recipient in a multiple domain environment. In another example, when a user signs in to their Active Directory account, the domain controller that is performing the authentication must contact a global catalog to check for universal group memberships before the user is authenticated.



In a single domain, all domain controllers should be configured to hold a copy of the global catalog; however, in a multiple domain environment, the infrastructure master should not be a global catalog server unless all the domain controllers in the domain are also global catalog servers. Which domain controllers are configured to hold a copy of the global catalog depends on replication traffic and network bandwidth. Many organizations are opting to make every domain controller a global catalog server.


Question: Should a domain controller be a global catalog?

The AD DS Sign-in Process

When a user signs in to AD DS, their system looks in DNS for service resource (SRV) records to locate the nearest suitable domain controller. *SRV records* specify information about available services, and are recorded in DNS for all domain controllers. By using DNS lookups, clients can locate a suitable domain controller to service their sign-in requests. If the sign-in is successful, the local security authority (LSA) builds an access token for the user that contains the SIDs for the user and any groups of which the user is a member. The token provides the access credentials for any process initiated by that user. For example, after signing in to AD DS, a user runs Microsoft® Word and attempts to open a file. Word uses the credentials in the user's access token to verify the level of the user's permissions for that file.

The AD DS sign-in process:

1. The User Account is authenticated to the domain controller.
2. The domain controller returns a ticket-granting ticket (TGT) back to client.
3. The client uses TGT to apply for access to the workstation.
4. The domain controller grants access to the workstation.
5. The client uses TGT to apply for access to the server.
6. The domain controller returns access to the server.



Note: A SID is a unique string in the form of S-R-X-Y1-Y2-Yn-1-Yn. For example, a user SID could be: S-1-5-21-322346712-1256085132-1900709958-500.

The parts of this SID are explained in this table:

Component	Definition	In The Example
S	Indicates that the string is a SID	S
R	Revision level	1
X	Identifier authority value	5 (NT Authority)
Y1-Y2-Yn-1	Domain identifier	21-322346712-1256085132-1900709958
Yn	RID	500

Every user and computer account and every group that you create has a unique SID. They only differ from each other by virtue of the unique RID. The SID in the example is a well-known SID for the domain administrator account. Well-known SIDs are used by the default accounts and groups, the Domain Administrator account's SID always ends with 500.

Sites

Sites are used by a client system when it needs to contact a domain controller. It starts by looking up SRV records in DNS. The response to the DNS query will include:

- A list of the domain controllers in the same site as the client system
- If there are no domain controllers available in the site the list will include the domain controllers from the next closest site that does not include an RODC, if enabled in the Try Next Closest Site Group Policy setting
- If no domain controller is found in the next closest site, a random list of available domain controllers in the domain

Administrators can define sites in AD DS. Sites usually align with the parts of the network that have good connectivity and bandwidth. For example, if a branch office is connected to the main data center by an unreliable wide area network (WAN) link, it would be better to define the data center and the branch office as separate sites in AD DS.

SRV records are registered in DNS by the Net Logon service that is running on each domain controller. If the SRV records are not entered in DNS correctly, you can trigger the domain controller to reregister those records by restarting the Net Logon service on that domain controller. This process only reregisters the SRV records; if you want to reregister the host (A) record information in DNS, you must run **ipconfig /registerdns** from a command prompt, just as you would for any other computer.

Although the sign-in process appears to the user as a single event, it is actually made up of two parts:

- The user provides credentials, usually a user account name and password, which are then checked against the AD DS database. If the user account name and the password match the information that is stored in the AD DS database, the user becomes an authenticated user and is issued a TGT by the domain controller. At this point, the user does not have access to any resources on the network.
- A secondary process in the background submits the TGT to the domain controller and requests access to the local machine. The domain controller issues a service ticket to the user, who is then able to interact with the local computer. At this point in the process, the user is authenticated to AD DS and signed in to the local machine.

When a user subsequently attempts to connect to another computer on the network, the secondary process is run again, and the TGT is submitted to the nearest domain controller. When the domain controller returns a service ticket, the user can access the computer on the network, which generates a logon event at that computer.



Note: A domain-joined computer also logs on to AD DS when it starts—a fact that is often overlooked. You do not see the transaction when the computer uses its computer account name and a password to log on to AD DS. Once authenticated, the computer becomes a member of the Authenticated Users group. Although the computer log on event does not have any visual confirmation in a GUI, it is recorded in the event log. Also, if auditing is enabled, additional events are recorded in the Security Log of the Event Viewer.

Demonstration: Viewing the SRV Records in DNS

The demonstration shows you how to display the various types of SRV records that the domain controllers register in DNS. These records are crucial to how AD DS operates because they are used to find domain controllers for signing in, changing passwords, and editing GPOs. SRV records are also used by domain controllers to find replication partners.

Demonstration Steps

View the SRV records by using DNS Manager

1. On LON-DC1, sign in with the user account **Adatum\Administrator** and the password **Pa\$\$w0rd**.
2. Open the DNS Manager window, and explore the underscore DNS domains.
3. View the SRV records that are registered by domain controllers.

These records provide alternate paths so that clients can discover them.

What Are Operations Masters?

Although all domain controllers contain the same information, there is some information that that can only be modified on one particular domain controller. There are the five operations masters, or flexible single master operations (FSMO) roles. The FSMO roles help prevent conflicts in AD DS due to replication latency. When making changes to data of any of the FSMO roles you must connect to the domain controller that holds the role.

The five FSMO roles are distributed as follows:

- Each forest has one schema master and one domain naming master.
- Each AD DS domain has one RID master, one infrastructure master, and one primary domain controller (PDC) emulator.

In the multi-master replication model, some operations must be single master

Many terms are used for single master operations in AD DS, including the following:

- Operations master (or operations master roles)
- Single master roles
- Flexible single master operations (FSMOs)

The five FSMOs are:

- | | |
|--|---|
| <ul style="list-style-type: none"> • Forest: <ul style="list-style-type: none"> • Domain naming master • Schema master | <ul style="list-style-type: none"> • Domain: <ul style="list-style-type: none"> • RID master • Infrastructure master • PDC Emulator master |
|--|---|

Forest Operations Masters

The following are single master roles found in a forest:

- Domain naming master. This is the domain controller that must be contacted when you add or remove a domain, or when you make domain name changes.

If the domain naming master is unavailable, you will not be able to add additional domains to the forest.

- Schema master. This is the domain controller where all schema changes are made. To make changes you would typically sign in to the schema master as a member of both the Schema Admins and Enterprise Admins groups. A user who is a member of both of these groups and who has the appropriate permissions could also edit the schema by using a script.

If the schema master is unavailable, you will be unable to make changes to the schema; this would prevent installation of applications that required schema changes, such as Microsoft Exchange Server.

Domain Operations Masters

The following are single master roles found in a domain:

- Relative ID (RID) master. Whenever an object is created in AD DS, the domain controller where the object is created assigns the object a unique identifying number known as a SID. To ensure that no two domain controllers assign the same SID to two different objects, the RID master allocates blocks of RIDs to each domain controller within the domain to use when building the SID.

If the RID master is unavailable, you can experience difficulties adding new objects to the domain. As domain controllers use their existing RID's they will eventually run out of RID's and be unable to create new objects.

- Infrastructure master. This role maintains inter-domain object references, such as when a group in one domain contains a member from another domain. In this situation, the infrastructure master is responsible for maintaining the integrity of this reference. For example, when you look at the security tab of an object, the system looks up the SIDs that are listed and translates them into names. In a multi-domain forest, the infrastructure master looks up SIDs from other domains.

If the infrastructure master is unavailable, domain controllers that are not global catalogs will be unable to check universal group memberships and will be unable to authenticate users.

The infrastructure role should not reside on a global catalog server, unless you have a single domain forest. The exception is when you follow best practices and make every domain controller a global catalog. In that case, the infrastructure role is not required because every domain controller knows about every object in the forest.

- Primary domain controller (PDC) emulator master. The domain controller that holds the PDC emulator is the time source for the domain. The PDC emulators in each domain in a forest synchronize their time with the PDC emulator in the forest root domain. You set the PDC emulator in the forest root domain to synchronize with an external atomic time source.

The PDC emulator is also the domain controller that receives urgent password changes. If a user's password is changed, the information is sent immediately to the domain controller holding the PDC emulator. This means that if the user subsequently tried to sign in and they were authenticated by a domain controller in a different location that had not yet received an update about the new password, the domain controller in the location to which the user tried to sign in would contact the domain controller holding the PDC emulator and check for recent changes.

If the PDC emulator is unavailable, users may have trouble signing on until their password change has replicated to all the domain controllers.

The PDC emulator is also used when editing GPOs. When a GPO other than a local GPO is opened for editing, the copy that is edited is the one stored on the PDC emulator. This is done to prevent conflicts if two administrators attempted to edit the same GPO at the same time on different domain controllers. However, you can choose to use a specific domain controller for editing GPOs. This is especially useful when editing GPOs in a remote office with a slow connection to the PDC emulator.



Note: The global catalog is not one of the operations master roles.



Note: The five FSMO roles are also known as:

- Schema operations master
- Domain naming operations master
- infrastructure operations master
- RID operations master
- PDC emulator operations master

Lesson 3

Installing a Domain Controller

Sometimes you need to install additional domain controllers in your Windows Server 2012 domain. There are several reasons you might do this:

- You need additional resources at a site because existing domain controllers are overworked.
- You are opening a new remote office that requires you to deploy one or more domain controllers.
- You are creating an off-site disaster recovery location.

The installation method that you use varies with the circumstances.

This lesson examines several ways to install additional domain controllers such as installing AD DS on a local machine and on a remote server by using Server Manager, installing AD DS on a Server Core installation, and installing AD DS by using a snapshot of the AD DS database that is stored on removable media. This lesson also examines how to upgrade a domain controller from an earlier Windows operating system to Windows Server 2012. Finally Windows Azure® Active Directory (Windows Azure AD) and how to install a domain controller in Windows Azure is discussed.

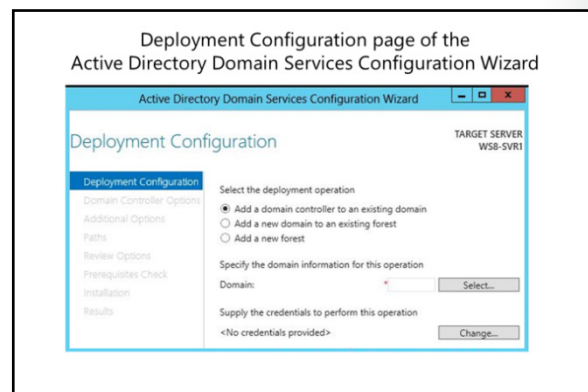
Lesson Objectives

After completing this lesson, you will be able to:

- Explain how to install a domain controller by using the GUI.
- Explain how to install a domain controller on a Server Core installation of Windows Server 2012.
- Explain how to upgrade a domain controller by using Install from Media.
- Explain how to install a domain controller by using Install from Media.
- Describe Windows Azure AD.
- Understand how to deploy domain controllers in Windows Azure.

Installing a Domain Controller from Server Manager

With Windows Server 2008 and earlier versions, it was common practice to start the Active Directory Domain Services Installation Wizard with the Dcpromo.exe tool to install domain controllers. But, beginning with Windows Server 2012, the Active Directory Domain Services Installation Wizard has been moved to Server Manager and the use of Dcpromo.exe is only supported for legacy automation. The domain controller promotion process is a two-step process. First you need to install the files that the domain controller role uses, then you install the domain controller role itself.



Note: The Active Directory Domain Services Installation Wizard (which can be opened from the command line by typing **dcpromo.exe**) is deprecated beginning in Windows Server 2012.

Before installing a new domain controller you need to have the answers to the following questions.

Question	Comments
Are you installing a new forest, a new tree or an additional domain controller for an existing domain?	Answering this question will determine what additional information you might need, such as the parent domain name.
What is the DNS name for the AD DS domain?	When you are creating the first domain controller for a domain you must specify the fully qualified domain name (FQDN). When you are adding a domain controller to an existing domain or forest the existing domain information will be provided in the wizard.
What will the forest functional level be set at?	The forest functional level determines the forest features that will be available and the supported domain controller operating system. This also sets the minimum domain functional level for the domains in the forest.
What will the domain functional level be set at?	The domain functional level determines the domain features that will be available and the supported domain controller operating system.
Will the domain controller be a DNS server?	Your DNS must be functioning well to support AD DS.
Will the domain controller host the global catalog?	This option is selected by default for the first domain controller in a forest and cannot be changed.
Will the domain controller be a RODC?	This option is not available for the first domain controller in a forest.
What will the Directory Services Restore Mode (DSRM) password be?	This is required to be able to recovery the active directory database from a backup.
What is the NetBIOS name for the AD DS domain?	When creating the first domain controller for a domain you must specify the NetBIOS name for the domain.
Where will the database, log files, and SYSVOL folders be created?	By default, the database and log files folder is C:\Windows\NTDS. By default the SYSVOL folder is C:\Windows\SYSVOL.

When you run Server Manager on the local system you install the AD DS role. At the end of the initial installation process, the AD DS files are installed but AD DS is not yet set up on that server.

To set up AD DS, you need to configure it by using the Active Directory Domain Services Configuration Wizard. You start the wizard by clicking the AD DS link in Server Manager. The wizard will allow you to do one of the following operations:

- Add a domain controller to an existing domain
- Add a new domain to an existing forest
- Add a new forest

The specific steps that you follow in the rest of the wizard depends on the operation you wish to complete.



Note: If you need to restore the AD DS database from a backup, restart the domain controller in DSRM. When the domain controller starts up it is not running the AD DS services, instead, it is running as a member server in the domain. To sign in to that server in the absence of AD DS, sign in using the Directory Services Recovery Mode password.



Note: Windows Server 2012 supports cloning AD DS servers. Before being cloned an AD DS sever must be a member of the Cloneable Domain Controllers group. Additionally the PDC emulator must be online and available to the cloned DC, and must be running Windows Server 2012.

Installing a Domain Controller on a Server Core Installation of Windows Server 2012

There are several options for configuring a Windows Server 2012 server that is running Server Core as a domain controller. Without the Server Manager GUI interface, you need to use alternate methods to install the files for the domain controller role and to install the domain controller role itself.

To install the AD DS files on the server you can do one of the following:

- Use Server Manager to connect remotely to the Server Core server and install the AD DS role as described in the previous topic.
- Use the Windows PowerShell command **Install-WindowsFeature AD-Domain-Services** to install the files.

Once you install the AD DS files, you can complete everything except the hardware the installation and configuration in one of the following ways:

- Use Server Manager to start the Active Directory Domain Services Configuration Wizard as described in the previous topic
- Load the Active Directory Domain Services Deployment module with the Windows PowerShell command **Import-Module ADDSDeployment**. Then run the Windows PowerShell command **Install-ADDSDomainController** and supply the required information on the command line.



Additional Reading: For complete details on using the Windows PowerShell command **Install-ADDSDomainController** see the article "Install Active Directory Domain Services (Level 100)" at <http://go.microsoft.com/fwlink/?LinkID=331087>.

Refer to the following site for details:

AD DS Deployment Cmdlets in Windows PowerShell,
<http://go.microsoft.com/fwlink/?LinkID=331089>

Installing AD DS is a two-step process regardless of which installation method you use.

- Method 1, use Server Manager on a Windows 2012 server with a GUI interface to connect to the system
 1. Install the files by installing the Active Directory Domain Services role
 2. Install the domain controller role by running the Active Directory Domain Services Configuration Wizard
- Method 2, Use Windows PowerShell locally
 1. Install the files by running the command **Install-WindowsFeature AD-Domain-Services**
 2. Install the domain controller role by running the command **Install-ADDSDomainController**

Upgrading a Domain Controller

You can upgrade to a Windows Server 2012 domain in one of two ways. You can upgrade the operating system on existing domain controllers that are running Windows Server 2008 or Windows Server 2008 R2. Alternatively, you can add Windows Server 2012 servers as domain controllers in a domain which already has domain controllers running previous versions of Windows Server. Of the two methods, the second is preferred because when you finish, you have a clean installation of the Windows Server 2012 operating system and the AD DS database.

Whenever a new domain controller is added the domain DNS records are updated and clients will find and use this domain controller immediately.

Options to upgrade AD DS to Windows Server 2012:

- In-place upgrade (from Windows Server 2008 or Windows Server 2008 R2)
 - Benefit: Except for the prerequisite checks, all the files and programs stay in-place and there is no additional work required
 - Risk: May leave legacy files and DLLs
- Introduce a new Windows Server 2012 server into the domain and promote it to be a domain controller
 - This option is usually the preferred choice
 - Benefit: The new server has no accumulated legacy files and settings
 - Risk: May need additional work to migrate administrators' files and settings

Upgrading to Windows Server 2012

To upgrade an AD DS domain that is running at an older Windows Server functional level to an AD DS domain running at Windows Server 2012 functional level, you must first upgrade all the domain controllers to the Windows Server 2012 operating system. You can perform this upgrade by upgrading all of the existing domain controllers to Windows Server 2012, or by introducing new domain controllers that are running Windows Server 2012, and then phasing out the existing domain controllers.

An in-place operating system upgrade does not perform automatic schema and domain preparation. To perform an in-place upgrade of a computer that has the AD DS role installed, you must first use the command-line commands **adprep.exe /forestprep** and **adprep.exe /domainprep** to prepare the forest and domain. 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 run the Windows Server 2012 operating system upgrade.

When you promote a Windows Server 2012 server to be a domain controller in an existing domain, and you are signed in as a member of the Schema Admins and Enterprise Admins groups, the AD DS schema will be updated automatically to Windows Server 2012. In this scenario, you do not need to run the Adprep.exe commands before starting the installation.

Deploying Windows Server 2012 Domain Controllers

To upgrade the operating system of a Windows Server 2008 domain controller to Windows Server 2012, perform the following steps:

1. Insert the installation disk for Windows Server 2012, and then run Setup.
The Windows Setup Wizard will eventually open.
2. After the **Language Selection** page, click **Install now**.
3. After the **Operating System Selection** page and the **License Acceptance** page, on the **Which type of installation do you want?** page, click **Upgrade: Install Windows and keep files, settings, and applications**.



Note: With this type of upgrade, there is no need to preserve users' settings and reinstall applications; everything is upgraded in-place. Remember to check for hardware and software compatibility before performing an upgrade.

To introduce a clean install of Windows Server 2012 as a domain controller, perform the following steps:

1. Deploy and configure a new installation of Windows Server 2012 and join it to the domain.
2. Promote the new server to be a domain controller in the domain by using Server Manager 2012 or one of the other methods described previously.
3. Update client DNS settings that are referring to the old domain controller(s) to use the new domain controller.



Note: You can upgrade directly from Windows Server 2008 and Windows Server 2008 R2 to Windows Server 2012.

Installing a Domain Controller by Using Install from Media

If you have an intervening network that is slow, unreliable, or costly, you might find it necessary to add another domain controller at a remote location or branch office. In this scenario, it is often better to deploy AD DS to a server by using the Install from Media (IFM) method.

For example, if you connect to a server which is in a remote office and use Server Manager to install AD DS, you need to copy the entire AD DS database and the SYSVOL folder to the new domain controller over a potentially unreliable WAN connection. As an alternative, and to significantly reduce the amount of traffic moving over the WAN link, you can make a backup of AD DS (perhaps to a USB drive) and take this backup to the remote location. When you are at the remote location and run Server Manager to install AD DS, you can select the option to Install From Media. Most of the copying is then done locally, and the WAN link is used only for security traffic and to ensure that the new domain controller receives any changes that were made to the central AD DS after you created the IFM backup.

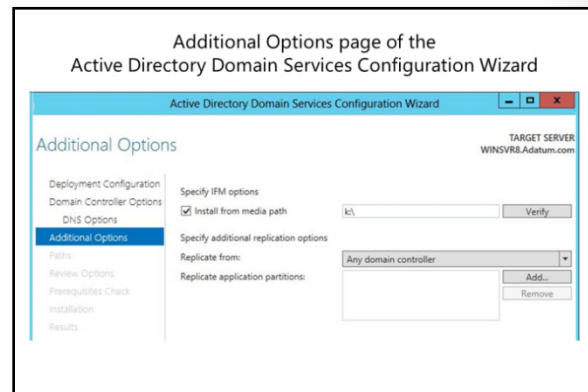
To install a domain controller by using IFM, browse to a domain controller that is not an RODC. Use the Ntdsutil tool to create a snapshot of the AD DS database, and then copy the snapshot to the server that will be promoted to a domain controller. Use Server Manager to promote the server to a domain controller by selecting the Install From Media option, and then providing the local path to the IFM directory that you created previously.

The procedure is as follows:

1. On the full domain controller, at an administrative command prompt, type the following commands (where C:\IFM is the destination directory that will contain the snapshot of the AD DS database):

```
Ntdsutil
Activate instance ntds
Ifm
create SYSVOL full C:\IFM
```

2. On the server that you are promoting to a domain controller, perform the following steps:
 - a. Use Server Manager to add the AD DS role.
 - b. Wait while the AD DS files install.



- c. In Server Manager, click the Notification icon and under Post-Deployment Configuration, click **Promote this server to a domain controller**.

The Active Directory Domain Services Configuration Wizard runs.

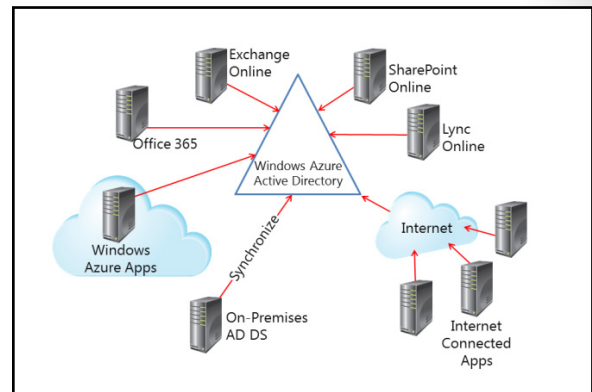
- d. On the appropriate page of the wizard, select the option to install from IFM, and then provide the local path to the snapshot directory.
3. AD DS then installs from the snapshot.
 4. When the domain controller restarts, it contacts other domain controllers in the domain and updates AD DS with any changes that were made since the snapshot was created.



Additional Reading: For more information about the steps necessary to install AD DS, see Install Active Directory Domain Services (Level 100) at <http://go.microsoft.com/fwlink/?LinkID=266739>.

What Is Windows Azure Active Directory?

Windows Azure Active Directory (Windows Azure AD) is a service that provides identity management and access control for your cloud-based applications. Windows Azure AD is used when you subscribe to Microsoft Office® 365, Exchange Online, Microsoft SharePoint® Online, or Microsoft Lync® Online. Additionally, Windows Azure AD can be used with Windows Azure Apps or Internet connected apps that require authentication. You can synchronize your on-premises AD DS with Windows Azure AD to allow your users to use the same identity across both internal resources and cloud-based resources.



Windows Azure AD does not include all the services available with an on-premises Windows Server 2012 AD solution. Windows Server 2012 AD supports five different services:

- Active Directory Domain Service (AD DS)
- Active Directory Lightweight Directory Service (AD LDS)
- Active Directory Federation Service (AD FS)
- Active Directory Certificate Service (AD CS)
- Active Directory Rights Management Service (AD RMS)

Windows Azure AD only includes:

- Windows Azure AD which supports identity management in the cloud
- Windows Azure Access Control Service which supports federation with external identity management services, including your on-premises AD DS

Windows Azure AD does not support Active Directory Integrated Applications, for applications to integrate with Windows Azure AD they need to be written for Windows Azure AD.

Deploying Domain Controllers in Windows Azure

Windows Azure provides Infrastructure as a Service (IaaS), which essentially is virtualization in the cloud. All the considerations for virtualizing applications and servers in on-premises infrastructure apply to deploying the same applications and servers to Windows Azure.



Note: Windows Server 2012 AD which has been deployed in Windows Azure is not the same as Windows Azure AD.

- Windows Server 2012 is cloud-ready and virtualization safe
- Considerations for deploying in Windows Azure include:
 - Rollback
 - Resource limitations
- Virtualization considerations for deploying AD DS
 - Time synchronization
 - Single point of failure

Windows Server 2012 AD which has been deployed in Windows Azure is your own roles and services (AD DS, AD LDS, AD FS, AD CS, and AD RMS) which you have deployed into Windows Azure.

When you deploy AD DS in Windows Azure you are responsible for maintaining everything except the hardware.

Windows Azure AD is a service that Microsoft has configured in the cloud. It does not have all of the functions that an on-premises AD DS has; it is primarily concerned with identity management and access control.

With Windows Azure AD you are only responsible for managing your data.

Windows Server 2012 has been designed to make it easy for you to integrate it into cloud-based systems. One of the most important decisions an administrator must make is whether the organization should use public cloud IaaS, private cloud virtualization technology, or continue to use physical servers.

When implementing AD DS in Windows Azure consider the following:

- **Rollback.** While Windows Azure does not provide rollback services to customers, Windows Azure servers may be rolled back as a regular part of maintenance. Domain controller replication depends on Update Sequence Numbers (USN), when an AD DS system is rolled back duplicate USNs could be created. To prevent this, Windows 2012 AD DS introduced a new identifier named *VM-Generation ID*. VM-Generation ID can detect a rollback, and prevents the virtualized domain controller from replicating changes outbound until the virtualized AD DS has converged with the other domain controllers in the domain.
- **Virtual machine limitations.** Windows Azure virtual machines are limited to 14 GB of RAM and one network adapter. Also, the snapshot feature is not supported.

Deploying Windows Server 2012 AD on Windows Azure virtual machines is subject to the same guidelines as running AD DS on-premises in a virtual machine. These guidelines include the following:

- **Time Synchronization.** A Windows-based AD DS domain infrastructure relies loosely on all communicating machines having the correct time. When domain controller clocks and domain member clocks have a time difference of more than five minutes, clients cannot sign in or access resources on the network. Therefore, Windows has the Windows Time Service (w32time). This service ensures that the time is synchronized across the domain in the following manner:
 - Domain members obtain the time from their domain controller.
 - Domain controllers use the PDC emulator from their own domain or from their parent domain.

- The PDC emulator of the root domain should be configured with an external time source, such as an internet time provider using the network time protocol (NTP).

In virtualized environments time synchronization is not as easy as on physical computers. The virtualization engine regulates the use of the virtualization host's central processing units (CPUs) and distributes it among the virtual machines as needed. The clock of the operating system relies on stable CPU cycles, which do not exist in virtual environments. Virtualization engines perform time synchronization with the guest computers by default. When virtualization hosts do not participate in time synchronization, it is likely the domain time and the virtualization host time will become out of synchronization. While the physical computers participate in the time synchronization, virtual machines are reset to the time on the virtualization host. You must configure the virtualization host to participate in time synchronization or disable the synchronization to the virtual domain controllers for the time synchronization to work properly.

- **Single Point of Failure.** Your AD DS domain controllers are the most important pieces of your infrastructure. If they fail, users are unable to sign in, access resources or applications, and certain services may not run as well as other applications or services.

When virtualizing domain controllers, it is very important that there is not a single point of failure of your domain controller AD DS infrastructure. Setting up all domain controllers as virtual servers on the same virtualization cluster is considered a single point of failure. Similarly, using replicated SANs between datacenters can become a single point of failure if corrupt data is replicated.

If domain controllers are distributed as mentioned in the Domain Membership section, you will ensure that there is not a single point of failure. This will save you from having to recover your entire forest if anything happens to your virtualization infrastructure.

Question: What is the reason to specify the DSRM password?

Lab: Installing Domain Controllers

Scenario

You have been asked by your manager to install a new domain controller in the data center to improve sign-in performance. You have been asked also to create a new domain controller for a branch office by using IFM.

Objectives

After performing this lab, you will be able to:

- Install a domain controller.
- Install a domain controller by using IFM.

Lab Setup

Estimated time: 45 minutes

Virtual machines	20410C-LON-DC1 20410C-LON-SVR1 20410C-LON-RTR 20410C-LON-SVR2
User name	Adatum\Administrator
Password	Pa\$\$w0rd

For this lab, you will use the available virtual machine environment. Before beginning 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 **20410C-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in using the following credentials:
 - User name: **Administrator**
 - Password: **Pa\$\$w0rd**
 - Domain: **Adatum**
5. Repeat steps 2 through 4 for 20410C-LON-SVR1, 20410C-LON-RTR, and 20410C-LON-SVR2.

Exercise 1: Installing a Domain Controller

Scenario

Users have been experiencing slow sign-ins in London during peak use times. The server team has determined that the domain controllers are overwhelmed when many users are authenticating simultaneously. To improve sign-in performance, you are adding a new domain controller in the London data center.

The main tasks for this exercise are as follows:

1. Add an Active Directory Domain Services (AD DS) role to a member server.
2. Configure a server as a domain controller.
3. Configure a server as a global catalog server.

► **Task 1: Add an Active Directory Domain Services (AD DS) role to a member server**

1. On LON-DC1, in Server Manager, add **LON-SVR1** to the server list.
2. Add the **Active Directory Domain Services** server role to **LON-SVR1**. Add all required features as prompted.

Installation will take several minutes.

3. When the installation completes, click **Close** to close the Add Roles and Features Wizard.

► **Task 2: Configure a server as a domain controller**

- On LON-DC1, use Server Manager to promote LON-SVR1 to a domain controller, and choose the following options:
 - Add a domain controller to the existing Adatum.com domain
 - Use the credentials **Adatum\Administrator** with the password **Pa\$\$w0rd**
 - For Domain Controller Options, install the **Domain Name System**, but remove the selection to install the global catalog
 - The DSRM password is **Pa\$\$w0rd**
 - For all other options, use the default options

► **Task 3: Configure a server as a global catalog server**

1. Sign in to LON-SVR1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Use Active Directory Sites and Services to make LON-SVR1 a global catalog server.

Results: After completing this exercise, you will have explored Server Manager and promoted a member server to be a domain controller.

Exercise 2: Installing a Domain Controller by Using IFM

Scenario

You have been assigned by management to manage one of the new branch offices that are being configured. A faster network connection is scheduled to be installed in a few weeks. Until that time, network connectivity is very slow.

It has been determined that the branch office requires a domain controller to support local sign-ins. To avoid problems with the slow network connection, you are using IFM to install the domain controller in the branch office.

The main tasks for this exercise are as follows:

1. Use the Ntdsutil tool to generate IFM.
2. Add the AD DS role to the member server.
3. Use IFM to configure a member server as a new domain controller.

► Task 1: Use the Ntdsutil tool to generate IFM

1. On LON-DC1, open an administrative command-line interface, and use **ntdsutil** to create an IFM backup of both the AD DS database and the SYSVOL folder. The commands to create the backup are as follows:

```
Ntdsutil
Activate instance ntds
Ifm
Create sysvol full c:\ifm
```

2. Wait for the IFM command to complete and then close the command prompt.

► Task 2: Add the AD DS role to the member server

1. Switch to **LON-SVR2**, and sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Open a command prompt and map the drive letter **K:** to **\\LON-DC1\C\$\IFM**.
3. Use Server Manager to install the AD DS server role on LON-SVR2.

► Task 3: Use IFM to configure a member server as a new domain controller

1. On LON-SVR2, at the command prompt copy the IFM backup from **K:** to **C:\ifm**.
2. On LON-SVR2, use Server Manager with the following options to perform the post-deployment configuration of AD DS:
 - Add a domain controller to the existing Adatum.com domain
 - Use **Adatum\Administrator** with the password **Pa\$\$w0rd** for credentials
 - Use **Pa\$\$w0rd** for the DSRM password
 - Use the IFM media to configure and install AD DS. Use the location **C:\IFM** for the IFM media
 - Accept all other defaults
3. Restart LON-SVR2 to complete the AD DS installation.

Results: After completing this exercise, you will have installed an additional domain controller for the branch office by using IFM.

► Prepare for the next module

When you have completed the lab, revert the virtual machines back 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 **20410C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for 20410C-LON-SVR1, 20410C-LON-RTR, and 20410C-LON-SVR2.

Lab Review Questions

Question: Why did you use Server Manager and not Dcpromo.exe when you promoted a server to be a domain controller?

Question: What are the three operations masters found in each domain?

Question: What are the two operations masters that are present in a forest?

Question: What is the benefit of performing an IFM install of a domain controller?

Module Review and Takeaways

Review Questions

Question: What are the two main purposes of OUs?

Question: Why would you need to deploy an additional tree in the AD DS forest?

Question: Which deployment method would you use if you had to install an additional domain controller in a remote location that had a limited WAN connection?

Question: If you needed to promote a Server Core installation of Windows Server 2012 to be a domain controller, which tool or tools could you use?

Module 3

Managing Active Directory Domain Services Objects

Contents:

Module Overview	3-1
Lesson 1: Managing User Accounts	3-2
Lesson 2: Managing Groups	3-11
Lesson 3: Managing Computer Accounts	3-20
Lesson 4: Delegating Administration	3-27
Lab: Managing Active Directory Domain Services Objects	3-31
Module Review and Takeaways	3-38

Module Overview

AD DS can help you manage your network more effectively in many ways. For instance, it allows you to manage user and computer accounts as part of groups instead of managing one account at a time. It also provides ways to delegate administrative tasks to various people to help you distribute workloads efficiently.

Managing computer identities is becoming more and more complex as more employees bring their own devices into the workplace. As bring your own device (BYOD) programs expand you will be managing computer accounts that run on many types of personal devices which in turn are running various operating systems. AD DS has many features that can make that easier.

This module describes how to manage user accounts and computer accounts, including how to manage BYOD programs. It covers how to manage an enterprise network by managing groups, instead of managing individual identities, and how to delegate administrative tasks to designated users or groups to ensure that enterprise administration is efficient and effective.

Objectives

After completing this module, you will be able to:

- Manage user accounts with graphical tools.
- Manage group accounts with graphical tools.
- Manage computer accounts.
- Delegate permissions to perform AD DS administration.

Lesson 1

Managing User Accounts

A user object in AD DS is far more than just properties related to the user's security identity, or account. It is the cornerstone of identity and access in AD DS. Therefore, consistent, efficient, and secure processes regarding the administration of user accounts are the cornerstone of enterprise security management.

Although users and computers, and even services, change over time, business roles and rules tend to be more stable. Your business probably has a finance role, which requires certain capabilities in the enterprise. The user or users who perform that role might change over time, but the role will remain relatively the same. For that reason, it is not sensible to manage an enterprise network by assigning rights and permissions to individual users, computers, or service identities. Instead, you should associate management tasks with groups. Consequently, it is important that you know how to use groups to identify administrative and user roles, to filter Group Policy Objects (GPOs), to assign unique password policies, and to assign rights and permissions.

Lesson Objectives

After completing this lesson, you will be able to:

- View AD DS objects by using various AD DS administration tools.
- Explain how to create user accounts that you can use in an enterprise network.
- Describe how to configure important user account attributes.
- Describe how to create user profiles.
- Explain how to manage user accounts.

AD DS Administration Tools

Before you can begin creating and managing user, group, and computer accounts, it is important that you understand which tools you can use to perform these various management tasks.

Active Directory Administration Snap-Ins

Most AD DS administration is performed with the following snap-ins:

- Active Directory Users and Computers. This snap-in manages most common day-to-day resources, including users, groups, computers, and organizational units (OUs). This is likely to be the most heavily used snap-in for an Active Directory administrator.
- Active Directory Sites and Services. This snap-in manages replication, network topology, and related services.
- Active Directory Domains and Trusts. This snap-in configures and maintains trust relationships and the forest functional level.
- Active Directory Schema. This snap-in examines and modifies the definition of Active Directory attributes and object classes. It is the blueprint for AD DS. It is rarely viewed, and even more rarely changed. Therefore, the Active Directory Schema snap-in is not registered, by default.

To manage AD DS objects, you can use the following graphical tools:

- Active Directory Administration snap-ins
- Active Directory Administrative Center



You can also use the following command-line tools:

- Active Directory module in Windows PowerShell
- Directory Service commands





Note: To register the Active Directory Schema snap-in, run the following command in an elevated command prompt:

```
regsvr32 schmmgmt.dll
```



Note: To administer AD DS from a computer that is not a domain controller, you must install Remote Server Administration Tools (RSAT). RSAT can be installed from the Features node of Server Manager on Windows Server® 2012.

You also can install RSAT on Windows clients, including Windows Vista® Service Pack 1 (or newer), Windows 7, and Windows 8. After you download the RSAT installation files from the Microsoft website, run the Setup Wizard, which steps you through the installation. After installing RSAT, you must turn on the tool or tools that you want to use. To do this, in Control Panel, on the Programs And Features category page, use Turn Windows Features On or Off.



Reference Links: To download the RSAT installation files, see the Microsoft Download Center at <http://go.microsoft.com/fwlink/?LinkID=266735>.

Active Directory Administrative Center

Windows Server 2012 provides another option for managing AD DS objects. The Active Directory Administrative Center provides a graphical user interface (GUI) built upon Windows PowerShell®. You can use this enhanced interface to perform AD DS object management by using task-oriented navigation. Tasks that you can perform by using the Active Directory Administrative Center include:

- Create and manage user, computer, and group accounts
- Create and manage OUs
- Connect to, and manage multiple domains within a single instance of the Active Directory Administrative Center
- Search and filter Active Directory data by building queries
- Manage Dynamic Access Control settings, such as Central Access Policies and Central Access Rules

Windows PowerShell

You can use the Active Directory module for Windows PowerShell (Active Directory module) to create and manage objects in AD DS. Windows PowerShell is not only a scripting language; it also enables you to run commands that perform administrative tasks, such as creating new user accounts, configuring services, deleting mailboxes, and similar functions.

Windows PowerShell is installed by default on Windows Server 2012, but the Active Directory module is only present when you do one of the following:

- Install the AD DS or Active Directory Lightweight Directory Services (AD LDS) server roles.
- Promote a computer to a domain controller
- Install RSAT.

Directory Service Command-Line Tools

You also can use the Directory Service command-line tools, in addition to Windows PowerShell. These tools enable you to create, modify, manage, and delete AD DS objects, such as users, groups, and computers. You can use the following commands:

- **dsadd.** Use to create new objects.
- **dsget.** Use to display objects and their properties.
- **dsmod.** Use to edit objects and their properties.
- **dsmove.** Use to move objects.
- **dsquery.** Use to query AD DS for objects that match criteria that you supply.
- **dsrm.** Use to delete objects.



Note: It is possible to pipe the results of the **dsquery** command to other Directory Service commands. For example, typing the following at a command prompt returns the office telephone number of all users that have a name starting with John:

```
dsquery user -name John* | dsget user -tel
```

Creating User Accounts

In AD DS, all users that require access to network resources must be configured with a user account. With this user account, users can authenticate to the AD DS domain and receive access to network resources.

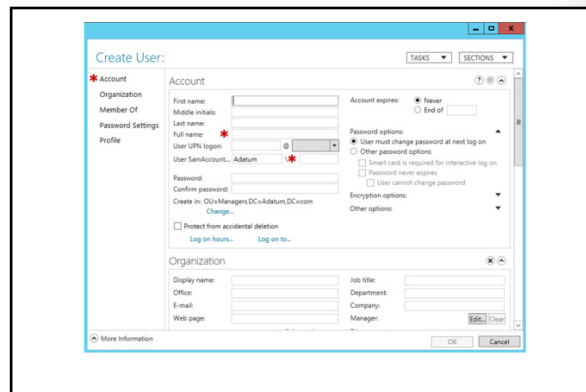
In Windows Server 2012, a *user account* is an object that contains all of the information that defines a user. A user account includes the user name, user password, and group memberships. A user account also contains many other settings that you can configure based upon your organizational requirements.

With a user account, you can:

- Allow or deny users permission to sign in to a computer based on their user account identity.
- Grant users access to processes and services for a specific security context.
- Manage users' access to resources such as AD DS objects and their properties, shared folders, files, directories, and printer queues.

A user account enables a user to sign in to computers and domains with an identity that the domain can authenticate. When creating a user account, you must provide a user logon name, which must be unique in the domain and forest in which the user account is created.

To maximize security, you should avoid multiple users sharing a single account, and instead ensure that each user who signs in to the network has a unique user account and password.



Note: Although AD DS accounts are the focus of this course, you also can store user accounts in the local security accounts manager (SAM) database of each computer, enabling local sign-in and access to local resources. Local user accounts are, for the most part, beyond the scope of this course.

Creating User Accounts

A user account includes the user name and password, which serve as the sign-in credentials for a user. A user object also includes several other attributes that describe and manage the user.

You can use Active Directory Users and Computers, Active Directory Administrative Center, Windows PowerShell, or the **dsadd.exe** command-line tool to create a user object. When you are creating user accounts, consider the following elements:

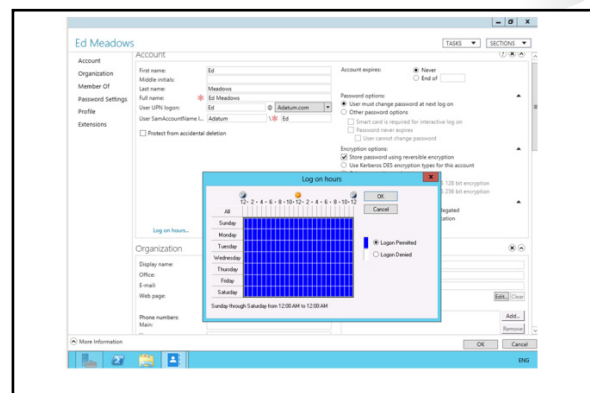
- The Full Name. The Full Name is used to create several attributes of a user object, most notably, the common name and display name attributes. The common name of a user is the name displayed in the details pane of the snap-in, and it must be unique within the container or OU. If you are creating a user object for a person with the same name as an existing user in the same OU or container, you need to give the new user object a unique Full Name.
- The User Principal Name (UPN) Logon. User UPN Logons follow the format *user logon name@(UPN)*.
 - User names in AD DS can contain special characters, including periods, hyphens, and apostrophes. These special characters let you generate accurate user names, such as O'Hare and Smith-Bates. However, certain programs and apps might have other restrictions, so we recommend that you use only standard letters and numerals until you fully test the applications in your enterprise environment for compatibility with special characters.
 - You can manage the list of available UPN suffixes by using the Active Directory Domains and Trusts snap-in. Right-click the root of the snap-in, click Properties, and then use the UPN Suffixes tab to add or remove suffixes. The Domain Name System (DNS) name of your AD DS domain is always available as a suffix, and you cannot remove it.

Note: It is important that you implement a user account naming strategy, especially in large networks where users might share the same full name. A combination of last name and first name, and where necessary, additional characters, should yield a unique user account name. Specifically, it is only the UPN name that must be unique within your AD DS forest. The Full name needs to be unique only within the OU where it resides, while the User SamAccountName name must be unique within that domain.

Configuring User Account Attributes

When you create a user account in AD DS, you also configure all the associated account properties, or attributes.

Note: The attributes that are associated with a user account are defined as part of the AD DS schema, which members of the Schema Admins security group can modify. Generally, the schema does not change frequently. However, when an enterprise-level program (such as Microsoft® Exchange Server) is introduced, many schema changes are required. These changes enable objects, including user objects, to have additional attributes.



When you create a new user object, you are not required to define many attributes beyond those required to allow the user to log on by using the account. Since you can associate a user object with many attributes, it is important that you understand what these attributes are, and how you can use them in your organization.

Attribute Categories

The attributes of a user object fall into several broad categories. These categories appear in the navigation pane of the User Properties dialog box in the Active Directory Administrative Center, and include the following:

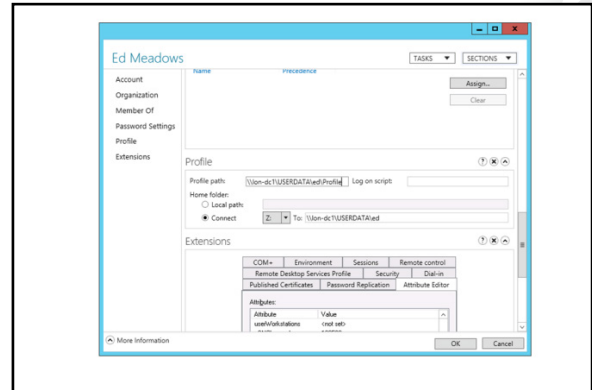
- **Account.** In addition to the user's name properties (First name, Middle initial, Last name, Full name) and the user's various logon names (User UPN logon, User SamAccountName logon), you can configure the following additional properties:
 - **Log on hours.** This property defines when the account can be used to access domain computers. You can use the weekly calendar style view to define Logon permitted hours and Logon denied hours.
 - **Log on to.** Use this property to define which computers a user can use to log on to the domain. Specify the computer's name and add it to a list of allowed computers.
 - **Account expires.** This value is useful when you want to create temporary-use user accounts. For example, you might want to create user accounts for interns that are used for just one year. You can use this value to set an account expiration date in advance. The account cannot be used after the expiration date until manually reconfigured by an administrator.
 - **User must change password at next log on.** This property enables you to force a user to reset their own password the next time they log on. This is typically something you might enable after you have reset a user's password.
 - **Smart card is required for interactive log on.** This value resets the user's password to a complex, random sequence of characters, and sets a property that requires that the user use a smart card to authenticate during logon.
 - **Password never expires.** This is a property that you normally use with service accounts; that is, those accounts that are not used by regular users but by services. By setting this value, you must remember to update the password manually on a periodic basis; however, you are not forced to do so at a predetermined interval. Consequently, the account can never be locked out due to password expiration—a feature that is particularly important for service accounts.
 - **User cannot change password.** Again, this option is generally used for service accounts.
 - **Store password using reversible encryption.** This policy provides support for programs that use protocols that require knowledge of the user's password for authentication purposes. Storing passwords using reversible encryption is essentially the same as storing plaintext versions of the passwords. For this reason, this policy should never be enabled unless program requirements outweigh the need to protect password information. This policy is required when using Challenge Handshake Authentication Protocol (CHAP) authentication through remote access or Internet Authentication Service (IAS). It is also required when using Digest Authentication in Internet Information Services (IIS).
 - **Account is trusted for delegation.** You can use this property to allow a service account to impersonate a standard user to access network resources on behalf of a user.
- **Organization.** This includes properties such as the user's Display name, Office, Email Address, various contact telephone numbers, managerial structure, department and company names, addresses and so forth.
- **Member of.** This section enables you to define the group memberships for the user.

- **Profile.** This section enables you to configure a location for the user's personal data, and to define a location in which to save the user's desktop profile when he or she logs out.
- **Extensions.** This section exposes many additional user properties, most of which do not normally require manual configuration.

Creating User Profiles

When users sign out, their desktop and app settings are saved to a subfolder that is created in the C:\Users folder on the local hard disk that matches their user name. This folder contains their user profile. Within this folder, subfolders contain documents and settings that represent the user's profile, including Documents, Videos, Pictures, and Downloads.

If a user is likely to sign in interactively at more than one client workstation, it is preferable for these settings and documents to be available on those other client workstations. As an administrator, you have a number of methods that you can use to ensure that users can access their profiles from multiple workstations.



Configuring User Account Properties to Manage Profiles

Using the user account settings in the Active Directory Administrative Center, you can configure the following properties relating to a user's desktop profile:

- **Profile path.** This path is either a local, or more usually, a Universal Naming Convention (UNC) path. The user's desktop settings are stored in the profile. If a user profile has a UNC path, then whichever domain computer services a user's sign-in, their desktop settings will be available. This is known as a *roaming profile*.



Note: As a best practice, use a subfolder of the user's home folder for the user's profile path.

- **Logon script.** This script is the name of a batch file that contains commands that execute when the user logs on. Typically, you use these commands to create drive mappings. Rather than use a logon script batch file, administrators typically implement logon scripts by using GPOs or Group Policy preferences. If you use a login script, this value should be in the form of a filename (with extension) only. Scripts should be stored in the C:\Windows\SYSVOL\domain\scripts folder on all domain controllers.
- **Home folder.** This value enables you to create a personal storage area in which users can save their personal documents. You can specify either a local path, or more usually, a UNC path to the user's folder. You also must specify a drive letter that is used to map a network drive to the specified UNC path. You can then configure a user's personal documents to this redirected home folder.



Note: When creating user accounts to use as templates, and using a common location for the profile path and home folder, it is recommended that you use the **%username%** variable in the path so that AD DS can automatically create these folders when the account is used as a template. For example, you could use the following paths where the fileserver is named LON-FS and shares have been created for the profiles and home folders, profile\$ and home\$ respectively:

Profile Path: \\LON-FS\profile\$\b>%username%

Home folder Connect H: to \\LON-FS\home\$\b>%username%

Using Group Policy to Manage Profiles

As an alternative to using the individual user account settings, you can also use GPOs to manage these settings. You can configure Folder Redirection settings by using the Group Policy Management Editor to open a GPO for editing, and then navigating to the User Configuration\Policies\Windows Settings node.

These settings contain the sub-nodes in the following table.

Sub-nodes in the Windows Settings node		
<ul style="list-style-type: none">• AppData (Roaming)• Desktop• Start Menu• Document	<ul style="list-style-type: none">• Pictures• Music• Videos• Favorites• Contacts	<ul style="list-style-type: none">• Downloads• Links• Searches• Saved Games

You can use these sub-nodes to configure all aspects of a user’s desktop profile and app settings. For a given sub-node, such as Documents, you can choose between Basic and Advanced redirection. In Basic redirection, all users affected by the GPO have their Documents folder redirected to an individual named subfolder off a common root folder defined by a UNC name, for example, \\LON-SVR1\Users\. In Advanced redirection you can use security group membership to specify where a user’s settings and documents will be stored.

Demonstration: Managing User Accounts

This demonstration shows you how to:

- Use the Active Directory Administrative Center to:
 - Delete a user account.
 - Create a new user account.
 - Move the user account.
- Use Windows PowerShell to:
 - Find inactive user accounts.
 - Find disabled user accounts.
 - Delete disabled user accounts.

Demonstration Steps

Delete a user account

1. Sign in to LON-DC1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On LON-DC1, open the **Active Directory Administrative Center**.
3. Locate Ed Meadows in the Managers OU, and delete the account.

Create a new user account

- Create a new user account named **Ed Meadows**. Ensure that the account is created with a strong password.

Move the user account

- Move the Ed Meadows account to the IT OU.

Find users that have not signed in during the last 30 days

1. Open **Windows PowerShell**.
2. Run the following commands:

```
$logonDate = (get-date).AddDays(-30)
Get-ADUser -Filter{lastLogon -le $logonDate}
```

Find and delete all disabled user accounts

1. To find all disabled user accounts, run the following command:

```
Get-ADUser -Filter{enabled -ne $True}
```

2. To delete the disabled user accounts, run the following command:

```
Get-ADUser -SearchBase "OU=Sales,DC=Adatum,DC=com" -Filter{enabled -ne $true} |
Remove-ADObject -Confirm:$False
```

3. To verify the disabled accounts have been deleted, run the following command:

```
Get-ADUser -Filter{enabled -ne $True}
```

Demonstration: Using Templates to Manage User Accounts

This demonstration shows you how to:

- Create a user template account.
- Use Windows PowerShell to create a user from the user template.
- Verify the properties of the new user account.

Demonstration Steps

Create a template account

1. On LON-DC1, open the **Active Directory Administrative Center**.
2. In the Sales OU, create a new user account named **_LondonSales Template** using the following settings:
 - First name: **_LondonSales**
 - Last name: **Template**
 - User UPN logon: **_LondonSales**
 - Check: **Protect from accidental deletion**
 - Department: **Sales**
 - Company: **A. Datum**
 - City: **London**
 - Description: **London Sales users**
 - Member of: **Sales**

Create a user from the _LondonSales template

- In the Windows PowerShell window, create a user from the _LondonSales template by using the following commands.

```
$LondonSales = Get-ADUser -Identity "_LondonSales" -Properties  
Department,Company,City  
  
New-ADUser -Name "Dan Park" -SamAccountName "Dan" -Path "OU=Sales,DC=Adatum,DC=com"  
-AccountPassword (ConvertTo-SecureString -AsPlainText "Pa$$w0rd" -Force) -GivenName  
"Dan" -Surname "Park" -DisplayName "Dan Park" -Enabled $True -UserPrincipalName  
"Dan@Adatum.com" -ChangePasswordAtLogon $true -Instance $LondonSales
```

Verify the user properties

1. Run the following command:

```
Get-ADUser -Identity "Dan" -Properties *
```

2. Examine the output to verify the properties were copied from the template.

Lesson 2

Managing Groups

While it might be practical to assign permissions and abilities to individual user accounts in small networks, it becomes impractical and inefficient in large enterprise networks. For example, if many users need the same level of access to a folder, it is more efficient to create a group that contains the required user accounts, and then assign the group the required permissions. This has the added benefit of enabling you to change a user's file permissions by adding or removing them from groups rather than editing the file permissions directly.

Before implementing groups in your organization, you must understand about the scope of various Windows Server group types, and how best to use these to manage access to resources or to assign management rights and abilities.

Lesson Objectives

After completing this lesson, you will be able to:

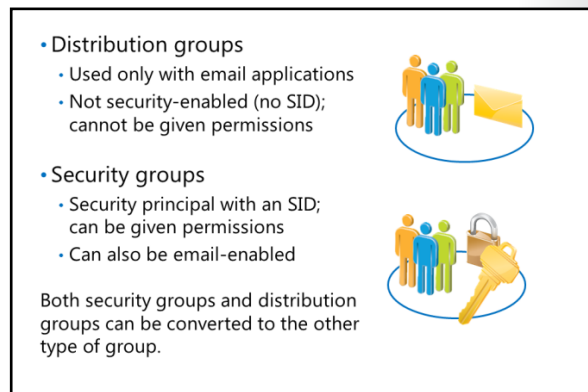
- Describe group types.
- Describe group scopes.
- Explain how to implement group management.
- Describe default groups.
- Describe special identities.
- Manage groups in Windows Server.

Group Types

In a Windows Server 2012 enterprise network, there are two types of groups: security, and distribution. When you create a group, you choose the group type and scope.

Distribution groups, which are not security-enabled, are used primarily by email apps. Distribution groups have security identifier (SIDs) however they have an Active Directory **groupType** attribute of 0x2 (ACCOUNT_GROUP), so they cannot be given permission to resources. Sending an email message to a distribution group sends the message to all group members.

Security groups are security-enabled, and are used to assign permissions to various resources. Security groups have SIDs, with an Active Directory **groupType** attribute of 0x80000002 (ACCOUNT_GROUP | SECURITY_ENABLED). You can therefore use these groups in permission entries in access control lists (ACLs) to control security for resource access. Because they also include the account group type, you also can use security groups as a means of distribution for email apps. If you want to use a group to manage security, it must be a security group.



Note: The default group type is security.

Because you can use security groups for both resource access and email distribution, many organizations use only security groups. However, we recommend that if a group is used only for email distribution, you should create the group as a distribution group. Otherwise, the group is assigned a SID, and the SID is added to the user's security access token, which can make the token unnecessarily large.

A security group can be converted to a distribution group at any time; when you do this the **groupType** attribute changes from 0x80000002 (ACCOUNT_GROUP | SECURITY_ENABLED) to 0x2 (ACCOUNT_GROUP). A security group which has been converted to a distribution group therefore loses all permissions assigned to it, even though the ACLs still contain the SID. When a distribution group is converted to a security group the reverse occurs, the **groupType** attribute changes from 0x2 (ACCOUNT_GROUP) to 0x80000002 (ACCOUNT_GROUP | SECURITY_ENABLED). If the distribution group was a previously converted security group you may inadvertently grant users access to rights and permissions that had been assigned to the group when it was previously a security group.



Note: Consider that when you add a user to a security group, the user's access token—which authenticates user processes—updates only when the user signs in. Therefore, if the user is currently signed in, they must sign out and sign back in to update their access token with any changed group memberships.



Note: The benefit of using distribution groups becomes more evident in large-scale Exchange Server deployments, especially where there is a need to nest these distribution groups across the enterprise.

Group Scopes

Windows Server 2012 supports group scoping. The scope of a group determines both the range of a group's abilities or permissions, and the group membership.

There are four group scopes:

- **Local.** This group type is for standalone servers or workstations, on domain member servers that are not domain controllers, or on domain member workstations. Local groups are truly local, which means that they are available only on the computer where they exist. The important characteristics of a local group are:
 - You can assign abilities and permissions only on local resources, meaning on the local computer.
 - Members can be from anywhere in the AD DS forest, and can include:
 - Any security principals from the domain: users, computers, global groups, or domain local groups.
 - Users, computers, and global groups from any domain in the forest.
 - Users, computers, and global groups from any trusted domain.
 - Universal groups defined in any domain in the forest.

Group scope	Members from same domain	Members from domain in same forest	Members from trusted external domain	Can be assigned permissions to resources
Local	U, C, GG, DLG, UG and local users	U, C, GG, UG	U, C, GG	On the local computer only
Domain Local	U, C, GG, DLG, UG	U, C, GG, UG	U, C, GG	Anywhere in the domain
Universal	U, C, GG, UG	U, C, GG, UG	N/A	Anywhere in the forest
Global	U, C, GG	N/A	N/A	Anywhere in the domain or a trusted domain

U	User	DLG	Domain Local Group
C	Computer	UG	Universal Group
GG	Global Group		

- Domain local. This group type is used primarily to manage access to resources or to assign management responsibilities (rights). Domain local groups exist on domain controllers in an AD DS forest, and consequently, the group's scope is localized to the domain in which they reside. The important characteristics of domain local groups are:
 - You can assign abilities and permissions only on domain local resources, meaning on all computers in the local domain.
 - Members can be from anywhere in the AD DS forest, and can include:
 - Any security principals from the domain: users, computers, global groups, or domain local groups.
 - Users, computers, and global groups from any domain in the forest.
 - Users, computers, and global groups from any trusted domain.
 - Universal groups defined in any domain in the forest.
- Global. This group type is used primarily to consolidate users that have similar characteristics. For example, global groups often are used to consolidate users that are part of a department or geographic location. The important characteristics of global groups are:
 - You can assign abilities and permissions anywhere in the forest.
 - Members can be only from the local domain, and can include:
 - Users, computers, and global groups from the local domain.
- Universal. This group type is most useful in multidomain networks because it combines the characteristics of both domain local groups and global groups. Specifically, the important characteristics of universal groups are:
 - You can assign abilities and permissions anywhere in the forest, as with global groups.
 - Members can be from anywhere in the AD DS forest, and can include:
 - Users, computers, and global groups from any domain in the forest.
 - Universal groups defined in any domain in the forest.
 - Properties of universal groups are propagated to the global catalog, and made available across the enterprise network on all domain controllers that host the global catalog role. This makes universal groups' membership lists more accessible, which is useful in multidomain scenarios. For example, if a universal group is used for email distribution purposes, the process for determining the membership list typically is quicker in distributed multidomain networks.

The following table summarizes and compares the basic properties of the four group scopes:

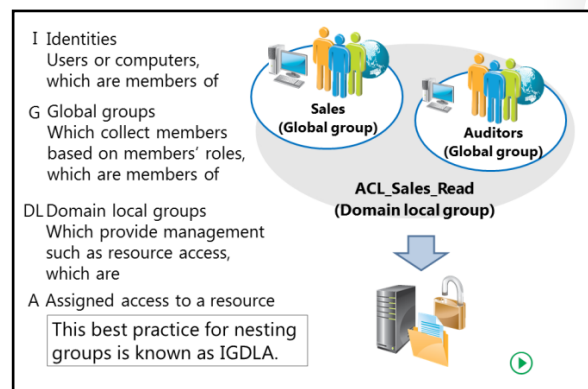
Group Scope	Can include members from these groups	Can be assigned permissions to these groups	Can be converted to these groups
Local	<ul style="list-style-type: none"> Domain Users, Domain Computers, global groups and universal groups from any domain in the forest Domain local groups from the same domain Local Users from the computer 	Local computer resources only	N/A
Domain local	<ul style="list-style-type: none"> Domain Users, Domain Computers, global groups and universal groups from any domain in the forest Domain local groups from the same domain 	Local domain resources only	Universal groups (as long as no other domain local groups exist as members)
Global	<ul style="list-style-type: none"> Domain Users, Domain Computers and global groups from the same domain 	Any domain resource in the forest	Universal groups (as long as it is not a member of any other global groups)
Universal	<ul style="list-style-type: none"> Domain Users, Domain Computers, global groups and universal groups from any domain in the forest 	Any domain resource in the forest	Domain local groups Global groups (as long as no other universal groups exist as members)

Implementing Group Management

Adding groups to other groups is a process called *nesting*. Nesting creates a hierarchy of groups that support your business roles and management rules.

A best practice for group nesting is known as IGDLA, which is an acronym for the following:

- Identities
- Global groups
- Domain local groups
- Access



These parts of IGDLA are related in the following way:

- Identities (user and computer accounts) are members of global groups, which represent business roles.
- Global groups (which are also known as role groups) are members of domain local groups, which represent management rules—for example, determining who has Read permission to a specific collection of folders.

- Domain local groups (which are also known as rule groups) are granted access to resources. In the case of a shared folder, access is granted by adding the domain local group to the folder's ACL, with a permission that provides the appropriate level of access.

In a multidomain forest, the best practice for group nesting is known as IGUDLA. The additional letter U stands for universal groups, which fit in between global and domain local groups as follows:

- Identities
- Global groups
- Universal groups
- Domain local groups
- Access

In this case, global groups from multiple domains are members of a single universal group. That universal group is a member of domain local groups in multiple domains.

IGDLA Example

This figure on the slide represents a group implementation that reflects the technical view of group management best practices (IGDLA), and the business view of role-based, rule-based management.

Consider the following scenario:

The sales force at Contoso, Ltd. has just completed its fiscal year. Sales files from the previous year are in a folder called Sales. The sales force needs Read access to the Sales folder. Additionally, a team of auditors from Woodgrove Bank, a potential investor, require Read access to the Sales folder to perform the audit. You would perform the following steps to implement the security required by this scenario:

1. Assign users with common job responsibilities or other business characteristics to role groups implemented as global security groups. Do this separately in each domain. Salespeople at Contoso are added to a Sales role group; Auditors at Woodgrove Bank are added to an Auditors role group.
2. Create a group to manage access to the Sales folders with Read permission. You implement this in the domain containing the resource that is being managed. In this case, the Sales folder resides in the Contoso domain. Therefore, you create the resource access management rule group as a domain local group named ACL_Sales Folders_Read.
3. Add the role groups to the resource access management rule group to represent the management rule. These groups can come from any domain in the forest or from a trusted domain, such as Woodgrove Bank. Global groups from trusted external domains, or from any domain in the same forest, can be members of a domain local group.
4. Assign the permission that implements the required level of access. In this case, grant the Allow Read permission to the domain local group.

This strategy results in two single points of management, reducing the management burden. One point of management defines who is in Sales, the other point of management defines who is an Auditor. Because these roles are likely to have access to a variety of resources beyond the Sales folder, you have another single point of management to determine who has Read access to the Sales folder. Furthermore, the Sales folder might not be a single folder on a single server; it could be a collection of folders across multiple servers, each of which assigns the Allow Read permission to the single domain local group.

Default Groups

Windows Server 2012 creates a number of groups automatically. These are called *default local groups*, and they include well-known groups such as Administrators, Backup Operators, and Remote Desktop Users. There are additional groups that are created in a domain, both in the Builtin and Users containers, including Domain Admins, Enterprise Admins, and Schema Admins.

Default Groups that Provide Administrative Privileges

The following list provides a summary of capabilities for the subset of default groups that have significant permissions and user rights related to the management of AD DS. Because of the rights that these groups have they are Protected groups.

- Carefully manage the default groups that provide administrative privileges, because these groups:
 - Typically have broader privileges than are necessary for most delegated environments
 - Often apply protection to their members

Group	Location
Enterprise Admins	Users container of the forest root domain
Schema Admins	Users container of the forest root domain
Administrators	Built-in container of each domain
Domain Admins	Users container of each domain
Server Operators	Built-in container of each domain
Account Operators	Built-in container of each domain
Backup Operators	Built-in container of each domain
Print Operators	Built-in container of each domain
Cert Publishers	Users container of each domain

- Enterprise Admins (in the Users container of the forest root domain). This group is a member of the Administrators group in every domain in the forest, which gives it complete access to the configuration of all domain controllers. It also owns the Configuration partition of the directory and has full control of the domain naming context in all forest domains.
- Schema Admins (Users container of the forest root domain). This group owns and has full control of the Active Directory schema.
- Administrators (Built-in container of each domain). Members of this group have complete control over all domain controllers and data in the domain naming context. They can change the membership of all other administrative groups in the domain, and the Administrators group in the forest root domain can change the membership of Enterprise Admins, Schema Admins, and Domain Admins. The Administrators group in the forest root domain is arguably the most powerful service administration group in the forest.
- Domain Admins (Users container of each domain). This group is added to the Administrators group of its domain. It therefore inherits all of the capabilities of the Administrators group. It is also, by default, added to the local Administrators group of each domain member computer, giving Domain Admins ownership of all domain computers.
- Server Operators (Built-in container of each domain). Members of this group can perform maintenance tasks on domain controllers. They have the right to sign in locally, start and stop services, perform backup and restore operations, format disks, create or delete shares, and shut down domain controllers. By default, this group has no members.
- Account Operators (Built-in container of each domain). Members of this group can create, modify, and delete accounts for users, groups, and computers located in any OU in the domain (except the Domain Controllers OU), and in the Users and Computers containers. Account Operator group members cannot modify accounts that are members of the Administrators or Domain Admins groups, nor can they modify those groups. Account Operator group members also can sign in locally to domain controllers. By default, this group has no members.
- Backup Operators (Built-in container of each domain). Members of this group can perform backup and restore operations on domain controllers, and sign in locally and shut down domain controllers. By default, this group has no members.

- Print Operators (Built-in container of each domain). Members of this group can maintain print queues on domain controllers. They also can sign in locally and shut down domain controllers.
- Cert Publishers (Users container of each domain). Members of this group are permitted to publish certificates to the directory.

Managing Groups that Provide Administrative Privileges

You need to carefully manage the default groups that provide administrative privileges because they typically have broader privileges than are necessary for most delegated environments, and because they often apply protection to their members.

The Account Operators group is a good example of this. If you examine the capabilities of the Account Operators group in the preceding list, you can see that members of this group have very broad rights—they can even sign in locally to a domain controller. In very small networks, such rights would probably be appropriate for one or two individuals who typically would be domain administrators anyway. In large enterprises, the rights and permissions granted to Account Operators usually are far too broad.

Additionally, the Account Operators group is, like the other administrative groups, a protected group.

Protected Groups

Protected groups are defined by the operating system and cannot be unprotected. Members of a protected group become protected by association and no longer inherit permissions (ACLs) from their OU, but rather receive a copy of an ACL from the protected group. This protected group ACL offers quite a lot of protection to the members. For example, if you add Jeff Ford to the Account Operators group, his account becomes protected, and the help desk, which can reset all other user passwords in the Employees OU, cannot reset Jeff Ford's password.

Custom Groups

You should try to avoid adding users to the groups that do not have members by default (Account Operators, Backup Operators, Server Operators, and Print Operators). Instead, create custom groups to which you assign permissions and user rights that achieve your business and administrative requirements.

For example, if Scott Mitchell should be able to perform backup operations on a domain controller, but should not be able to perform restore operations that could lead to database rollback or corruption, and should not be able to shut down a domain controller, do not put Scott in the Backup Operators group. Instead, create a group and assign it only the Backup Files And Directories user right, and then add Scott as a member.

Special Identities

Windows and AD DS also support special identities, which are groups for which membership is controlled by the operating system. You cannot view the groups in any list (in Active Directory Users and Computers, for example), you cannot view or modify the membership of these special identities, and you cannot add them to other groups. You can, however, use these groups to assign rights and permissions.

- Special identities:
 - Are groups for which membership is controlled by the operating system
 - Can be used by the Windows Server operating system to provide access to resources:
 - Based on the type of authentication or connection
 - Not based on the user account
- Important special identities include:
 - Anonymous Logon
 - Authenticated Users
 - Everyone
 - Interactive
 - Network
 - Creator Owner

The most important special identities—often referred to as *groups* (for convenience)—are described in the following list:

- **Anonymous Logon.** This identity represents connections to a computer and its resources that are made without supplying a user name and password. Before Windows Server 2003, this group was a member of the Everyone group. Beginning with Windows Server 2003, this group is no longer a default member of the Everyone group.
- **Authenticated Users.** This represents identities that have been authenticated. This group does not include Guest, even if the Guest account has a password.
- **Everyone.** This identity includes Authenticated Users and the Guest account. (On computers that are running versions of the Windows Server operating system that precede Windows Server 2003, this group includes Anonymous Logon.)
- **Interactive.** This represents users accessing a resource while logged on locally to the computer that is hosting the resource, as opposed to accessing the resource over the network. When a user accesses any given resource on a computer to which the user is logged on locally, the user is added to the Interactive group automatically for that resource. Interactive also includes users who log on through a Remote Desktop connection.
- **Network.** This represents users accessing a resource over the network, as opposed to users who are logged on locally at the computer that is hosting the resource. When a user accesses any given resource over the network, the user is added automatically to the Network group for that resource.
- **Creator Owner.** This represents the security principal that created an object.

The importance of these special identities is that you can use them to provide access to resources based on the type of authentication or connection, rather than the user account. For example, you could create a folder on a system that allows users to view its contents when they are logged on locally to the system, but that does not allow the same users to view the contents from a mapped drive over the network. You could achieve this by assigning permissions to the Interactive special identity.

Demonstration: Managing Groups

This demonstration shows you how to:

- Create a new group.
- Add members to the group.
- Add a user to the group.
- Change the group type and scope.
- Modify the group's Managed By property.

Demonstration Steps

Create a new group

1. On LON-DC1, open the Active Directory Administrative Center.
2. Create a new global security group in the IT OU called **IT Managers**.

Add members to the group

- Add multiple users to the new group.

Add a user to the group

- Add **Ed Meadows** to the **IT Managers** group.

Change the group type and scope

- In the Properties of the IT Managers group, change the Group Scope to **Universal**, and the Type to **Distribution**.

Modify the group's Managed By property

- Add Ed Meadows to the Managed By list, and then grant him the **Manager can update membership list** permission.

Lesson 3

Managing Computer Accounts

Computers, like users, are security principals:

- They have an account with a logon name and password that Windows Server changes automatically on a periodic basis.
- They authenticate with the domain.
- They can belong to groups, have access to resources, and you can configure them by using Group Policy.

A computer account begins its life cycle when you create it and join it to your domain. Thereafter, day-to-day administrative tasks include the following:

- Configuring computer properties
- Moving the computer between OUs
- Managing the computer itself
- Renaming, resetting, disabling, enabling, and eventually deleting the computer object

It is important that you know how to perform these various computer-management tasks so you can configure and maintain the computer objects within your organization.

Lesson Objectives

After completing this lesson, you will be able to:

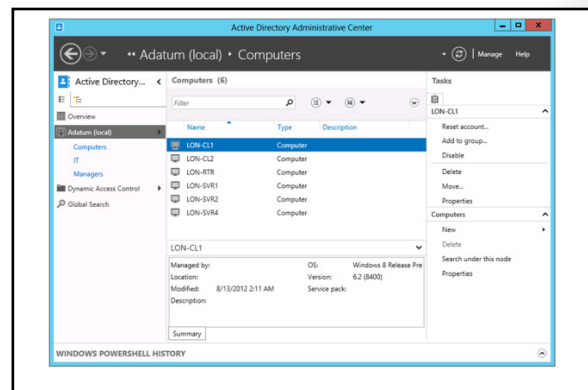
- Explain the purpose of the AD DS Computers container.
- Describe how to configure the location of computer accounts.
- Explain how to control who has permission to create computer accounts.
- Describe how to perform an offline domain join.
- Describe computer accounts and the secure channel.
- Explain how to reset the secure channel.

What Is the Computers Container?

Before you create a computer object in the Directory Service, you must have a place to put it.

When you create a domain, the Computers container is created by default (common name (the **cn** attribute)=Computers). This container is not an OU; instead, it is an object of the container class.

There are subtle but important differences between a container and an OU. You cannot create an OU within a container, so you cannot subdivide the Computers container. You also cannot link a GPO to a container. Therefore, we recommend that you create custom OUs to host computer objects, instead of using the Computers container.



Specifying the Location of Computer Accounts

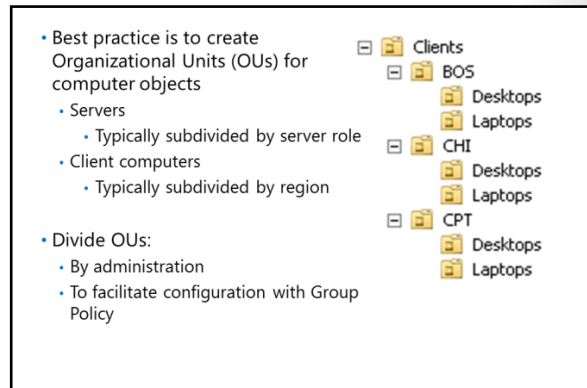
Most organizations create at least two OUs for computer objects: one for servers, and another to host computer accounts for client computers, such as desktops, laptops, and other user systems. These two OUs are in addition to the Domain Controllers OU that is created by default during the AD DS installation.

Computer objects are created in both OUs. There is no technical difference between a computer object in a client's OU and a computer object in a server's or domain controller's OU; computer objects are computer objects. However, separate OUs typically are created to provide unique scopes of management, so that you can delegate management of client objects to one team and management of server objects to another.

Your administrative model might require you to divide your client and server OUs into smaller groups. Many organizations create sub-OUs beneath a server OU, to classify and manage specific types of servers. For example, you might create an OU for file and print servers, and an OU for database servers. By doing so, you can delegate permissions to manage computer objects in the appropriate OU to the team of administrators for each type of server. Similarly, geographically-distributed organizations with local desktop support teams often divide a parent OU for clients into sub-OUs for each site. This approach enables each site's support team to create computer objects in the site for client computers, and to join computers to the domain by using those computer objects.

These specific examples aside, what is most important is that your OU structure reflects your administrative model so that your OUs can provide single points of management for the delegation of administration.

Additionally, by using separate OUs, you can create various baseline configurations by using different GPOs that are linked to the client and the server OUs. With Group Policy, you can specify configuration for collections of computers by linking GPOs that contain configuration instructions to OUs. It is common for organizations to separate clients into desktop and laptop OUs. You then can link GPOs that specify desktop or laptop configuration to the appropriate OUs.



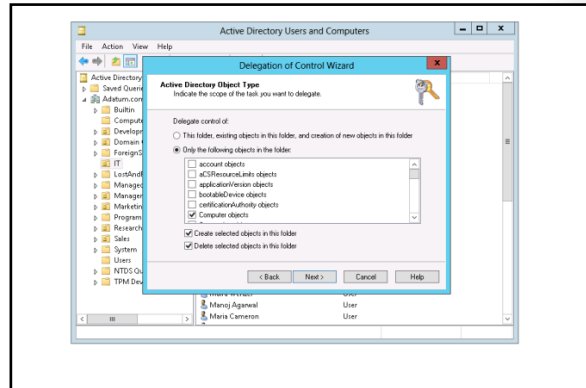
Note: You can use the **Redircmp.exe** command-line tool to reconfigure the default container for computers. For example, if you want to change the default container for computers to an OU called mycomputers, use the following syntax:


```
redircmp ou=mycomputers,DC=contoso,dc=com
```


Controlling Permissions to Create Computer Accounts

Before joining a computer to a domain you should first create a computer object in the appropriate OU. To join a computer to an AD DS domain, three conditions must be met:

- You must have appropriate permissions on the computer object that allow you to join a physical computer with a name that matches the name of the object in AD DS to the domain.
- You must be a member of the local Administrators group on the computer. This allows you to change the computer's domain or workgroup membership.
- You must not have exceeded the maximum number of computer accounts that you can add to the domain. By default, users can add a maximum of ten computers to the domain; this value is known as the *machine account quota* and is controlled by the MS-DS-MachineQuota value. You can modify this value by using the Active Directory Service Interfaces Editor (ADSI Edit) snap-in.



 **Note:** You do not have to create a computer object in the directory service, but it is recommended. Many administrators join computers to a domain without first creating a computer object. However, when you do this, Windows Server attempts to join the domain to an existing object. When Windows Server does not find the object, it fails and creates a computer object in the default Computers container.

The process of creating a computer account in advance is called *pre-staging a computer*. There are two major advantages of pre-staging a computer:

- The account is placed into the correct OU, and is therefore delegated according to the security policy defined by the ACL of the OU.
- The computer is within the scope of GPOs linked to the OU, before the computer joins the domain.

After you have been given permission to create computer objects, you can do so by right-clicking the OU and from the New menu clicking Computer. Next, enter the computer name, following the naming convention of your enterprise, and select the user or group that is allowed to join the computer to the domain with this account. The computer names that you type into the two fields—Computer Name and Computer Name (pre-Windows 2000)—should be the same. Very rarely is there a justification for configuring them separately.

Delegating Permissions

By default, the Enterprise Admins, Domain Admins, Administrators, and Account Operators groups have permission to create computer objects in any new OU. However, as discussed earlier, we recommend that you tightly restrict membership in the first three groups, and that you do not add Administrators to the Account Operators group.

Instead, you should delegate the permission to create computer objects (called Create Computer Objects) to appropriate administrators or support personnel. This permission, which is assigned to the group to which you are delegating administration, allows group members to create computer objects in a specified OU. For example, you might allow your desktop support team to create computer objects in the clients OU, and allow your file server administrators to create computer objects in the file servers OU.

To delegate permissions to create computer accounts, you can use the Delegate Control Wizard to choose a custom task to delegate.

When you delegate permissions to manage computer accounts, you might consider granting additional permissions beyond those required to create computer accounts. For example, you might decide to allow a delegated administrator to manage the properties of existing computer accounts, to delete the computer account, or to move the computer account.



Note: If you want to allow a delegated administrator to move computer accounts, consider that they require the appropriate permissions both in the source AD DS container (where the computer currently exists), and in the target container (where the computer will be moved to). Specifically, they must have Delete Computer permissions in the source container and Create Computer permissions in the target container.

Performing an Offline Domain Join

Typically when you want to join a computer to a domain the computer needs to be able to communicate with an online domain controller. Beginning with Windows Server 2008 R2, Microsoft introduced *offline domain join* which makes it possible for you to join a computer which is running Windows 7, Windows 8, Windows Server 2008 R2, or Windows Server 2012 to a domain without directly communicating with an online domain controller. You can use the command-line tool, Djoin.exe, to join a computer to the domain without contacting an online domain controller; this includes generating a domain join file and then importing it on the client computer.

Offline Domain join can be used to join computers to a Domain when they cannot contact a domain controller.

- Create a domain join file using:

```
djoin.exe /requestODJ /LoadFile <filepath>
/WindowsPath <path to the Windows directory of
the offline image>
```

- Import the domain join file using:

```
djoin.exe /requestODJ /LoadFile <filepath>
/WindowsPath <path to the Windows directory of
the offline image>
```

Starting with Windows 8 and Windows Server 2012 offline domain join can also configure DirectAccess for offsite computers. The DirectAccess policies need to be created before the offline domain join process can include Direct Access settings.

Offline domain join requires the following steps:

1. Djoin.exe is run to provision a computer account in the domain and create the domain join file. The command is:

```
djoin.exe /Provision /Domain <DomainName> /Machine <MachineName> /SaveFile <filepath>
```

If the computer account is not pre-staged it will be created in the Computers container. If the computer is pre-staged then Djoin.exe will also require the **/reuse** option.

2. Djoin.exe is run on the provisioned computer to import the information in the domain join file. The command is:

```
djoin.exe /requestODJ /LoadFile <filepath> /WindowsPath <path to the Windows
directory of the offline image>
```

Optionally you can perform the import on an online operating system using the **/localOS** option. If using the **/localOS** option the **/WindowsPath** option should be set to **%systemroot%** or **%windir%**.

3. Start or Reboot the computer.

When performing a DirectAccess offline domain join, the computer account should be added to the DirectAccessClients group before initially running Djoin.exe and creating the domain join file. Additional Djoin.exe options are available if DirectAccess has been configured to use public key infrastructure (PKI).



Note: To apply the domain join file to a Windows image (.wim) file or virtual hard disk (.vhd or .vhdx) file, first use the Deployment Image Servicing and Management (DISM.exe) command-line tool to mount the image to a file system, and then use the **djoin** command to apply the domain join file. After the Djoin.exe process is complete, you can then use DISM.exe to unmount the image file and prepare the .wim file for deployment

Computer Accounts and Secure Channels

Every member computer in an AD DS domain maintains a computer account with a user name (SamAccountName) and password, just like a user account does. The computer stores its password in the form of a local security authority (LSA) secret, and changes its password with the domain approximately every 30 days. The NetLogon service uses the credentials to log on to the domain, which establishes the secure channel with a domain controller.

Computer accounts and the secure relationships between computers and their domain are robust.

Nevertheless, certain scenarios might arise in which a computer is no longer able to authenticate with the domain. When this happens, users are unable to sign in and the computer cannot access resources, such as Group Policy. Examples of scenarios where this can happen include:

- Computers have accounts
 - sAMAccountName and password
 - Used to create a secure channel between the computer and a domain controller
- Scenarios where a secure channel can be broken
 - Reinstalling a computer, even with same name, generates a new SID and password
 - Restoring a computer from an old backup, or rolling back a computer to an old snapshot
 - Computer and domain disagree about what the password is

- After reinstalling the operating system on a workstation, the workstation is unable to authenticate, even though the technician used the same computer name as was used in the previous installation. Because the new installation generated a new SID, and because the new computer does not know the original computer account password in the domain, it does not belong to the domain and cannot authenticate to the domain.
- A computer has not been used for an extended period, perhaps because the user is on vacation or working away from the office. Computers change their passwords every 30 days, and AD DS remembers the current and previous password. If the computer is unused within this period, authentication can fail.
- A computer's LSA secret gets out of synchronization with the password that the domain knows. You can think of this as the computer forgetting its password. Although it did not forget its password, it just disagrees with the domain over what the password really is. When this happens, the computer cannot authenticate, and the secure channel cannot be created.

The following topic discusses the steps to take when one of these scenarios happens.

Resetting the Secure Channel

Occasionally, the security relationship between a computer account and its domain can be broken; this results in a number of symptoms and errors. The most common signs of computer account problems are:

- Messages at sign-in indicate that a domain controller cannot be contacted, that the computer account might be missing, that the password on the computer account is incorrect, or that the trust relationship (another way of saying *the secure relationship*) between the computer and the domain has been lost.
- Error messages or events in the event log indicate similar problems or suggest that passwords, trusts, secure channels, or relationships with the domain or a domain controller have failed. One such error is "NETLOGON Event ID 3210: Failed To Authenticate," which appears in the computer's event log.
- A computer account is missing in AD DS.

- Do not delete a computer from the domain and rejoin
 - This process creates a new account, resulting in new SID and lost group memberships
- Options for resetting the secure channel
 - Active Directory Users and Computers
 - DsMod.exe
 - NetDom.exe
 - NLTest.exe
 - Windows PowerShell

When the secure channel fails, you must reset it. Many administrators do this by removing the computer from the domain, putting it in a workgroup, and then rejoining the domain. When you remove the computer from the domain, the computer account in AD DS is disabled. When you rejoin the computer to the domain, the same computer account is reused and activated. Do not rename the computer when you join it to the domain.

You can also reset the secure channel between a domain member and the domain by using the following:

- Active Directory Users and Computers
- Active Directory Administrative Center
- the Dsmod.exe command line tool
- the NetDom.exe command line tool
- the Nltest.exe command line tool

If you reset the account, the computer's SID remains the same, and it maintains its group memberships.

To reset the secure channel by using Active Directory Users and Computers:

1. Right-click a computer, and then click **Reset Account**.
2. Click **Yes** to confirm your choice.
3. Rejoin the computer to the domain, and then restart the computer.

To reset the secure channel by using Active Directory Administrative Center:

1. Right-click a computer, and then click **Reset Account**.
2. Click **Yes** to confirm your choice.
3. Rejoin the computer to the domain, and then restart the computer.

To reset the secure channel by using Dsmod:

1. At a command prompt, type the following command:

```
dsmod computer "ComputerDN" -reset
```

2. Rejoin the computer to the domain, and then restart the computer.

To reset the secure channel by using NetDom.exe, type the following command at a command prompt, where the credentials belong to the local Administrators group of the computer:

```
netdom reset MachineName /domain DomainName /User0 UserName /Password0 {Password | *}

```

This command resets the secure channel by attempting to reset the password on both the computer and the domain, so it does not require rejoining or rebooting.

To reset the secure channel by using NLTest.exe, on the computer that has lost its trust, type the following command at a command prompt:

```
nltest /server:servername /sc_reset:domain\domaincontroller

```

You also can use Active Directory module for Windows PowerShell to reset a computer account. The following example demonstrates how to reset the secure channel between the local computer and the domain to which it is joined.

You must run this command on the local computer:

```
Test-ComputerSecureChannel -Repair

```



Note: You also can reset a remote computer's password with Windows PowerShell:

```
invoke-command -computername Workstation1 -scriptblock {reset-computermachinepassword}

```

Bring Your Own Device

With the proliferation of bring your own device programs in corporate environments security considerations must include consumer devices. Typically security for BYOD programs relies on user based security. Windows Server 2012 R2 introduces enhancements to Active Directory Federation Services that you can use to apply security to consumer devices. The four enhancements include:

- **Workplace Join.** With workplace join you can create AD DS objects for consumer devices. The devices do not join the domain; instead, the devices are issued certificates to represent the AD DS objects. Once the AD DS objects are created they can be used to manage security like any other AD DS objects.
- **Multi-factor Access Control.** With multi-factor access control you can configure claims-based authorization rules.
- **Multi-Factor Authentication.** With multi factor authentication you can require users to use more than one form of authentication, for example, forms based authentication and a smartcard based certificate.
- **Web Application Proxy.** Web Application Proxy is a reverse proxy for allowing internal apps to be securely accessed from the web.

AD FS has been enhanced to support bring your own device programs

- **Workplace Join** – Creates an AD DS object for consumer devices
- **Multi-Factor Access Control** – Increases security by using claims-based authorization rules
- **Multi-Factor Authentication** – Increases security by requiring more than one form of authentication
- **Web Application Proxy** – Allows apps to be securely publish to the Internet

Lesson 4

Delegating Administration

Although a single person can manage a small network with a few user and computer accounts, as the network grows, the volume of work that relates to network management grows too. At some point, teams with particular specializations evolve, each with responsibility for some specific aspect of network management. In AD DS environments, it is common practice to create OUs for different departments and geographical regions, and to delegate control of those OUs to different people. It is important that you know why and how to create OUs, and how to delegate administrative tasks to users on objects within those OUs.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the use of OUs.
- Describe AD DS permissions.
- Determine a user's effective AD DS permissions on an AD DS object.
- Delegate administrative control of an AD DS object to a specified user or group of users.

Considerations for Using Organizational Units

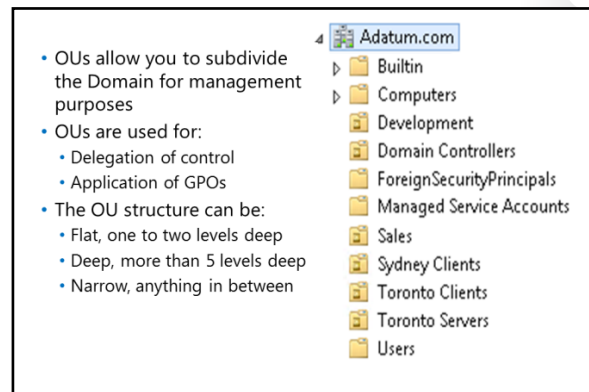
By default the AD DS only contains a single OU, the Domain Controllers OU. While it is possible to manage a small organization without creating additional OUs, even small organizations typically create an OU hierarchy. An OU hierarchy allows you to subdivide the administration of your domain for management purposes. The best practice for creating an OU hierarchy is to create OUs to organize objects for:

- Delegation of control
- Application of GPOs
- Both delegation of control and application of GPOs

When you design an OU hierarchy there are many different strategies that you can follow. You can create a flat, wide structure that has OUs that are only one or two levels deep; you can create a deep, narrow structure nesting five or more levels of OUs; or you can create anything in between. The key factor in designing your OU hierarchy is that it should help your manage your organization.

How you design your OU hierarchy could be based on any of the following:

- Geographic location. There could be local IT staff for delegating management, local regulations that require specific policies, or a myriad of other factors.
- Departmental characteristics. Typically different departments are managed differently and have unique requirements.
- Resource type. Some organizations create separate OUs for different resources. File servers are typically managed differently than SQL servers and require different policies applied to them.



- Management structure. Some organizations want their OU hierarchy to mirror their management structure.
- Any combination of the above. There is no one right way to design your OU hierarchy.

For example, you might have a diverse organization with offices in many geographic locations, and there are sufficient IT staff at most locations. You could create top level OUs based on these geographic locations and delegate control of these OUs to the local IT staff. Each of these could have child OUs based on the departments in those locations with GPOs applied to those OUs to enforce departmental settings. Another design for the same organization could have the top level OUs representing the departmental structure with child OUs representing locations.

AD DS Permissions

All AD DS objects, such as users, computers, and groups, can be secured by using a list of permissions. The permissions on an object are called access control entries (ACEs), and they are assigned to users, groups, or computers, which are also known as security principals. ACEs are saved in the object's discretionary access control list (DACL), which is part of the object's ACL. The ACL contains the system access control list (SACL) that includes auditing settings.

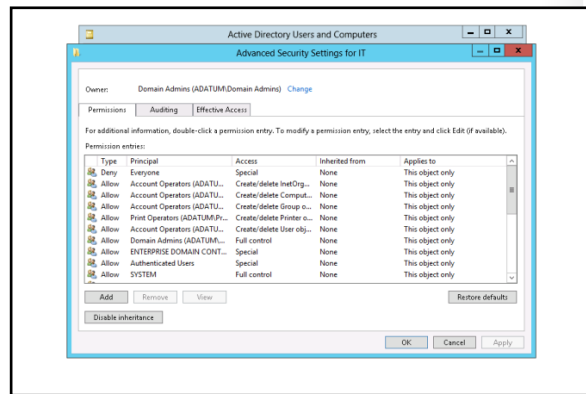
Each object in AD DS has its own ACL. If you have sufficient permissions, you can modify the permissions to control the level of access on a specific AD DS object. If you have sufficient permissions, you can also delegate administrative control—that is, just as you can give a group the ability to change files in a folder, you can give a group the ability to reset passwords on user objects, for example.

You can also use the DACL of an object to assign permissions to an object's specific properties. For example, you can allow (or deny) permission to **Read phone and email options** or **Write phone and email options**. While you can do this with a single checkbox, this is actually a property set that includes multiple, specific properties. Using property sets, you can easily manage permissions to commonly used collections of properties. However, you can also assign more granular permissions and allow or deny permission to change just some of the information, such as the mobile telephone number or the street address.

Assigning the help desk permission to reset passwords for each individual user object is tedious. Even so, in AD DS, it is not a good practice to assign permissions to individual objects. Instead, you should assign permissions at the level of OUs.

The permissions that you assign to an OU are inherited by all objects in the OU. Therefore, if you give the help desk permission to reset passwords for user objects and attach that permission to the OU that contains the users, all user objects within that OU will inherit that permission. In just one step, you have delegated that administrative task.

Child objects inherit the permissions of the parent container or OU. That container or OU in turn inherits its permissions from its parent container or OU. If it is a first-level container or OU, it inherits the permissions from the domain itself. The reason child objects inherit permissions from their parents is that, by default, each new object is created with the Include inheritable permissions from this object's parent option enabled.



Effective AD DS Permissions

Effective permissions are the resulting permissions for a security principal (such as a user or group), based on the cumulative effect of each inherited and explicit ACE. Your ability to reset a user's password, for example, might be due to your membership in a group that is allowed the Reset Password permission on an OU several levels above the user object. The inherited permission assigned to a group to which you belong results in an effective permission of Allow: Reset Password. Your effective permissions can be complicated when you consider Allow and Deny permissions, explicit and inherited ACEs, and the fact that you might belong to multiple groups, each of which might be assigned different permissions.

Permissions assigned to users and groups accumulate

Best practice is to assign permissions to groups, not to individual users

In the event of conflicts:

- Deny permissions override Allow permissions
- Explicit permissions override Inherited permissions
- Explicit Allow overrides Inherited Deny

To evaluate effective permissions, you can use:

- The Effective Access tab
- Manual analysis

Permissions, whether assigned to your user account or to a group to which you belong, are equivalent. This means that ultimately an ACE applies to you, the user. The best practice is to manage permissions by assigning them to groups, but it is also possible to assign ACEs to individual users or computers. A permission that has been assigned directly to you, the user, is neither more important nor less important than a permission assigned to a group to which you belong.

The Allow permissions, which allow access, are cumulative. When you belong to several groups, and when those groups have been granted permissions that allow a variety of tasks, you will be able to perform all of the tasks assigned to all of those groups, and tasks assigned directly to your user account.

Deny permissions, which deny access, override equivalent Allow permissions. If you are in one group that has been allowed the permission to reset passwords, and you are also another group that has been denied permission to reset passwords, the Deny permission prevents you from resetting passwords.



Note: Use Deny permissions rarely. In fact, it is unnecessary to assign Deny permissions, because if you do not assign an Allow permission, users cannot perform the task. Before assigning a Deny permission, check to see if you could achieve your goal instead by removing an Allow permission. For example, if you want to delegate an Allow permission to a group, but exempt only one member from that group, you can use a Deny permission on that specific user account while the group still has an Allow permission.

Every permission is granular. Even if you have been denied the ability to reset passwords, you might still have the ability through other Allow permissions to change the user's logon name or email address.

Because child objects inherit the inheritable permissions of parent objects by default, and because explicit permissions can override inheritable permissions, an explicit Allow permission will actually override an inherited Deny permission.

Unfortunately, the complex interaction of user, group, explicit, inherited, Allow, and Deny permissions can make evaluating effective permissions tedious. You can use the permissions reported by the **dsaccs** command, or listed on the Effective Access tab of the Advanced Security Settings dialog box to begin evaluating effective permissions, but it is still a manual task.

Demonstration: Delegating Administrative Permissions

This demonstration shows you how to:

- Create an OU.
- Move Objects into an OU.
- Delegate a standard task.
- Delegate a custom task.
- View AD DS permissions resulting from these delegations.

Demonstration Steps

Create an OU

1. Open **Active Directory Users and Computers**.
2. Create an OU named **Executives** in the Adatum.com root.



Note: Discuss the **Protect Container From Accidental Deletion** setting.

Move users into the Executives OU

- Move all the users from **Carol Troup** to **Euan Garden** into the **Executives OU**.

Delegate a standard task

- Use the Delegate Control Wizard to grant the IT group permissions to perform the following standard management tasks on the Executives OU:
 - **Create, delete, and manage user accounts**
 - **Reset user passwords and force password change at next logon**
 - **Read all user information**

Delegate a custom task

- Use the Delegate Control Wizard to grant the following permissions on the Executives OU to the IT group:
 - **Full Control on computer objects**
 - **Create computer objects**
 - **Delete computer objects**

View AD DS permissions resulting from these delegations

1. Enable the **Advanced Features** view in Active Directory Users and Computers.
2. View the **Properties** of the Executives OU.
3. Use the **Security** tab to verify the assigned permissions.
4. Close all open windows.

Lab: Managing Active Directory Domain Services Objects

Scenario

You have been working for A. Datum as a desktop support specialist and have visited desktop computers to troubleshoot app and network problems. You have recently accepted a promotion to the server support team. One of your first assignments is configuring the infrastructure service for a new branch office.

To begin deployment of the new branch office, you are preparing AD DS objects. As part of this preparation, you need to create an OU for the branch office and delegate permission to manage it. Then you need to create users and groups for the new branch office. Finally, you need to reset the secure channel for a computer account that has lost connectivity to the domain in the branch office.

Objectives

After completing this lab, you will be able to:

- Delegate administration for a branch office.
- Create and configure user accounts in AD DS.
- Manage computer objects in AD DS.

Lab Setup

Estimated time: 60 minutes

Virtual machines	20410C-LON-DC1 20410C-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 **20410C-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in using the following credentials:
 - User name: **Administrator**
 - Password: **Pa\$\$w0rd**
 - Domain: **Adatum**
5. Repeat steps 2 to 4 for 20410C-LON-CL1.

Exercise 1: Delegating Administration for a Branch Office

Scenario

A. Datum delegates management of each branch office to a specific group. This allows an employee who works onsite to be configured as an administrator when required. Each branch office has a branch administrators group that is able to perform full administration within the branch office OU. There is also a branch office help desk group that is able to manage users in the branch office OU, but not other objects. You need to create these groups for the new branch office and delegate permissions to the groups.

The main tasks for this exercise are as follows:

1. Delegate administration for Branch Administrators.
2. Delegate a user administrator for the Branch Office Help Desk.
3. Add a member to the Branch Administrators.
4. Add a member to the Branch Help Desk group.

► Task 1: Delegate administration for Branch Administrators

1. On LON-DC1, open Active Directory Users and Computers, and create in the Adatum.com domain a new OU named **Branch Office 1**.
2. Create the following global security groups in the **Branch Office 1** OU:
 - **Branch 1 Help Desk**
 - **Branch 1 Administrators**
 - **Branch 1 Users**
3. Move **Holly Dickson** from the **IT** OU to the **Branch Office 1** OU.
4. Move the following users to the **Branch Office 1** OU:
 - **Development\Bart Duncan**
 - **Managers\Ed Meadows**
 - **Marketing\Connie Vrettos**
 - **Research\Barbara Zighetti**
 - **Sales\Arlene Huff**
5. Move the LON-CL1 computer to the **Branch Office 1** OU, and then restart the LON-CL1 computer.
6. Sign in to LON-CL1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
7. On LON-DC1, in Active Directory Users and Computers, use the Delegate Control Wizard to delegate administration of the **Branch Office 1** OU to the **Branch 1 Administrators** security group by delegating the following common and custom tasks:
 - a. Delegate the following common tasks:
 - **Create, delete, and manage user accounts**
 - **Reset user passwords and force password change at next logon**
 - **Read all user information**
 - **Create, delete and manage groups**
 - **Modify the membership of a group**
 - **Manage Group Policy links**

b. Delegate the following custom tasks:

- **Create and delete computer objects in the current OU**
- **Full control of computer objects in the current OU**

► **Task 2: Delegate a user administrator for the Branch Office Help Desk**

1. On LON-DC1, in Active Directory Users and Computers, use the Delegate Control Wizard to delegate administration of the Branch Office 1 OU to the Branch 1 Help Desk security group.
2. Delegate the following common tasks:
 - **Reset user passwords and force password change at next logon**
 - **Read all user information**
 - **Modify the membership of a group**

► **Task 3: Add a member to the Branch Administrators**

1. On LON-DC1, add **Holly Dickson** to the **Branch 1 Administrators** global group.
2. Add the **Branch 1 Administrators** global group to the **Server Operators** domain local group.
3. Sign out from LON-DC1.
4. Sign in as **Adatum\Holly** with the password **Pa\$\$w0rd**.

You can sign in locally at a domain controller because Holly belongs indirectly to the Server Operators domain local group.

5. In Server Manager, open **Active Directory Users and Computers**.
Confirm Holly's current credentials in the **User Account Control** dialog box.
6. Attempt to delete **Sales\Aaren Ekelund**.
You are unsuccessful, because Holly lacks the required permissions.
7. Try to delete **Branch Office 1\Ed Meadows**.
You are successful, because Holly has the required permissions.

► **Task 4: Add a member to the Branch Help Desk group**

1. On LON-DC1, add **Bart Duncan** to the **Branch 1 Help Desk** global group.
2. Close Active Directory Users and Computers, and then close Server Manager.
3. Open Server Manager, and then open **Active Directory Users and Computers**.
4. In the **User Account Control** dialog box, specify **Adatum\Administrator** and **Pa\$\$w0rd** as the required credentials.



Note: To modify the Server Operators membership list, you must have permissions beyond those available to the Branch 1 Administrators group.

5. Add the **Branch 1 Help Desk** global group to the **Server Operators** domain local group.
6. Sign out from LON-DC1.
7. Sign in as **Adatum\Bart** with the password **Pa\$\$w0rd**.

You can sign in locally at a domain controller because Bart belongs, indirectly, to the Server Operators domain local group.

8. Open Server Manager, and then open **Active Directory Users and Computers**. Confirm your current credentials in the **User Account Control** dialog box.
9. Try to delete **Branch Office 1\Connie Vrettos**.
You are unsuccessful, because Bart lacks the required permissions.
10. Reset Connie's password to **Pa\$\$w0rd**.
11. After confirming the password reset is successful, sign out from LON-DC1.
12. Sign in to LON-DC1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.

Results: After completing this exercise, you will have successfully created an OU, and delegated administration of it to the appropriate group.

Exercise 2: Creating and Configuring User Accounts in AD DS

Scenario

You have been given a list of new users for the branch office, and you need to begin creating user accounts for them.

The main tasks for this exercise are as follows:

1. Create a user template for the branch office.
2. Configure the template settings.
3. Create a new user for the branch office, based on the template.
4. Sign in as a user to test account settings.

► Task 1: Create a user template for the branch office

1. On LON-DC1, create a folder called **C:\branch1-userdata**, and then share it.
2. Modify the shared folder permissions so that the **Everyone** group has **Full Control Allow** permissions.
3. In Server Manager, open **Active Directory Users and Computers**, and then create a new user with the following properties in the **Branch Office 1** OU:
 - Full name: **_Branch_template**
 - User logon name: **_Branch_template**
 - Password: **Pa\$\$w0rd**
 - **Account is disabled**

► Task 2: Configure the template settings

- On LON-DC1, modify the following properties of the **_Branch_template** account:
 - City: **Slough**
 - Group: **Branch 1 Users**
 - Home folder: **\\lon-dc1\branch1-userdata\%username%**

► Task 3: Create a new user for the branch office, based on the template

1. On LON-DC1, copy the **_Branch_template** user account, and configure the following properties:
 - First name: **Ed**
 - Last name: **Meadows**
 - Password: **Pa\$\$w0rd**
 - **User must change password at next logon** is cleared
 - **Account is disabled** is cleared
2. Verify that the following properties have been copied during account creation:
 - City: **Slough**
 - Home folder path: **\\lon-dc1\branch1-userdata\Ed**
 - Group: **Branch 1 Users**
3. Sign out from LON-DC1.

► Task 4: Sign in as a user to test account settings

1. Switch to LON-CL1 and sign off.
2. Sign in to LON-CL1 as **Adatum\Ed** with the password **Pa\$\$w0rd**.
You are able to sign in successfully.
3. Verify that you have a drive mapping for drive Z to Ed's home folder on LON-DC1.
4. Sign out of LON-CL1.

Results: After completing this exercise, you will have successfully created and tested a user account created from a template.

Exercise 3: Managing Computer Objects in AD DS

Scenario

A workstation has lost its connectivity to the domain and cannot authenticate users properly. When users attempt to access resources from this workstation, access is denied. You need to reset the computer account to recreate the trust relationship between the client and the domain.

The main tasks for this exercise are as follows:

1. Reset a computer account.
2. Observe the behavior when a client logs on.
3. Rejoin the domain to reconnect the computer account.

► Task 1: Reset a computer account

1. On LON-DC1, sign in as **Adatum\Holly** with the password **Pa\$\$w0rd**.
2. Open **Active Directory Users and Computers**.
3. Confirm your credentials in the **User Account Control** dialog box.
4. Navigate to **Branch Office 1**.
5. Reset the **LON-CL1** computer account.

► **Task 2: Observe the behavior when a client logs on**

1. Switch to LON-CL1 and attempt to sign in as **Adatum\Ed** with the password **Pa\$\$w0rd**.

A message appears stating that **The trust relationship between this workstation and the primary domain failed**.

2. Click **OK** to acknowledge the message.

► **Task 3: Rejoin the domain to reconnect the computer account**

1. Sign in to LON-CL1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Open Control Panel, switch to **Large icons** view, and then open **System**.
3. View the **Advanced system settings**, and then click the **Computer Name** tab.
4. In the **System Properties** dialog box, use the **Network ID** button to rejoin the computer to the domain.
5. Complete the wizard using the following settings:
 - User name: **Administrator**
 - Password: **Pa\$\$w0rd**
 - Domain: **Adatum**
 - Do you want to enable a domain user account on this computer: **No**
6. When prompted, restart the computer.
7. Sign in as **Adatum\Ed** with the password of **Pa\$\$w0rd**.

You are successful because the computer had been successfully rejoined.

Results: After completing this exercise, you will have successfully reset a trust relationship.

► **Prepare for the next module**

When you have completed the lab, revert the virtual machines back 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 **20410C-LON-CL1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for 20410C-LON-DC1.

Lab Review Questions

Question: What are the options for modifying the attributes of new and existing users?

Question: What types of objects can be members of global groups?

Question: What types of objects can be members of domain local groups?

Question: What are the two credentials that are necessary for any computer to join a domain?

Module Review and Takeaways

Module Review Questions

Question: A company with branches in multiple cities has members of a sales team that travel frequently between domains. Each of these domains has their own printers that are managed by using domain local groups. How can you provide these members with access to the various domains printers?

Question: You are responsible for managing accounts and access to resources for your group members. A user in your group transfers to another department within the company. What should you do with the user's account?

Question: What is the main difference between the Computers container and an OU?

Question: When should you reset a computer account? Why is it better to reset the computer account rather than to disjoin and then rejoin it to the domain?

Question: A project manager in your department is starting a group project that will continue for the next year. Several users from your department and other departments will be dedicated to the project during this time. The project team must have access to the same shared resources. The project manager must be able to manage the user accounts and group accounts in AD DS; however, you do not want to give the project manager permission to manage anything else in AD DS. What is the best way to do this?

Question: You are working as an IT technician in Contoso, Ltd. You are managing the Windows Server-based infrastructure. You have to find a method for joining new Windows 8.1-based computers to a domain during the installation process, without intervention of a user or an administrator. What is the best way to do this?

Best Practices

Best Practices for User Account Management

- Do not let users share user accounts. Always create a user account for each individual, even if that person will not be with your organization for long.
- Educate users about the importance of password security.
- Ensure that you choose a naming strategy for user accounts that enables you to identify the user to whom the account relates. Also ensure that your naming strategy uses unique names within your domain.

Best Practices for Group Management

- When managing access to resources, try to use both domain local group and role groups.
- Use universal groups only when necessary because they add weight to replication traffic.
- Use Windows PowerShell with Active Directory Module for batch jobs on groups.
- Avoid adding users to built-in and default groups.

Best Practices Related to Computer Account Management

- Always provision a computer account before joining computers to a domain, and then place them in appropriate OU.
- Redirect the default Computers container to another location.
- Reset the computer account, instead of disjoining and rejoining.
- Integrate the offline domain join functionality with unattended installations.

Tools

Tool	Use	Where to find it
Active Directory Users and Computers	Manage groups	Administrative Tools
Active Directory module for Windows Power Shell	Manage groups	Installed as Windows Feature
Active Directory module for Windows PowerShell	Computer account management	Administrative Tools
Djoin.exe	Offline domain join	Command line
Redircmp.exe	Change default computer container	Command line
DsacIs	View and modify AD DS permissions	Command line

Module 4

Automating Active Directory Domain Services Administration

Contents:

Module Overview	4-1
Lesson 1: Using Command-line Tools for AD DS Administration	4-2
Lesson 2: Using Windows PowerShell for AD DS Administration	4-8
Lesson 3: Performing Bulk Operations with Windows PowerShell	4-14
Lab: Automating AD DS Administration by Using Windows PowerShell	4-25
Module Review and Takeaways	4-30

Module Overview

You can use command-line tools and Windows PowerShell® to automate Active Directory® Domain Services (AD DS) administration. Automating administration speeds up processes that you might otherwise perform manually. Windows PowerShell includes cmdlets for performing AD DS administration and for performing bulk operations. You can use bulk operations to change many AD DS objects in a single step rather than updating each object manually.

Objectives

After completing this module, you will be able to:

- Use command-line tools for AD DS administration.
- Use Windows PowerShell cmdlets for AD DS administration.
- Perform bulk operations by using Windows PowerShell.

Lesson 1

Using Command-line Tools for AD DS Administration

Windows Server® 2012 R2 includes several command-line tools that you can use to perform AD DS administration. Many organizations create scripts that use command-line tools to automate the creation and management of AD DS objects, such as user accounts and groups. You must understand how to use these command-line tools to ensure that if required, you can modify the scripts that your organization uses.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the benefits of using command-line tools for AD DS administration.
- Describe how and when to use **csvde**.
- Describe how and when to use **ldifde**.
- Describe how and when to use Directory Services (DS) commands.

Benefits of Using Command-Line Tools for AD DS Administration

Many administrators prefer to use graphical tools for AD DS administration whenever possible. Graphical tools, such as the Active Directory Users and Computers snap-in, are intuitive to use because they represent information visually and provide options in the form of radio buttons and dialog boxes. When information is represented graphically, you do not need to memorize syntax.

Graphical tools work well in many situations, but they cannot be automated. To automate AD DS administration, you require command-line tools. Command-line tools can be used in scripts, or they can be used by other apps and programs.

Some benefits of using command-line tools are:

- Faster implementation of bulk operations. For example, you can export a list of new user accounts from a human resources app. You can then use a command-line tool or script to create the new user accounts based on the exported information. This is much faster than creating each new user account manually.
- Customized processes for AD DS administration. You can use a customized graphical program to gather information about a proposed new group, and then create the new group. When the information is gathered, the graphical program can verify that the information format—such as the naming convention—is correct. Then, the graphical program uses a command-line tool to create the new group. This process allows company-specific rules to be enforced.
- AD DS administration on Server Core. The Server Core installation of Windows Server cannot run graphical administration tools such as Active Directory Users and Computers. However, you can use command-line tools on Server Core.

Command-line tools allow you to automate AD DS administration

Benefits of using command-line tools:

- Faster implementation of bulk operations
- Customized processes for AD DS administration
- AD DS administration on server core

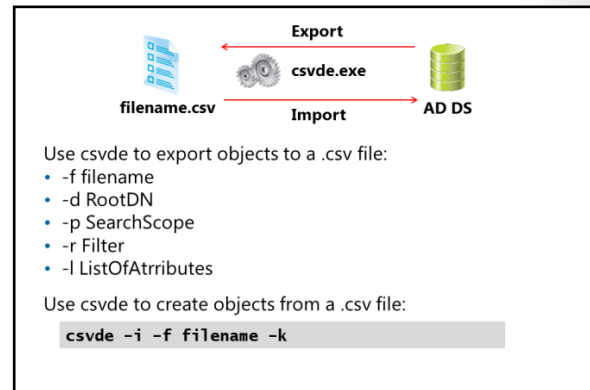


Note: You can also administer Server Core remotely by using graphical tools.

What Is Csvde?

Csvde is a command-line tool that exports or imports AD DS objects to or from a comma-separated values (.csv) file. Many programs and apps are capable of exporting or importing data from .csv files. This makes csvde useful for interoperability with other programs and apps, such as databases or spreadsheets.

The main limitation of csvde is that it cannot modify existing Active Directory objects; it can only create new objects. For example, you can use csvde to create a set of new user accounts, but you cannot use it to modify the properties of the user accounts after they are created. You can also use csvde to export object properties, such as a list of users and their email addresses.



Export Objects by Using csvde

To export objects by using csvde, as a minimum, you need to specify the filename of the .csv file to which data will be exported. With only the filename specified, all objects in the domain will be exported.

The basic syntax to use csvde for export is:

```
csvde -f filename
```

Other options that you can use with csvde are listed in the following table.

Option	Description
-d RootDN	Specifies the distinguished name of the container from which the export will begin. The default is the domain.
-p SearchScope	Specifies the scope of the search relative to the container specified by the option -d . The SearchScope option can be: <ul style="list-style-type: none"> • base (this object only) • onelevel (objects within this container) • subtree (this container and all subcontainers). This is the default value.
-r Filter	Limits the objects returned to those that match the filter. The filter is based on Lightweight Directory Access Protocol (LDAP) query syntax.
-l ListOfAttributes	Specifies the attributes to be exported. Use the LDAP name for each attribute, and separate them with commas.

After the export completes, the .csv file will contain a header row and one row for each object that was exported. The header row is a comma-separated list with the names of the attributes for each object.

Create Objects by Using csvde

The basic syntax for using csvde to create objects is:

```
csvde -i -f filename -k
```

The **-i** parameter specifies import mode. The **-f** parameter identifies the file name from which to import. The **-k** parameter instructs csvde to ignore error messages, including the "Object Already Exists" error message. The **suppress errors** option is useful when importing objects to ensure that all of the objects possible are created, instead of stopping when partially complete.

The .csv file that is being used for an import must have a header row that contains names of LDAP attributes for the data in the .csv file. Each row must contain exactly the correct number of items as specified in the header row.

You cannot use csvde to import passwords, because passwords in a .csv file are not protected. As a result, user accounts that you create with csvde have a blank password and are disabled.



Note: For more information about parameters for csvde, at a command prompt, type **csvde /?**, and then press Enter.

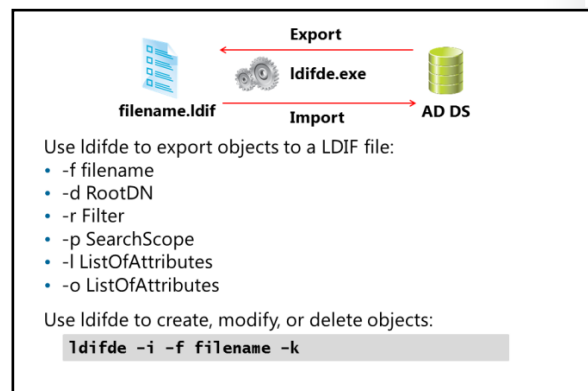


Additional Reading: For more information about LDAP query syntax, see LDAP Query Basics at <http://go.microsoft.com/fwlink/?LinkId=168752>.

What Is Ldifde?

Ldifde is a command-line tool that you can use to export, create, modify, or delete AD DS objects. Like csvde, Ldifde uses data that is stored in a file. The file must be in LDAP Data Interchange Format (LDIF). Most programs and apps cannot export or import data in LDIF format. It is more likely that you will obtain data in LDIF format from another directory service.

An LDIF file is text-based, with blocks of lines composing a single operation such as creating or modifying a user object. Each line within the operation specifies something about the operation, such as an attribute or the type of operation. A blank line separates multiple operations within the LDIF file.



The following is an example of an LDIF file that creates a single user:

```
dn: CN=Bonnie Kearney,OU=Employees,OU=User Accounts,DC=adatum,DC=com
changetype: add
objectClass: top
objectClass: person
objectClass: organizationalPerson
objectClass: user
cn: Bonnie Kearney
sn: Kearney
title: Operations
description: Operations (London)
givenName: Bonnie
displayName: Kearney, Bonnie
company: Contoso, Ltd.
sAMAccountName: bonnie.kearney
userPrincipalName: bonnie.kearney@adatum.com
mail: bonnie.kearney@adatum.com
```

For each operation in an LDIF file, the **changetype** line defines the operation to be performed. The valid values are **add**, **modify**, or **delete**.

Export Objects by Using Ldifde

When using Ldifde to export objects, the minimum information you must provide is a filename to hold the data. When no other options are selected, all objects in the domain are exported.

The basic syntax for exporting objects by using Ldifde is:

```
ldifde -f filename
```

Some other options you can use when exporting objects Ldifde are listed in the following table.

Option	Description
-d RootDN	The root of the LDAP search. The default is the root of the domain.
-r Filter	An LDAP search filter that limits the results returned.
-p SearchScope	The scope, or depth, of the search. This can be: <ul style="list-style-type: none"> • subtree (the container and all child containers) • base (the immediate child objects of the container only) • onelevel (the container and its immediate child containers)
-l ListOfAttributes	A comma-separated list of attributes to include in the export.
-o ListOfAttributes	A comma-separated list of attributes to exclude in the export.

Import Objects by Using Ldifde

When you use Ldifde to import objects, you must specify the operation to perform on the object. For each operation in an LDIF file, the **changetype** line defines the operation to be performed.

The basic syntax for using Ldifde to import objects is:

```
ldifde -i -f filename -k
```

The **-i** parameter specifies import mode. The **-f** parameter identifies the file name from which to import. The **-k** parameter instructs `ldifde` to ignore errors, including the “*Object Already Exists*” error. The option **suppress errors** is useful when importing objects to ensure that all objects possible are created, instead of stopping when partially complete.

You cannot use `ldifde` to import passwords, because passwords in an LDIF file are not secure. As a result, user accounts created by `ldifde` have a blank password and are disabled.

What Are DS Commands?

Windows Server 2012 includes command-line tools called *DS commands*, which are suitable for use in scripts. You can use DS command-line tools to create, view, modify, and remove AD DS objects. The following table describes DS command-line tools.

Windows Server 2012 includes command-line tools that are suitable for use in scripts

- Examples
- To modify the department of a user account, type:
`Dsmod user "cn=Joe Healy,ou=Managers,dc=adatum,dc=com" -dept IT`
- To display the email of a user account, type:
`Dsget user "cn=Joe Healy,ou=Managers,dc=adatum,dc=com" -email`
- To delete a user account, type:
`Dsrm "cn=Joe Healy,ou=Managers,dc=adatum,dc=com"`
- To create a new user account, type:
`Dsadd user "cn=Joe Healy,ou=Managers,dc=adatum,dc=com"`

Tool	Description
DSadd	Creates AD DS objects.
DSget	Displays properties of AD DS objects.
DSquery	Searches for AD DS objects.
DSmod	Modifies AD DS objects.
DSrm	Removes AD DS objects.
DSmove	Moves AD DS objects.

User Management Command Examples

The following are examples of DS commands that you could type at a command prompt.

To modify the department of a user account, type:

```
dsmod user "cn=Joe Healy,ou=Managers,dc=adatum,dc=com" -dept IT
```

To display the email of a user account, type:

```
dsget user "cn=Joe Healy,ou=Managers,dc=adatum,dc=com" -email
```

To delete a user account, type:

```
dsrm "cn=Joe Healy,ou=Managers,dc=adatum,dc=com"
```

To create a new user account, type:

```
dsadd user "cn=Joe Healy,ou=Managers,dc=adatum,dc=com"
```

Question: What criteria would you use to select between using csvde, ldifde, and the DS commands?

Lesson 2

Using Windows PowerShell for AD DS Administration

Windows PowerShell is the preferred scripting environment in Windows Server 2012. It is much easier to use than previous scripting languages such as Microsoft® Visual Basic Scripting Edition (VBScript). Windows PowerShell includes an extensive list of cmdlets to manage AD DS objects. You can use cmdlets to create, modify, and remove user accounts, groups, computer accounts, and organizational units (OUs).

Lesson Objectives

After completing this lesson, you will be able to:

- Use Windows PowerShell cmdlets to manage user accounts.
- Use Windows PowerShell cmdlets to manage groups.
- Use Windows PowerShell cmdlets to manage computer accounts.
- Use Windows PowerShell cmdlets to manage OUs.

Using Windows PowerShell Cmdlets to Manage User Accounts

You can use Windows PowerShell cmdlets to create, modify, and delete user accounts. You can use these cmdlets for individual operations or as part of a script to perform bulk operations. Some of the cmdlets for managing user accounts are in the following table.

Cmdlet	Description
New-ADUser	Creates user accounts
Set-ADUser	Modifies properties of user accounts
Remove-ADUser	Deletes user accounts
Set-ADAccountPassword	Resets the password of a user account
Set-ADAccountExpiration	Modifies the expiration date of a user account
Unlock-ADAccount	Unlocks a user account after it has become locked after too many incorrect login attempts
Enable-ADAccount	Enables a user account
Disable-ADAccount	Disables a user account


```
New-ADUser "Sten Faerch" -AccountPassword (Read-Host "Enter password") -Department IT
```

Cmdlet	Description
New-ADUser	Creates user accounts.
Set-ADUser	Modifies properties of user accounts.
Remove-ADUser	Deletes user accounts.
Set-ADAccountPassword	Resets the password of a user account.
Set-ADAccountExpiration	Modifies the expiration date of a user account.
Unlock-ADAccount	Unlocks a user account when it is locked after exceeding the accepted number of incorrect login attempts.
Enable-ADAccount	Enables a user account.
Disable-ADAccount	Disables a user account.

Create New User Accounts

When you use the **New-ADUser** cmdlet to create new user accounts, you can set most user properties including a password. For example:

- If you do not use the **-AccountPassword** parameter, no password is set and the user account is disabled. The **-Enabled** parameter cannot be set as **\$true** when no password is set.
- If you use the **-AccountPassword** parameter to specify a password, then you must specify a variable that contains the password as a secure string, or choose to be prompted for the password. A secure string is encrypted in memory. If you set a password then you can enable the user account by setting the **-Enabled** parameter as **\$true**.

Some commonly used parameters for the **New-ADUser** cmdlet are listed in the following table.

Parameter	Description
AccountExpirationDate	Defines the expiration date for the user account.
AccountPassword	Defines the password for the user account.
ChangePasswordAtLogon	Requires the user account to change passwords at the next logon.
Department	Defines the department for the user account.
Enabled	Defines whether the user account is enabled or disabled.
HomeDirectory	Defines the location of the home directory for a user account.
HomeDrive	Defines the drive letters that are mapped to the home directory for a user account.
GivenName	Defines the first name for a user account.
Surname	Defines the last name for a user account.
Path	Defines the OU or container where the user account will be created.

The following is an example of a command that you could use to create a user account with a prompt for a password:

```
New-ADUser "Sten Faerch" -AccountPassword (Read-Host -AsSecureString "Enter password")
-Department IT
```

Question: Are all cmdlet parameters that you use to manage user accounts the same?

Using Windows PowerShell Cmdlets to Manage Groups

You can use Windows PowerShell to create, modify, and delete groups. You can use these cmdlets for individual operations or as part of a script to perform bulk operations. Some of the cmdlets for managing groups are listed in the following table.

Cmdlet	Description
New-ADGroup	Creates new groups
Set-ADGroup	Modifies properties of groups
Get-ADGroup	Displays properties of groups
Remove-ADGroup	Deletes groups
Add-ADGroupMember	Adds members to groups
Get-ADGroupMember	Displays membership of groups
Remove-ADGroupMember	Removes members from groups
Add-ADPrincipalGroupMembership	Adds group membership to objects
Get-ADPrincipalGroupMembership	Displays group membership of objects
Remove-ADPrincipalGroupMembership	Removes group membership from an object


```
New-ADGroup -Name "CustomerManagement" -Path
"ou=managers,dc=adatum,dc=com" -GroupScope Global
-GroupCategory Security

Add-ADGroupMember -Name "CustomerManagement"
-Members "Joe"
```

Cmdlet	Description
New-ADGroup	Creates new groups.
Set-ADGroup	Modifies properties of groups.
Get-ADGroup	Displays properties of groups.
Remove-ADGroup	Deletes groups.
Add-ADGroupMember	Adds members to groups.
Get-ADGroupMember	Displays membership of groups.
Remove-ADGroupMember	Removes members from groups.
Add-ADPrincipalGroupMembership	Adds group membership to objects.
Get-ADPrincipalGroupMembership	Displays group membership of objects.
Remove-ADPrincipalGroupMembership	Removes group membership from an object.

Create New Groups

You can use the **New-ADGroup** cmdlet to create groups. However, when you create groups using the **New-ADGroup** cmdlet, you must use the **GroupScope** parameter in addition to the group name. This is the only required parameter. The following table lists commonly used parameters for **New-ADGroup**.

Parameter	Description
Name	Defines the name of the group.
GroupScope	Defines the scope of the group as DomainLocal , Global , or Universal . You must provide this parameter.
DisplayName	Defines the LDAP display name for the object.
GroupCategory	Defines whether it is a security group or a distribution group. If you do not specify either, a security group is created.
ManagedBy	Defines a user or group that can manage the group.
Path	Defines the OU or container in which the group is created.
SamAccountName	Defines a name that is backward compatible with older operating systems.

The following command is an example of what you could type at a Windows PowerShell prompt to create a new group:

```
New-ADGroup -Name "CustomerManagement" -Path "ou=managers,dc=adatum,dc=com" -GroupScope Global -GroupCategory Security
```

Manage Group Membership

There are two sets of cmdlets that you can use to manage group membership: ***-ADGroupMember**, and ***-ADPrincipalGroupMembership**. The distinction between these two sets of cmdlets is the perspective used when modifying group membership. They are:

- The ***-ADGroupMember** cmdlets modify the membership of a group. For example, you add or remove members of a group.
 - You cannot pipe a list of members to these cmdlets.
 - You can pass a list of groups to these cmdlets.
- The ***-ADPrincipalGroupMembership** cmdlets modify the group membership of an object such as a user. For example, you can modify a user account to add it as a member of a group.
 - You can pipe a list of members to these cmdlets.
 - You cannot provide a list of groups to these cmdlets.



Note: Piping is a common process in scripting languages that allow you to use the output of one cmdlet as input for the next cmdlet in the command. For example, the command below creates a user account, and then enables the account:

```
New-ADUser -Name "Sten Faerch" -AccountPassword (Read-Host -AsSecureString "Enter password") | Enable-Account
```

The following is a command you could use to add a member to a group:

```
Add-ADGroupMember -Name "CustomerManagement" -Members "Joe"
```

Using Windows PowerShell Cmdlets to Manage Computer Accounts

You can use Windows PowerShell to create, modify, and delete computer accounts. You can use these cmdlets for individual operations or as part of a script to perform bulk operations. Some of the cmdlets for managing computer accounts are listed in the following table.

Cmdlet	Description
New-ADComputer	Creates new computer accounts
Set-ADComputer	Modifies properties of computer accounts
Get-ADComputer	Displays properties of computer accounts
Remove-ADComputer	Deletes computer accounts
Test-ComputerSecureChannel	Verifies or repairs the trust relationship between a computer and the domain
Reset-ComputerMachinePassword	Resets the password for a computer account

```
New-ADComputer -Name "LON-SVR8" -Path "ou=marketing,dc=adatum,dc=com" -Enabled $true
```

```
Test-ComputerSecureChannel -Repair
```

Cmdlet	Description
New-ADComputer	Creates a new computer account.
Set-ADComputer	Modifies properties of a computer account.
Get-ADComputer	Displays properties of a computer account.
Remove-ADComputer	Deletes a computer account.
Test-ComputerSecureChannel	Verifies or repairs the trust relationship between a computer and the domain.
Reset-ComputerMachinePassword	Resets the password for a computer account.

Create New Computer Accounts

You can use the **New-ADComputer** cmdlet to create a new computer account before the computer is joined to the domain. You do this so you can create the computer account in the correct OU before deploying the computer.

The following table lists commonly used parameters for **New-ADComputer**.

Parameter	Description
Name	Defines the name of the computer account.
Path	Defines the OU or container where the computer account will be created.
Enabled	Defines whether the computer account is enabled or disabled. By default, the computer account is enabled and a random password is generated.

The following is an example that you can use to create a computer account:

```
New-ADComputer -Name LON-SVR8 -Path "ou=marketing,dc=adatum,dc=com" -Enabled $true
```

Repair the Trust Relationship for a Computer Account

You can use the **Test-ComputerSecureChannel** cmdlet with the **-Repair** parameter to repair a lost trust relationship between a computer and the domain. You must run the cmdlet on the computer with the lost trust relationship.

The following is a command that you could use to repair the trust relationship for a computer account:

```
Test-ComputerSecureChannel -Repair
```

Using Windows PowerShell Cmdlets to Manage OUs

You can use Windows PowerShell cmdlets to create, modify, and delete OUs. You can use these cmdlets for individual operations or as part of a script to perform bulk operations. Some of the cmdlets for managing OUs are listed in the following table.

Cmdlet	Description
New-ADOrganizationalUnit	Creates organizational units (OUs)
Set-ADOrganizationalUnit	Modifies properties of OUs
Get-ADOrganizationalUnit	Views properties of OUs
Remove-ADOrganizationalUnit	Deletes OUs
New-ADOrganizationalUnit	Creates OUs
Set-ADOrganizationalUnit	Modifies properties of OUs
Get-ADOrganizationalUnit	Views properties of OUs

```
New-ADOrganizationalUnit -Name "Sales"
-Path "ou=marketing,dc=adatum,dc=com"
-ProtectedFromAccidentalDeletion $true
```

Cmdlet	Description
New-ADOrganizationalUnit	Creates OUs.
Set-ADOrganizationalUnit	Modifies properties of OUs.
Get-ADOrganizationalUnit	Displays properties of OUs.
Remove-ADOrganizationalUnit	Deletes OUs.

Create New OUs

You can use **New-ADOrganizationalUnit** cmdlet to create a new OU to represent departments or physical locations within in your organization.

The following table shows commonly used parameters for the **New-ADOrganizationalUnit** cmdlet.

Parameter	Description
Name	Defines the name of the new OU.
Path	Defines the location of the new OU.
ProtectedFromAccidentalDeletion	Prevents the OU from being deleted accidentally. The default value is \$true .

The following is an example you can use when you want to create a new OU:

```
New-ADOrganizationalUnit -Name Sales -Path "ou=marketing,dc=adatum,dc=com"
-ProtectedFromAccidentalDeletion $true
```

Question: In the slide example, is the **ProtectedFromAccidentalDeletion** parameter required?

Lesson 3

Performing Bulk Operations with Windows PowerShell

Windows PowerShell is a powerful scripting environment that you can use to perform bulk operations, which are normally tedious to perform manually. You can also perform some bulk operations in graphical tools.

To perform bulk operations using Windows PowerShell, you must first understand how to create queries for a list of AD DS objects, and how to work with .csv files. Then you can create scripts that perform the bulk operations that you require.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe bulk operations.
- Use graphical tools to perform bulk operations.
- Query AD DS objects by using Windows PowerShell.
- Modify AD DS objects by using Windows PowerShell.
- Use .csv files with Windows PowerShell.
- Modify and execute Windows PowerShell scripts to perform bulk operations.

What Are Bulk Operations?

A *bulk operation* is a single action that changes multiple objects. Performing a bulk operation is much faster than changing many objects individually. It might also be more accurate, because performing many individual actions increases the likelihood of making a typographical mistake. However, due to the nature of Bulk Operations, if a mistake is introduced, it may multiply itself to every single object being changed. Therefore, ensure you test your bulk operation on a smaller object set before executing it on all the elements you want to modify.

- A bulk operation is a single action that changes multiple objects
- Sample bulk operations
 - Create user accounts based on data in a spreadsheet
 - Disable all accounts not used in 6 months
 - Rename the department for many users
- You can perform bulk operations by using:
 - Graphical tools
 - Command-line tools
 - Script

The operations below are examples of common bulk operations in a Windows Server environment:

- Create new user accounts based on information from a spreadsheet.
- Disable all user accounts that have not been used in the past 6 months.
- Change the department name for all user belonging to a given department.

You can perform bulk operations with graphical tools, at a command prompt, or by using scripts. Each method for performing bulk operations has different capabilities.

For example:

- Graphical tools tend to be limited in the properties that they can modify.
- Command-line tools tend to be more flexible than graphical tools when defining queries, and they have more options for modifying object properties.
- Scripts can combine multiple command-line actions for the most complexity and flexibility.

Demonstration: Using Graphical Tools to Perform Bulk Operations

You can use the graphical tools Active Directory Administrative Center and Active Directory Users and Computers to modify the properties of multiple objects simultaneously.



Note: When you use graphical tools to modify multiple user accounts simultaneously, you are limited to modifying the properties that display in the user interface.

To perform a bulk operation using graphical tools, perform the following steps:

1. Perform a search or create a filter to display the objects that you want to modify.
2. Select the objects.
3. Examine the properties of the objects.
4. Modify the properties that you want to change.

In this demonstration, you will see how to:

- Create a query for all users.
- Configure the Company attribute for all users.
- Verify that the Company attribute has been modified.

Demonstration Steps

Create a query for all users

1. Sign in to LON-DC1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On LON-DC1, open Active Directory Administrative Center.
3. Go to Global Search, and add the criteria **Object type is user/inetOrgPerson /computer/group/organization unit**.
4. Verify that the criteria that you added is for the type **User**, and perform the search.

Configure the Company attribute for all users

1. Select all the user accounts and modify their properties.
2. Make the Company name **A. Datum**.

Verify that the Company attribute has been modified

- Open the properties of Adam Barr, and verify that the Company is **A. Datum**.

Querying Objects with Windows PowerShell

In Windows PowerShell, you use the **Get-*** cmdlets to obtain lists of objects, such as user accounts. You can also use these cmdlets to generate queries for objects on which you can perform bulk operations. The following table lists commonly used parameters with the **Get-AD*** cmdlets.

Parameter	Description
SearchBase	Defines the AD DS path to begin searching.
SearchScope	Defines at what level below the SearchBase a search should be performed.
ResultSetSize	Defines how many objects to return in response to a query.
Properties	Defines which object properties to return and display.
Filter	Defines a filter by using PowerShell syntax
LDAPFilter	Defines a filter by using LDAP query syntax

Descriptions of operators

-eq	Equal to	-gt	Greater than
-ne	Not equal to	-ge	Greater than or equal to
-lt	Less than	-like	Uses wildcards for pattern matching
-le	Less than or equal to		

Parameter	Description
SearchBase	Defines the AD DS path to begin searching, for example, the domain or an OU.
SearchScope	Defines at what level below the SearchBase that the search should be performed. You can choose to search only in the base, one level down, or the entire subtree.
ResultSetSize	Defines how many objects to return in response to a query. To ensure that all objects are returned, set this to \$null .
Properties	Defines which object properties to return and display. To return all properties, type an asterisk (*). You do not need to use this parameter to use a property for filtering.

Create a Query

You can use the **Filter** parameter or the **LDAPFilter** parameter to create queries for objects with the **Get-AD*** cmdlets. The **Filter** parameter is used for queries written in Windows PowerShell. The **LDAPFilter** parameter is used for queries written as LDAP query strings.

Windows PowerShell is preferred because:

- It is easier to write queries in Windows PowerShell.
- You can use variables inside the queries.
- There is automatic conversion of variable types, when it is required.

The following table lists commonly used operators that you can use in Windows PowerShell.

Operator	Description
-eq	Equal to
-ne	Not equal to
-lt	Less than
-le	Less than or equal to
-gt	Greater than
-ge	Greater than or equal to
-like	Uses wildcards for pattern matching

Filter

As previously mentioned, you can use the **Filter** parameter to filter data retrieved by a **Get-*** cmdlet. The **Filter** parameter uses the same syntax as the **Where-Object** cmdlet in Windows PowerShell. As an example, the command below retrieves all user accounts that have Smith as their last name, followed by the output of the command:

```
> Get-ADUser -Filter {sn -eq "Smith"}

DistinguishedName : CN=Denise Smith,OU=Marketing,DC=Adatum,DC=com
Enabled           : True
GivenName        : Denise
Name             : Denise Smith
ObjectClass      : user
ObjectGUID       : 1ff1b2cb-38c1-4bc7-bda7-511e19744d2a
SamAccountName   : Denise
SID              : S-1-5-21-322346712-1256085132-1900709958-1407
Surname          : Smith
UserPrincipalName : Denise@adatum.com

DistinguishedName : CN=Tony Smith,OU=IT,DC=Adatum,DC=com
Enabled           : True
GivenName        : Tony
Name             : Tony Smith
ObjectClass      : user
ObjectGUID       : ac7eb8db-3cf1-4e6d-91d3-7527e540c284
SamAccountName   : Tony
SID              : S-1-5-21-322346712-1256085132-1900709958-1408
Surname          : Smith
UserPrincipalName : Tony@adatum.com
```

One of the characteristics of all **Get-*** cmdlets that are used to retrieve data from Active Directory is that they do not always return all properties for the objects they retrieve. For instance, looking at the output above you only see some of the properties that a user account has in Active Directory. You do not see, for instance, the **mail** property. You can use the **-Properties** parameters to retrieve properties not returned by default when running **Get-*** cmdlets.

For example, the code below returns the same list of users, but this time with the mail and **PasswordLastSet** properties:

```
> Get-ADUser -Filter {sn -eq "Smith"} -Properties mail,passwordlastset

DistinguishedName : CN=Denise Smith,OU=Marketing,DC=Adatum,DC=com
Enabled           : True
GivenName        : Denise
Name             : Denise Smith
ObjectClass      : user
ObjectGUID       : 1ff1b2cb-38c1-4bc7-bda7-511e19744d2a
PasswordLastSet  : 7/9/2013 12:51:30 PM
SamAccountName   : Denise
SID              : S-1-5-21-322346712-1256085132-1900709958-1407
Surname          : Smith
UserPrincipalName : Denise@adatum.com

DistinguishedName : CN=Tony Smith,OU=IT,DC=Adatum,DC=com
Enabled           : True
GivenName        : Tony
Name             : Tony Smith
ObjectClass      : user
ObjectGUID       : ac7eb8db-3cf1-4e6d-91d3-7527e540c284
PasswordLastSet  : 7/9/2013 12:51:30 PM
SamAccountName   : Tony
SID              : S-1-5-21-322346712-1256085132-1900709958-1408
Surname          : Smith
UserPrincipalName : Tony@adatum.com
```

The following is a command that you use to display all of the properties for a user account:

```
Get-ADUser -Name "Administrator" -Properties *
```

The following is a command that you use to return all the user accounts in the Marketing OU, and all of its child OUs:


```
Get-ADUser -Filter * -SearchBase "ou=Marketing,dc=adatum,dc=com" -SearchScope subtree
```

The following is a command that you use to show all of the user accounts with a last logon date older than a specific date:

```
Get-ADUser -Filter {lastlogondate -lt "January 1, 2012"}
```

The following is a command that you use to show all of the user accounts in the Marketing department that have a last logon date older than a specific date:

```
Get-ADUser -Filter {(lastlogondate -lt "January 1, 2012")} -and (department -eq "Marketing")}
```

 **Additional Reading:** For more information about filtering with **Get-AD*** cmdlets, see the article about `_ActiveDirectory_Filter` at <http://go.microsoft.com/fwlink/?LinkID=266740>.

LDAPFilter

You can also use the **LDAPFilter** parameter to filter data retrieved by a **Get-*** cmdlet. The **LDAPFilter** parameter takes a string value that uses the same syntax you use to build an LDAP query. For example, the command below retrieves all users whose last name is Smith:

```
> Get-ADUser -LDAPFilter "(sn=Smith)"
```

Search-ADAccount

One of the downsides of the **Get-AD*** cmdlets is how they deal with the **userAccountControl** property. This property is a 4-byte bitmap, where each bit corresponds to a different property linked to an Active Directory account. The table below shows some of the bits in the bit map:

Hexadecimal value	Decimal value	Identifier	Description
0x00000001	1	ADS_UF_SCRIPT	Logon script is executed
0x00000002	2	ADS_UF_ACCOUNTDISABLE	User account is disabled
0x00000008	8	ADS_UF_HOMEDIR_REQUIRED	A home folder is required
0x00000010	16	ADS_UF_LOCKOUT	User account is locked out



Additional Reading: For the full list of flags in the **UserAccountControl** property, visit <http://go.microsoft.com/fwlink/?LinkID=331075>.

When you read the **userAccountControl**, you receive a numerical value. Not a group of true/false values per individual flag. For example, the code below shows the **userAccountControl** property for all users whose last name begin with S.

```
> Get-ADUser -Filter {sn -like "Sm*"} -Properties userAccountControl | Select
Name,userAccountControl | FT -AutoSize
```

Name	userAccountControl
-----	-----
Denise Smith	66048
Tony Smith	66048

Imagine that you need to retrieve a list of all disabled accounts in Active Directory. To do that, you need to retrieve all accounts where the **userAccountControl** property has the second to last bit enabled. You can do this by using the code below:

```
> Get-ADUser -LDAPFilter "(userAccountControl:1.2.840.113556.1.4.803:=2)" | Select Name
```

Memorizing, or even looking up, flags in the **userAccountControl** property is a time consuming job. Of course, you can create scripts that already have the code built-in and just reuse them. But it would be much better to have a command that abstracts all that work for you—and that is exactly what the **Search-ADAccount** cmdlet does. Here is how you retrieve a list of disabled accounts by using the **Search-ADAccount** cmdlet:

```
> Search-ADAccount -AccountDisabled | Select name
```

That is much easier than using **Get-ADUser**. The **Search-ADAccount** cmdlet uses the following parameters:

Parameter	Description
AccountDisabled	Retrieves a list of disabled accounts.
AuthType	Specifies the authentication type used when running this command.
AccountExpired	Retrieves a list of expired accounts that have.
AccountExpiring	Retrieves a list of accounts that expire within a given time span.
AccountInactive	Retrieves a list of accounts that will become inactive within a given time span.
LockedOut	Retrieves a list of accounts that are locked out.
PasswordExpired	Retrieves a list of accounts that have their password expired.
PasswordExpiring	Retrieves a list of accounts that have their password about to expire within a time span.
PasswordNeverExpires	Retrieves a list of accounts whose password never expires.
ComputersOnly	Retrieves computer accounts.
UsersOnly	Retrieves user accounts.
TimeSpan	Used in conjunction with PasswordExpiring, AccountExpiring and AccountInactive to specify the time span for those parameters.
DateTime	Used in conjunction with PasswordExpiring, AccountExpiring and AccountInactive to specify the expiry date for those parameters.
SearchBase	Specifies the base LDAP container for the search.
SearchScope	Specifies the scope for the search.
Server	Specifies the server to connect to.

Here are some examples of the **Search-ADAccount** cmdlet:

```
# Retrieve all disabled user accounts
Search-ADAccount -AccountDisabled -UsersOnly
```

```
# Retrieve all user accounts inactive for the last 5 days
Search-ADAccount -AccountInactive -TimeSpan -5 -UsersOnly
```

```
# Retrieve all user accounts whose password will expire on 7/4/2014
Search-ADAccount -AccountExpiring -DateTime "4/7/2014" -UsersOnly
```

```
# Retrieve all computer accounts that are locked out
Search-ADAccount -ComputersOnly -LockedOut
```

Question: What is the difference between using **-eq** and **-like** when comparing strings?

Modifying Objects with Windows PowerShell

To perform a bulk operation, you need to pass the list of objects that you have queried to another cmdlet to modify the objects. In most cases, you use the **Set-AD*** cmdlets to modify the objects.

To pass the list of queried objects to another cmdlet for further processing, you use the pipe (|) character. The pipe character passes each object from the query to a second cmdlet, which then performs a specified operation on each object.

The following is a command that you could use for those accounts that do not have the Company attribute set. It generates a list of user accounts and sets the Company attribute to **A. Datum**.

```
Get-ADUser -Filter {company -notlike "*"} | Set-ADUser -Company "A. Datum"
```



Additional Reading: For more information on the Set-ADUser cmdlet, visit <http://go.microsoft.com/fwlink/?LinkID=331074>.

The following is a command that you could use to generate a list of user accounts that have not logged on since a specific date, and then disable them:

```
Get-ADUser -Filter {lastlogondate -lt "January 1, 2012"} | Disable-ADAccount
```

Use Objects from a Text File

Instead of using a list of objects from a query to perform a bulk operation, you can use a list of objects in a text file. This is useful when you have a list of objects to modify or remove, and it is not possible to generate that list by using a query. For example, the human resources department might generate a list of user accounts to be disabled. There is no query that can identify a list of users that have left the organization.

When you use a text file to specify a list of objects, the text file needs to have the name of each object on a single line.

The following is a command that you could use to disable the user accounts that are listed in a text file:

```
Get-Content C:\users.txt | Disable-ADAccount
```

Question: Which attributes of a user account can you use when creating a query by using the **Filter** parameter?

Use the pipe character (|) to pass a list of objects to a cmdlet for further processing

```
Get-ADUser -Filter {company -notlike "*"} |  
Set-ADUser -Company "A. Datum"
```

```
Get-ADUser -Filter {lastlogondate -lt "January 1,  
2012"} | Disable-ADAccount
```

```
Get-Content C:\users.txt | Disable-ADAccount
```

Working with CSV Files

A .csv file can contain much more information than a simple list. Similar to a spreadsheet, a .csv file can have multiple rows and columns of information. Each row in the .csv file represents a single object, and each column in the .csv file represents a property of the object. This is useful for bulk operations such as creating user accounts where multiple pieces of information about each object are required.

You can use the **Import-CSV** cmdlet to read the contents of a .csv file into a variable, and then work with the data. After the data is imported into the variable, you can then refer to each individual row of data and each individual column of data. Each column of data has a name that is based on the header row (the first row) of the .csv file. You can refer to each column by name.

The following is an example a .csv file with a header row:

```
FirstName,LastName,Department
Greg,Guzik,IT
Robin,Young,Research
Qiong,Wu,Marketing
```

The first line of a .csv file defines the names of the columns

```
FirstName,LastName,Department
Greg,Guzik,IT
Robin,Young,Research
Qiong,Wu,Marketing
```

A foreach loop processes the contents of a .csv that have been imported into a variable

```
$users=Import-CSV -LiteralPath "C:\users.csv"
foreach ($user in $users) {
    Write-Host "The first name is:"
    $user.FirstName
}
```

Use Foreach to Process CSV Data

In many cases, you will be creating script that will be reused for multiple .csv files, and you will not know how many rows there are in each .csv file. In these cases, you can use a **foreach** loop to process each row in a .csv file. This type of loop does not require that you know how many rows there are. During each iteration of the **foreach** loop, a row from the .csv is imported into a variable that is then processed.

The following is a command that you could use to import a .csv file into a variable, and use a **foreach** loop to display the first name from each row in a .csv file:

```
$users=Import-CSV -LiteralPath "C:\users.csv"
foreach ($user in $users)
{
    Write-Host "The first name is:" $user.FirstName
}
```

Question: In the **foreach** loop, how does **\$i** change?

Demonstration: Performing Bulk Operations with Windows PowerShell

You can use a script to combine multiple Windows PowerShell commands to perform more complex tasks. Within a script, you often use variables and loops to process data. Windows PowerShell scripts have a .ps1 extension.

The execution policy on a server determines whether scripts are able to run. The default execution policy on Windows Server 2012 is **RemoteSigned**. This means that local scripts can run without being digitally signed. You can control the execution policy by using the **Set-ExecutionPolicy** cmdlet.

In this demonstration, you will see how to:

- Configure a department for users.
- Create an OU.
- Run a script to create new user accounts.
- Verify that new user accounts were created.

Demonstration Steps

Configure a department for users

1. On LON-DC1, open a **Windows PowerShell** Command Prompt window.
2. At the Windows PowerShell prompt, search for user accounts in the Research OU using the following command:

```
Get-ADUser -Filter * -SearchBase "ou=Research,dc=adatum,dc=com"
```

3. Set the **department** attribute of all users in the Research OU using the following command:

```
Get-ADUser -Filter * -SearchBase "ou=Research,dc=adatum,dc=com" | Set-ADUser -Department Research
```

4. Display a table-formatted list of users in the Research department. Display the **distinguished name** and **department** by using the following command:

```
Get-ADUser -Filter 'department -eq "Research"' | Format-Table DistinguishedName,Department
```

5. Use the **Properties** parameter in the previous command so that the department is displayed correctly in the output. Use the following command:

```
Get-ADUser -Filter 'department -eq "Research"' -Properties Department | Format-Table DistinguishedName,Department
```

Create an OU

- At the Windows PowerShell prompt, create a new OU named **LondonBranch** using the following command:

```
New-ADOrganizationalUnit LondonBranch -Path "dc=adatum,dc=com"
```

Run a script to create new user accounts

1. Open **E:\Labfiles\Mod04\DemoUsers.csv**, and read the header row.
2. Edit **DemoUsers.ps1**, and review the contents of the script. Note that the script:
 - Refers to the location of the .csv file.
 - Uses a **foreach** loop to process the .csv file contents.
 - Refers to the columns defined by the header in the .csv file.
3. At the Windows PowerShell prompt, change to the E:\Labfiles\Mod04 directory, and then run the following command:

```
.\DemoUsers.ps1
```

Verify that new user accounts were created

1. In Server Manager, open **Active Directory Administrative Center**.
2. In Active Directory Administrative Center, go to **Adatum (local)>LondonBranch**, and verify that the user accounts were created.

Note that the passwords are disabled because no password was set during creation.

Lab: Automating AD DS Administration by Using Windows PowerShell

Scenario

You have been working for A. Datum for several years as a desktop support specialist. In this role, you visited desktop computers to troubleshoot app and network problems. You have recently accepted a promotion to the server support team. One of your first assignments is configuring the infrastructure service for a new branch office.

As part of configuring a new branch office, you need to create user and group accounts. Creating multiple users with graphical tools is inefficient, so, you will use Windows PowerShell.

Objectives

After completing this lab, you will be able to:

- Create user accounts and groups by using Windows PowerShell.
- Use Windows PowerShell to create user accounts in bulk.
- Use Windows PowerShell to modify user accounts in bulk.

Lab Setup

Estimated time: 45 minutes

Virtual machines	20410C-LON-DC1 20410C-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 **20410C-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in using the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
5. Repeat steps 2 and 3 for 20410C-LON-CL1. Do not sign in to LON-CL1 until directed to do so.

Exercise 1: Creating User Accounts and Groups by Using Windows PowerShell

Scenario

A. Datum Corporation has a number of scripts that have been used in the past to create user accounts by using command-line tools. It has been mandated that all future scripting will be done by using Windows PowerShell. As the first step in creating scripts, you need to identify the syntax required to manage AD DS objects in Windows PowerShell.

The main tasks for this exercise are as follows:

1. Create a user account by using Windows PowerShell.
2. Create a group by using Windows PowerShell.

► Task 1: Create a user account by using Windows PowerShell

1. On LON-DC1, open a **Windows PowerShell** Command Prompt window.
2. At the Windows PowerShell prompt, create a new OU named **LondonBranch** by typing the following command:

```
New-ADOrganizationalUnit LondonBranch
```

3. Create a new user account for **Ty Carlson** in the LondonBranch OU using the following command:

```
New-ADUser -Name Ty -DisplayName "Ty Carlson" -GivenName Ty -Surname Carlson -Path "ou=LondonBranch,dc=adatum,dc=com"
```

4. Change the blank password for the new account to **Pa\$\$w0rd**, using the following command:

```
Set-ADAccountPassword Ty
```

5. Enable the new user account using the following command:

```
Enable-ADAccount Ty
```

6. On LON-CL1, sign in as **Ty** with the password **Pa\$\$w0rd**.
7. Verify that the sign-in is successful, and then sign out of LON-CL1.

► Task 2: Create a group by using Windows PowerShell

1. On LON-DC1, at the Windows PowerShell prompt, create a new global security group for users in the London branch office, using the following command:

```
New-ADGroup LondonBranchUsers -Path "ou=LondonBranch,dc=adatum,dc=com" -GroupScope Global -GroupCategory Security
```

2. At the Windows PowerShell prompt, add **Ty** as a member of LondonBranchUsers, using the following command:

```
Add-ADGroupMember LondonBranchUsers -Members Ty
```

- At the Windows PowerShell prompt, confirm that Ty has been added as a member of LondonBranchUsers, using the following command:

```
Get-ADGroupMember LondonBranchUsers
```

Results: After completing this exercise, you will have created user accounts and groups by using Windows PowerShell.

Exercise 2: Using Windows PowerShell to Create User Accounts in Bulk

Scenario

You have been given a .csv file that contains a large list of new users for the branch office. It would be inefficient to create these users individually with graphical tools. Instead, you will use a Windows PowerShell script to create the users. A colleague that is experienced with scripting has provided you with a script that she created. You need to modify the script to match the format of your .csv file.

The main tasks for this exercise are as follows:

- Prepare the .csv file.
- Prepare the script.
- Run the script.

► Task 1: Prepare the .csv file

- On LON-DC1, read the contents in **E:\Labfiles\Mod04\LabUsers.ps1** to identify the header requirements for the .csv file.
- Edit the contents in **C:\Labfiles\Mod04\LabUsers.csv**, and add the appropriate header.

► Task 2: Prepare the script

- On LON-DC1, use Windows PowerShell Integrated Scripting Environment (ISE) to modify the variables in **LabUsers.ps1**:
 - \$csvfile: **E:\Labfiles\Mod04\labUsers.csv**
 - \$OU: **"ou=LondonBranch,dc=adatum,dc=com"**
- Save the modified **LabUsers.ps1**.
- Review the contents of the script.

► Task 3: Run the script

- On LON-DC1, open a **Windows PowerShell** Command Prompt window, and run **E:\Labfiles\Mod04\LabUsers.ps1**.
- At the Windows PowerShell prompt use the following command to verify that the users were created:

```
Get-ADUser -Filter * -SearchBase "ou=LondonBranch,dc=adatum,dc=com"
```

- On LON-CL1, sign in as **Luka** with the password **Pa\$\$w0rd**.

Results: After completing this exercise, you will have used Windows PowerShell to create user accounts in bulk.

Exercise 3: Using Windows PowerShell to Modify User Accounts in Bulk

Scenario

You have received a request to update all user accounts in the new branch office OU with the correct address of the new building. You have also been asked to ensure that all of the new user accounts in the branch office are configured to force users to change their passwords the next time they sign in.

The main tasks for this exercise are as follows:

1. Force all user accounts in LondonBranch to change their passwords at next sign in.
2. Configure the address for user accounts in LondonBranch.

► Task 1: Force all user accounts in LondonBranch to change their passwords at next sign-in

1. On LON-DC1, open a **Windows PowerShell** Command Prompt window.
2. At the Windows PowerShell prompt, create a query for user accounts in the LondonBranch OU using the following command:

```
Get-ADUser -Filter * -SearchBase "ou=LondonBranch,dc=adatum,dc=com" | Format-Wide DistinguishedName
```

3. At the Windows PowerShell prompt, modify the previous command to force all users to change their password the next time they sign in by using the following command:

```
Get-ADUser -Filter * -SearchBase "ou=LondonBranch,dc=adatum,dc=com" | Set-ADUser -ChangePasswordAtLogon $true
```

► Task 2: Configure the address for user accounts in LondonBranch

1. On LON-DC1, open **Active Directory Administrative Center**.
2. Open the properties for all user accounts in LondonBranch.
3. Set the address for multiple users as follows:
 - Street: **Branch Office**
 - City: **London**
 - Country/Region: **United Kingdom**

Results: After completing this exercise, you will have modified user accounts in bulk.

► Prepare for the next module

When you finish the lab, revert all virtual machines back to their initial state by performing the following steps:

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20410C-LON-CL1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20410C-LON-DC1**.

Lab Review Questions

Question: By default, are new user accounts enabled or disabled when you create them by using the **New-ADUser** cmdlet?

Question: What file extension do Windows PowerShell scripts use?

Module Review and Takeaways

Review Questions

- Question:** A colleague is creating a Windows PowerShell script that creates user accounts from data in a .csv file. However, his script is experiencing errors when attempting to set a default password. Why might this be happening?
- Question:** You are an administrator for a school district that creates 20,000 new user accounts for students each year. The administration system for students can generate a list of the new students and then export it as a .csv file. After the data has been exported to a .csv file, what information do you need to work with the data in a script?
- Question:** The Research department in your organization has been renamed "Research and Development." You need to update the Department property of users in the Research department to reflect this change.
- You have created a query for user accounts with the **department** property set to **Research**, by using the **Get-ADUser** cmdlet and the **-Filter** parameter. What is the next step to update the department property to Research and Development?

Tools

Tool	Use	Where to find it
Csvde	Csvde is a command-line tool that exports or imports AD DS objects to or from a comma-separated values (.csv) file.	In Windows Server 2012.
Ldifde	Ldifde is a command-line tool that you can use to export, create, modify, or delete AD DS objects. Like csvde, ldifde uses data that is stored in a file.	In Windows Server 2012.
DS command-line tools	You can use DS command-line tools to create, view, modify, and remove AD DS objects. These tools are suitable for scripts and include: DSadd, DSget, DSquery, DSmod, DSrm and DSmove.	In Windows Server 2012

Module 5

Implementing IPv4

Contents:

Module Overview	5-1
Lesson 1: Overview of TCP/IP	5-2
Lesson 2: Understanding IPv4 Addressing	5-6
Lesson 3: Subnetting and Supernetting	5-11
Lesson 4: Configuring and Troubleshooting IPv4	5-16
Lab: Implementing IPv4	5-26
Module Review and Takeaways	5-31

Module Overview

IPv4 is the network protocol used on the Internet and local area networks. To ensure that you can understand and troubleshoot network communication, it is essential that you understand how IPv4 is implemented. In this module, you will see how to implement an IPv4 addressing scheme, and determine and troubleshoot network-related problems.

Objectives

After completing this module, you will be able to:

- Describe the TCP/IP protocol suite.
- Describe IPv4 addressing.
- Determine a subnet mask necessary for subnetting or supernetting.
- Configure IPv4 and troubleshoot IPv4 communication.

Lesson 1

Overview of TCP/IP

TCP/IP is an industry standard suite of protocols that provides communication in a heterogeneous network. This lesson provides an overview of IPv4 and how it relates to other protocols to enable network communication. It also covers the concept of sockets, which applications use to accept network communications. Combined together, this lesson provides a foundation for understanding and troubleshooting network communication.

Lesson Objectives

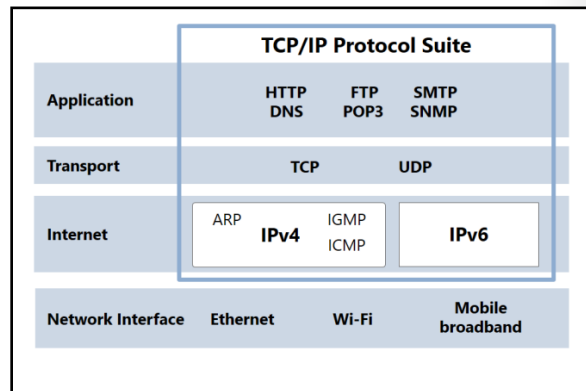
After completing this lesson, you will be able to:

- Describe the elements of the TCP/IP suite of protocols.
- Describe the individual protocols that make up the TCP/IP suite.
- Describe TCP/IP application layer protocols.
- Describe a socket, and identify port numbers for specified protocols.

The TCP/IP Protocol Suite

The tasks performed by TCP/IP in the communication process are distributed between protocols. These protocols are organized into four distinct layers within the TCP/IP stack:

- Application layer. Applications use the application layer protocols to access network resources.
- Transport layer. The transport layer protocols control data transfer reliability on the network.
- Internet layer. The internet layer protocols control packet movement between networks.
- Network interface layer. The network interface layer protocols define how datagrams from the Internet layer are transmitted on the media.



Benefits of Architecture Layers

Rather than creating a single protocol, dividing the network functions into a stack of separate protocols provides several benefits:

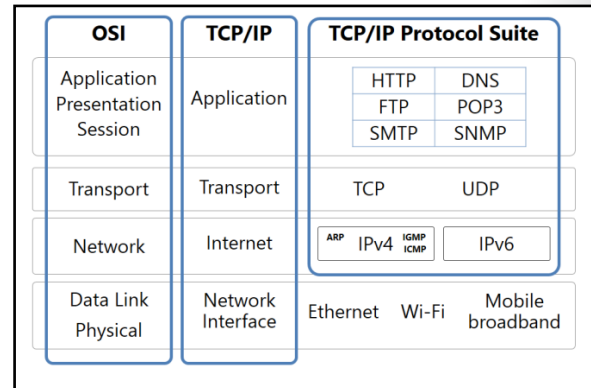
- Separate protocols make it easier to support a variety of computing platforms.
- Creating or modifying protocols to support new standards does not require modification of the entire protocol stack.
- Having multiple protocols operating at the same layer makes it possible for applications to select the protocols that provide only the level of service required.
- Because the stack is split into layers, the development of the protocols can proceed simultaneously by personnel who are uniquely qualified in the operations of the particular layers.

Protocols in the TCP/IP Suite

The Open Systems Interconnection (OSI) model defines distinct layers related to packaging, and sending and receiving data transmissions over a network. The layered suite of protocols that form the TCP/IP stack carry out these functions.

Application Layer

The application layer of the TCP/IP model corresponds to the application, presentation, and session layers of the OSI model. This layer provides services and utilities that enable applications to access network resources.



Transport Layer

The transport layer corresponds to the transport layer of the OSI model and is responsible for end-to-end communication using TCP or User Datagram Protocol (UDP). The TCP/IP protocol suite offers application programmers the choice of TCP or UDP as a transport layer protocol:

- **TCP.** Provides connection-oriented reliable communications for applications. Connection-oriented communication confirms that the destination is ready to receive data before it sends the data. To make communication reliable, TCP confirms that all packets are received. Reliable communication is desired in most cases, and is used by most applications. Web servers, File Transfer Protocol (FTP) clients, and other applications that move large amounts of data use TCP.
- **UDP.** Provides connectionless and unreliable communication. When using UDP, reliable delivery is the responsibility of the application. Applications use UDP for faster communication with less overhead than TCP. Applications such as streaming audio and video use UDP so that a single missing packet will not delay playback. UDP is also used by applications that send small amounts of data, such as Domain Name System (DNS) name lookups.

The transport layer protocol that an application uses is determined by the developer of an application, and is based on the communication requirements of the application.

Internet Layer

The Internet layer corresponds to the network layer of the OSI model and consists of several separate protocols, including: IP; Address Resolution Protocol (ARP); Internet Group Management Protocol (IGMP); and Internet Control Message Protocol (ICMP). The protocols at the Internet layer encapsulate transport layer data into units called *packets*, address them, and then route them to their destinations.

The Internet layer protocols are:

- **IP.** IP is responsible for routing and addressing. The Windows® 8 operating system and the Windows Server® 2012 operating system implement a dual-layer IP protocol stack, including support for both IPv4 and IPv6.
- **ARP.** ARP is used by IP to determine the media access control (MAC) address of local network adapters—that is, adapters installed on computers on the local network—from the IP address of a local host. ARP is broadcast-based, meaning that ARP frames cannot transit a router and are therefore localized. Some implementations of TCP/IP provide support for Reverse ARP (RARP) in which the MAC address of a network adapter is used to determine the corresponding IP address.
- **IGMP.** IGMP provides support for multitasking applications over routers in IPv4 networks.
- **ICMP.** ICMP sends error messages in an IP-based network.

Network Interface Layer

The *network interface layer* (sometimes referred to as the *link layer* or *data link layer*) corresponds to the data link and physical layers of the OSI model. The network interface layer specifies the requirements for sending and receiving packets on the network media. This layer is often not formally considered part of the TCP/IP protocol suite because the tasks are performed by the combination of the network adapter driver and the network adapter.

TCP/IP Applications

Applications use application layer protocols to communicate over the network. A client and server must be using the same application layer protocol to communicate. The following table lists some common application layer protocols.

Some common application layer protocols:

- HTTP
- HTTPS
- FTP
- RDP
- SMB
- SMTP
- POP3


Protocol	Description
HTTP	Used for communication between web browsers and web servers.
HTTP/Secure (HTTPS)	A version of HTTP that encrypts communication between web browsers and web servers.
FTP	Used to transfer files between FTP clients and servers.
Remote Desktop Protocol (RDP)	Used to remotely control a computer that is running Windows operating systems over a network.
Server Message Block (SMB)	Used by servers and client computers for file and printer sharing.
Simple Mail Transfer Protocol (SMTP)	Used to transfer email messages over the Internet.
Post Office Protocol version 3 (POP3)	Used to retrieve messages from some email servers.
Internet Message Application Protocol (IMAP)	Used to retrieve messages from some email servers.

What Is a Socket?

When an application wants to establish communication with an application on a remote host, it creates a TCP or a UDP socket, as appropriate. A socket identifies the following as part of the communication process:

- The transport protocol that the application uses, which could be TCP or UDP
- The TCP or UDP port numbers that the applications are using
- The IPv4 or IPv6 address of the source and destination hosts

A socket is a combination of an IP address, a transport protocol, and a port



This combination of transport protocol, IP address, and port creates a socket.

Well-Known Ports

Applications are assigned a port number between 0 and 65,535. The first 1,024 ports are known as *well-known ports* and have been assigned to specific applications. Applications listening for connections use consistent port numbers to make it easier for client applications to connect. If an application listens on a non-standard port number, then you need to specify the port number when connecting to it. Client applications typically use a random source port number above 1,024. The following table identifies some of these well-known ports.

Port	Protocol	Application
80	TCP	HTTP used by a web server
443	TCP	HTTPS for a secure web server
110	TCP	POP3 used for email retrieval
143	TCP	IMAP used for email retrieval
25	TCP	SMTP used for sending email messages
53	UDP	DNS used for most name resolution requests
53	TCP	DNS used for zone transfers
20, 21	TCP	FTP used for file transfers

You need to know the port numbers that applications use so you can configure firewalls to allow communication. Most applications have a default port number for this purpose, but it can be changed when required. For example, some web-based applications run on a port other than port 80 or port 443.

Question: Are there other well-known ports that you can think of?

Lesson 2

Understanding IPv4 Addressing

Understanding IPv4 network communication is critical to ensuring that you can implement, troubleshoot, and maintain IPv4 networks. One of the core components of IPv4 is addressing. Understanding addressing, subnet masks, and default gateways allows you to identify the proper communication between hosts. To identify IPv4 communication errors, you need to understand how the communication process is supposed to work.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe IPv4 Addressing.
- Identify public and private IPv4 addresses.
- Explain how dotted decimal notation relates to binary numbers.
- Describe a simple IPv4 network with classfull addressing.
- Describe a more complex IPv4 network with classless addressing.

IPv4 Addressing

To configure network connectivity, you must be familiar with IPv4 addresses and how they work. Network communication for a computer is directed to the IPv4 address of that computer. Therefore, each networked computer must be assigned a unique IPv4 address.

Each IPv4 address is 32 bits long. To make IP addresses more readable, they are displayed in dotted decimal notation. Dotted decimal notation divides a 32-bit IPv4 address into four groups of 8 bits, which are converted to a decimal number between zero and 255. The decimal numbers are separated by a period (dot). Each decimal number is called an *octet*, for example, the following IP address: 172.16.0.10.

- Each networked computer must be assigned a unique IPv4 address
- Network communication for a computer is directed to the IPv4 address of the computer
- Each IPv4 address contains:
 - ✓ Network ID, identifying the network
 - ✓ Host ID, identifying the computer
- The subnet mask identifies which part of the IPv4 address is the network ID (255) and the host ID (0)

IP address	172	16	0	10
Subnet mask	255	255	0	0
Network ID	172	16	0	0
Host ID	0	0	0	10



Subnet Mask

Each IPv4 address is composed of a network ID and a host ID. The *network ID* identifies the network on which the computer is located. The *host ID* uniquely identifies the computer on that specific network. A *subnet mask* identifies which part of an IPv4 address is the network ID, and which part is the host ID.

In the simplest scenarios, each octet in a subnet mask is either 255 or 0. A 255 represents an octet that is part of the network ID, while a 0 represents an octet that is part of the host ID. For example, a computer with an IP address of 172.16.0.10 and a subnet mask of 255.255.0.0 has a network ID of 172.16.0.0 and a host ID of 0.0.0.10.



Note: The terms network, subnet, and VLAN (Virtual Local Area Network) are often used interchangeably. A large network is often subdivided into subnets, and VLANs are configured on routers or on Layer 3 switches to represent subnets.

Default Gateway

A *default gateway* is a device (usually a router), on a TCP/IP network that forwards IP packets to other networks. The multiple internal networks in an organization can be referred to as an *intranet*.

On an intranet, any given network might have several routers that connect it to other networks, both local and remote. You must configure one of the routers as the default gateway for local hosts. This enables the local hosts to communicate with hosts on remote networks.

Before a host sends an IPv4 packet, it uses its own subnet mask to determine whether the destination host is on the same network, or on a remote network. If the destination host is on the same network, the sending host transmits the packet directly to the destination host. If the destination host is on a different network, the host transmits the packet to a router for delivery.

When a host transmits a packet to a remote network, IPv4 consults the internal routing table to determine the appropriate router for the packet to reach the destination subnet. If the routing table does not contain any routing information about the destination subnet, IPv4 forwards the packet to the default gateway. The host assumes that the default gateway contains the required routing information. The default gateway is used in most cases.

Client computers usually obtain their IP addressing information from a Dynamic Host Configuration Protocol (DHCP) server. This is more straightforward than assigning a default gateway manually on each host. Most servers have a static IP configuration that is assigned manually.

Question: How is network communication affected if a default gateway is configured incorrectly?

Public and Private IPv4 Addresses



Devices and hosts that connect directly to the Internet require a public IPv4 address. Hosts and devices that do not connect directly to the Internet do not require a public IPv4 address.

Public IPv4 Addresses

Public IPv4 addresses must be unique. Internet Assigned Numbers Authority (IANA) assigns public IPv4 addresses to regional Internet registries (RIRs). RIRs then assign IPv4 addresses to Internet service providers (ISPs). Usually your ISP allocates you one or more public addresses from its address pool. The number of addresses that your ISP allocates to you depends upon how many devices and hosts that you have to connect to the Internet.

Private IPv4 Addresses

The pool of IPv4 addresses is becoming smaller, so RIRs are reluctant to allocate superfluous IPv4 addresses. Technologies such as network address translation (NAT) enable administrators to use a relatively small number of public IPv4 addresses and at the same time, enable local hosts to connect to remote hosts and services on the Internet.

Public	Private
<ul style="list-style-type: none"> • Required by devices and hosts that connect directly to the Internet • Must be globally unique • Routable on the Internet • Must be assigned by IANA/RIR 	<ul style="list-style-type: none"> • Not routable on the Internet <ul style="list-style-type: none"> • 10.0.0.0/8 • 172.16.0.0/12 • 192.168.0.0/16 • Can be assigned locally by organization • Must be translated to access the Internet
	

IANA defines the address ranges in the following table as private. Internet-based routers do not forward packets originating from, or destined to, these ranges.

Network	Range
10.0.0.0/8	10.0.0.0-10.255.255.255
172.16.0.0/12	172.16.0.0-172.31.255.255
192.168.0.0/16	192.168.0.0-192.168.255.255

How Dotted Decimal Notation Relates to Binary Numbers

When you assign IP addresses, you use dotted decimal notation. Dotted decimal notation is based on the decimal number system. However, in the background, computers use IP addresses in binary. To understand how to choose a subnet mask for complex networks, you must understand IP addresses in binary.

Within an 8-bit octet, each bit position has a decimal value. A bit that is set to 0 always has a zero value. A bit that is set to 1 can be converted to a decimal value. The *low-order bit*—the rightmost bit in the octet—represents a decimal value of 1. The *high-order bit*—the leftmost bit in the octet—represents a decimal value of 128. If all bits in an octet are set to 1, then the octet’s decimal value is 255 (that is: 128 + 64 + 32 + 16 + 8 + 4 + 2 + 1). That is the highest possible value of an octet.

Dotted decimal notation is based on the decimal number system, but computers use IP addresses in binary

- Within an 8-bit octet, each bit position has a decimal value
 - A bit that is set to 0 always has a zero value
 - A bit that is set to 1 can be converted to a decimal value
 - The low-order bit represents a decimal value of 1
 - The high-order bit represents a decimal value of 128
- If all bits in an octet are set to 1, then the octet’s decimal value is 255, the highest possible value of an octet:
 - 128 + 64 + 32 + 16 + 8 + 4 + 2 + 1

Most of the time, you can use a calculator to convert decimal numbers to binary and vice versa. The Calculator application included in Windows operating systems can perform decimal-to-binary conversions, as shown in the following example.

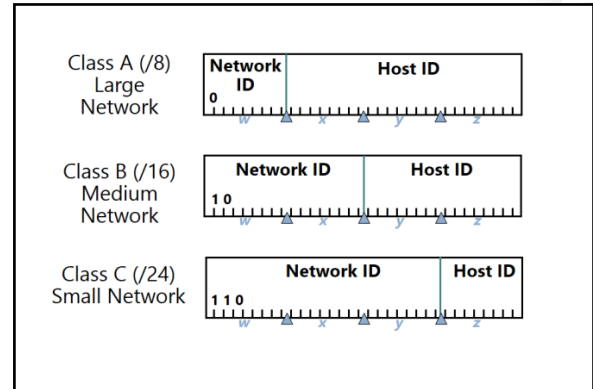
Binary	Dotted decimal notation
10000011 01101011 00000011 00011000	131.107.3.24

Simple IPv4 Implementations

IPv4 Address Classes

The IANA organizes IPv4 addresses into classes. Each class of address has a different default subnet mask that defines the number of valid hosts on the network. IANA has named the IPv4 address classes from *Class A* through *Class E*.

Classes A, B, and C are IP networks that you can assign to IP addresses on host computers. Class D addresses are used by computers and applications for multicasting. The IANA reserves Class E for experimental use. The following table lists the characteristics of each IP address class.



Class	First octet	Default subnet mask	Number of networks	Number of hosts per network
A	1-127	255.0.0.0	126	16,777,214
B	128-191	255.255.0.0	16,384	65,534
C	192-223	255.255.255.0	2,097,152	254



Note: The Internet no longer uses routing based on the default subnet mask of IPv4 address classes.

Simple IPv4 Networks

You can use subnetting to divide a large network into multiple smaller networks. In simple IPv4 networks, the subnet mask defines full octets as part of the network ID and host ID. A 255 represents an octet that is part of the network ID, and a 0 represents an octet that is part of the host ID. For example, you can use the 10.0.0.0 network with a subnet mask of 255.255.0.0 to create 256 smaller networks.



Note: The IPv4 address 127.0.0.1 is used as a loopback address; you can use this address to test the local configuration of the IPv4 protocol stack. Consequently, the network address 127 is not permitted for configuring IPv4 hosts.

More Complex IPv4 Implementations

In complex networks, subnet masks might not be simple combinations of 255 and 0. Rather, you might subdivide one octet with some bits that are for the network ID, and some that are for the host ID. This allows you to have the specific number of subnets and hosts that you require.

The following example shows a subnet mask that can be used to divide a class B network into 16 subnets: **172.16.0.0/255.255.240.0**


In many cases, rather than using a dotted decimal representation of the subnet mask, the number of bits in the network ID is specified instead. This is called *Classless Interdomain Routing (CIDR)*.

The following is an example of CIDR notation: **172.16.0.0/20**

Variable Length Subnet Masks

Modern routers support the use of variable length subnet masks (VLSMs). VLSMs allow you to create subnets of different sizes when you subdivide a larger network. For example, you could subdivide a small network with 256 addresses into three smaller networks with 128 addresses, 64 addresses, and 64 addresses. This allows you to use IP addresses in a network more efficiently.

Question: Does your organization use simple or complex networking?



Lesson 3

Subnetting and Supernetting

In most organizations, you need perform subnetting to divide your network into smaller subnets and allocate those subnets for specific purposes or locations. To do this you need to understand how to select the correct number of bits to include in the subnet masks. In some cases, you may also need to combine multiple networks into a single larger network through supernetting.

Lesson Objectives


After completing this lesson, you will be able to:

- Describe how bits are used in a subnet mask or prefix length.
- Identify when to use subnetting.
- Calculate a subnet mask that supports a specific number of subnet addresses.
- Calculate a subnet mask that supports a specific number of host addresses.
- Identify an appropriate subnet mask for a scenario.
- Describe supernetting.

How Bits Are Used in a Subnet Mask or Prefix Length

In simple networks, subnet masks are composed of four octets, and each octet has a value of 255 or 0. If the octet is 255, that octet is part of the network ID. If the octet is 0, that octet is part of the host ID.

In complex networks, you can convert the subnet mask to binary, and evaluate each bit in the subnet mask. A subnet mask is composed of contiguous 1s and 0s. The 1s start at the leftmost bit and continue uninterrupted until the bits change to all 0s.



Note: Windows PowerShell® cmdlets for configuring IPv4 use a prefix length value rather than a subnet mask to define the number of network bits. The prefix length is the same number of bits used by CIDR notation.

You can identify the network ID of a subnet mask by the 1s. You can identify the host ID by the 0s. Any bits taken from the host ID and allocated to the network ID must be contiguous with the original network ID:

- Each 1 bit is part of the network ID.
- Each 0 bit is part of the host ID.

The mathematical process that is used to compare an IP address and a subnet mask is called *ANDing*.

When you use more bits for the subnet mask, you can have more subnets, but fewer hosts on each subnet. Using more bits than you need allows for subnet growth, but limits growth for hosts. Using fewer bits than you need allows for growth in the number of hosts you can have, but limits growth in subnets.

The Benefits of Using Subnetting

When you subdivide a network into subnets, you must create a unique ID for each subnet. These unique IDs are derived from the main network ID—you allocate some of the bits in the host ID to the network ID. This enables you to create more networks.

By using subnets, you can:

- Use a single, large network across multiple physical locations.
- Reduce network congestion by segmenting traffic and reducing broadcasts on each segment.
- Increase security by dividing the network and using firewalls to control communication.
- Overcome limitations of current technologies, such as exceeding the maximum number of hosts that each segment can have.

When you subdivide a network into subnets, you create a unique ID for each subnet that is derived from the main network ID

By using subnets, you can:

- Use a single network address across multiple locations
- Reduce network congestion by segmenting traffic
- Increase security by using firewalls
- Overcome limitations of current technologies

Calculating Subnet Addresses

Before you define a subnet mask, estimate how many subnets and hosts for each subnet you may require. This enables you to use the appropriate number of bits for the subnet mask.

You can calculate the number of subnet bits that you need in the network. Use the formula 2^n , where n is the number of bits. The result is the number of subnets that your network requires.

The following table indicates the number of subnets that you can create by using a specific number of bits.

Number of bits (n)	Number of subnets (2^n)
1	2
2	4
3	8
4	16
5	32
6	64

When determining subnet addresses you should:

- Choose the number of subnet bits based on the number of subnets required
- Use 2^n to determine the number of subnets available from n bits

For five locations, the following three subnet bits are required:

- 5 locations = 5 subnets required
- $2^2 = 4$ subnets (not enough)
- $2^3 = 8$ subnets

To determine the subnet addresses quickly, you can use the lowest value bit in the subnet mask. For example, if you choose to subnet the network 172.16.0.0 by using 3 bits, this means the subnet mask is 255.255.224.0. The decimal 224 is 11100000 in binary, and the lowest bit has a value of 32, so that is the increment between each subnet address.

The following table shows the subnet addresses for this example; the 3 bits that you have chosen to use to subnet the network are shown in bold type.

Binary network number	Decimal network number
172.16. 000 00000.00000000	172.16.0.0
172.16. 001 00000.00000000	172.16.32.0
172.16. 010 00000.00000000	172.16.64.0
172.16. 011 00000.00000000	172.16.96.0
172.16. 100 00000.00000000	172.16.128.0
172.16. 101 00000.00000000	172.16.160.0
172.16. 110 00000.00000000	172.16.192.0
172.16. 111 00000.00000000	172.16.224.0



Note: You can use a subnet calculator to determine the appropriate subnets for your network, rather than calculating them manually. Subnet calculators are widely available on the Internet.

Calculating Host Addresses

To determine host bits in the mask, determine the required number of bits for the supporting hosts on a subnet. Calculate the number of host bits required by using the formula $2^n - 2$, where n is the number of bits. This result must be at least the number of hosts that you need for your network, and the maximum number of hosts that you can configure on that subnet.

On each subnet, two host IDs are allocated automatically and cannot be used by computers. An address with the host ID of all 0s represents the network. An address with the host ID of all 1s is the broadcast address for that network.

When determining host addresses you should:

- Choose the number of host bits based on the number of hosts that you require on each subnet
- Use $2^n - 2$ to determine the number of hosts that are available on each subnet

For subnets with 100 hosts, seven host bits are required:

- $2^6 - 2 = 62$ hosts (not enough)
- $2^7 - 2 = 126$ hosts

The following table shows how many hosts a class C network has available based on the number of host bits.

Number of bits (n)	Number of hosts (2^n-2)
1	0
2	2
3	6
4	14
5	30
6	62

You can calculate each subnet's range of host addresses by using the following process:

1. The first host is one binary digit higher than the current subnet ID.
2. The last host is two binary digits lower than the next subnet ID.

The following table shows examples of calculating host addresses.

Network	Host range
172.16.64.0/19	172.16.64.1 – 172.16.95.254
172.16.96.0/19	172.16.96.1 – 172.16.127.254
172.16.128.0/19	172.16.128.1 – 172.16.159.254

To create an appropriate addressing scheme for your organization, you must know how many subnets you need, and how many hosts you need on each subnet. Once you have that information, you can calculate an appropriate subnet mask.

Discussion: Creating a Subnetting Scheme for a New Office

Read the following scenario and answer the questions on the slide.

You are identifying an appropriate network configuration for a new campus. You have been allocated the 10.34.0.0/16 network that you can subnet as required.

There are four buildings on the new campus, and each should have its own subnet to allow for routing between the buildings. Each building will have up to 700 users. Each building will also have printers. The typical ratio of users to printers is 50 to 1.

You also need to allocate a subnet for the server data center that will hold up to 100 servers.

- How many subnets are required?
- How many bits are required to create that number of subnets?
- How many hosts are required on each subnet?
- How many bits are required to allow that number of hosts?
- What is an appropriate subnet mask to meet these needs?



20 minutes

What Is Supernetting?

Supernetting combines multiple small networks into a single large network. This may be appropriate when you have a small network that has grown and you need to expand the address space. For example, a branch office that is using the network 192.168.16.0/24 and that exhausts all of its IP addresses could be allocated the additional network 192.168.17.0/24. If you use the default subnet mask of 255.255.255.0 for these networks then you must perform routing between them. You can use supernetting to combine them into a single network.

- Supernetting combines multiple small networks into a larger network
- The networks that you are combining must be contiguous
- The following table shows an example of supernetting two class C networks

Network	Range
192.168. 00010000 .00000000/24	192.168.16.0 - 192.168.16.255
192.168. 00010001 .00000000/24	192.168.17.0 - 192.168.17.255
192.168. 00010000 .00000000/23	192.168.16.0 - 192.168.17.255

To perform supernetting, the networks that you are combining must be contiguous. For example, 192.168.16.0/24 and 192.168.17.0/24 can be supernetted, but you cannot supernet 192.168.16.0/24 and 192.168.54.0/24.

Supernetting is the opposite of subnetting. When you perform supernetting, you allocate bits from the network ID to the host ID. The following table shows how many networks that you can combine by using a specific number of bits.

Number of bits	Number of networks combined
1	2
2	4
3	8
4	16

The following table shows an example of supernetting two class C networks. The portion of the subnet mask that you are using as part of the network ID is shown in bold type.

Network	Range
192.168. 00010000 .00000000/24	192.168.16.0-192.168.16.255
192.168. 00010001 .00000000/24	192.168.17.0-192.168.17.255
192.168. 00010000 .00000000/23	192.168.16.0-192.168.17.255

Lesson 4

Configuring and Troubleshooting IPv4

If IPv4 is configured incorrectly, then it affects the availability of services that are running on a server. To ensure the availability of network services, you need to understand how to configure and troubleshoot IPv4. Windows Server 2012 introduces the ability to configure IPv4 by using Windows PowerShell which is useful for scripting.

The troubleshooting tools in Windows Server 2012 are similar to previous versions of Windows client operating systems and server operating systems. However, you may not be familiar with Network Monitor, which you can use to perform detailed analysis of your network communication.

Lesson Objectives

After completing this lesson, you will be able to:

- Configure IPv4 manually to provide a static configuration for a server.
- Configure a server so that it obtains an IPv4 configuration automatically.
- Explain how to use IPv4 troubleshooting tools.
- Explain how to use Windows PowerShell cmdlets for troubleshooting IPv4.
- Describe the troubleshooting process used to resolve fundamental IPv4 problems.
- Describe the function of Network Monitor.
- Use Network Monitor to capture and analyze network traffic.

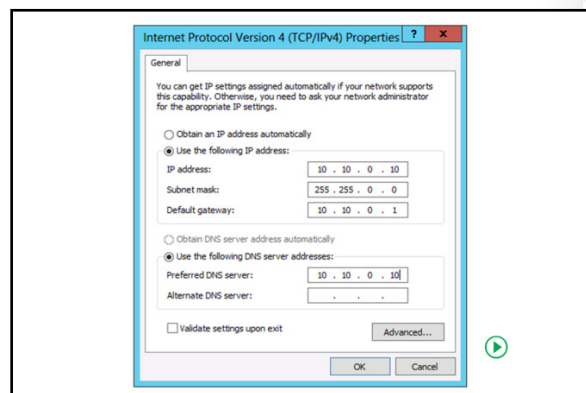
Configuring IPv4 Manually

You typically configure servers with a static IP address. This is done to ensure that you know and can document the IP addresses that you use for various services on your network. For example, a DNS server is accessed at a specific IP address that should not change.

IPv4 configuration includes:

- IPv4 address
- Subnet mask
- Default gateway
- DNS servers

Static configuration requires that you visit each computer and input the IPv4 configuration manually. This method of computer management is reasonable for servers, but it is very time consuming for client computers. Manually entering a static configuration also increases the risk of configuration mistakes.



Configuring a Static IP Address by Using Windows PowerShell

Windows Server 2012 includes Windows PowerShell cmdlets that you can use to manage network configuration. The following table describes some of the Windows PowerShell cmdlets that are available for configuring IPv4.

Cmdlet	Description of IPv4 configuration uses
New-NetIPAddress	Creates a new IP address and binds it to a network adapter. You cannot modify an existing IP address; you must remove an existing IP address and then create a new IP address.
Set-NetIPAddress	Changes the configuration of an IP address.
Set-NetIPInterface	Enables or disables DHCP for an interface.
New-NetRoute	Creates routing table entries, including the default gateway (0.0.0.0). You cannot modify the next hop of an existing route; instead, you must remove an existing route and create a new route with the correct next hop.
Set-DNSClientServerAddress	Configures the DNS server that is used for an interface.

The following code is an example of the Windows PowerShell cmdlets that you can use to configure the interface Local Area Connection with the following parameters:

- Static IP address 10.10.0.10
- Subnet mask 255.255.255.0
- Default gateway 10.10.0.1

Local Area Connection is also configured to use DNS servers of 10.12.0.1 and 10.12.0.2.

```
New-NetIPAddress -InterfaceAlias "Local Area Connection" -IPAddress 10.10.0.10
-PrefixLength 24 -DefaultGateway 10.10.0.1

Set-DNSClientServerAddress -InterfaceAlias "Local Area Connection" -ServerAddresses
10.12.0.1,10.12.0.2
```

Configuring a Static IP Address by Using Netsh

You also can configure a static IP address either in the properties of the network connection or by using the netsh command-line tool. For example, the following command configures the interface Local Area Connection with the following parameters:

- Static IP address 10.10.0.10
- Subnet mask 255.255.255.0
- Default gateway 10.10.0.1

```
Netsh interface ipv4 set address name="Local Area Connection" source=static
addr=10.10.0.10 mask=255.255.255.0 gateway=10.10.0.1
```



Additional Reading: For more information about net TCP/IP cmdlets in Windows PowerShell, go to <http://go.microsoft.com/fwlink/?LinkId=269708>.

Question: Do any computers or devices in your organization have static IP addresses?

Configuring IPv4 Automatically

DHCP for IPv4 enables you to assign automatic IPv4 configurations for large numbers of computers without having to assign each one individually. The DHCP service receives requests for IPv4 configuration from computers that you configure to obtain an IPv4 address automatically. It also assigns additional IPv4 settings from scopes that you define for each of your network's subnets. The DHCP service identifies the subnet from which the request originated and assigns IP configuration from the relevant scope.

DHCP helps simplify the IP configuration process; however, you must be aware that if you use DHCP to assign IPv4 information and the service is business-critical, you must do the following:

- Include resilience in your DHCP service design so that the failure of a single server does not prevent the service from functioning.
- Configure the scopes on the DHCP server carefully. If you make a mistake, it can affect the entire network and prevent communication.

If you use a laptop to connect to multiple networks (such as at work and at home), each network might require a different IP configuration. Windows operating systems support the use of Automatic Private IP Addressing (APIPA) or an alternate static IP address for this situation.

When you configure Windows-based computers to obtain an IPv4 address from DHCP, use the Alternate Configuration tab to control the behavior if a DHCP server is not available. By default, Windows uses APIPA to assign itself an IP address automatically from the 169.254.0.0 to 169.254.255.255 address range, but with no default gateway or DNS server; this enables limited functionality.

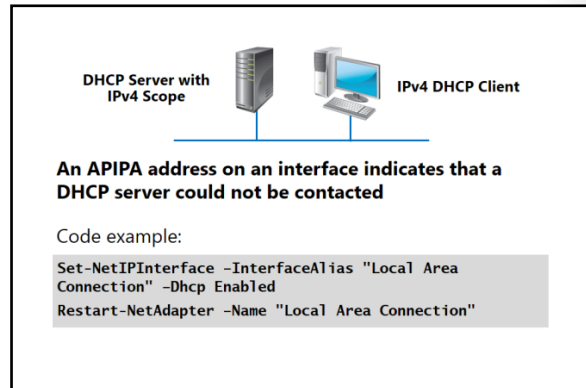
APIPA is useful for troubleshooting DHCP; if the computer has an address from the APIPA range, it is an indication that the computer cannot communicate with a DHCP server.

Windows Server 2012 also has Windows PowerShell cmdlets that you can use to enable DHCP for an interface. The following table describes some of the available Windows PowerShell cmdlets that are available for configuring DHCP on an interface.

Cmdlet	Description
Get-NetIPInterface	Obtains a list of interfaces and their configuration. This does not include IPv4 configuration of the interface.
Set-NetIPInterface	Enables or disables DHCP for an interface.
Get-NetAdapter	Obtains a list of network adapters in a computer.
Restart-NetAdapter	Disables and re-enables a network adapter. This forces a DHCP client to obtain a new DHCP lease.

The following code is an example of how you can enable DHCP for the adapter Local Area Connection, and ensure that it receives an address:

```
Set-NetIPInterface -InterfaceAlias "Local Area Connection" -Dhcp Enabled
Restart-NetAdapter -Name "Local Area Connection"
```



Using Windows PowerShell Cmdlets to Troubleshoot IPv4

You can use command-line tools or Windows PowerShell cmdlets in Windows Server 2012 to configure your network and perform troubleshooting. Although you could use Windows PowerShell in earlier versions of Windows Server to perform network troubleshooting and configuration, it required you to use Windows Management Instrumentation (WMI) objects, which are more difficult to use than native Windows PowerShell cmdlets.

The following table lists some of the Windows PowerShell cmdlets that you can use.

For network configuration and troubleshooting in Windows Server 2012 and earlier versions, use:

- Command-line tools
- Windows PowerShell cmdlets

In earlier versions of Windows Server:

- Windows PowerShell required using Windows Management Instrumentation (WMI) objects
- WMI objects are not as easy to use as native Windows PowerShell cmdlets

Cmdlet	Purpose
Get-NetAdapter	Obtains a list of network adapters in a computer.
Get-NetIPv4Protocol	Gets information about the IP version 4 (IPv4) protocol configuration. (Get-NetIP6Protocol returns the same information for the IP version 6 (IPv6) protocol).
Restart-NetAdapter	Disables and re-enables a network adapter.
Get-NetIPInterface	Obtains a list of interfaces and their configuration.
Get-NetIPAddress	Obtains a list of IP addresses that are configured for interfaces.
Get-NetRoute	Obtains the list of routes in the local routing table.
Get-NetConnectionProfile	Obtains the type of network (public, private, domain) to which a network adapter is connected.
Get-DnsClient	Retrieves configuration details specific to the different network interfaces on a specified computer
Get-DNSClientCache	Obtains the list of resolved DNS names that are stored in the DNS client cache.
Get-DnsClientGlobalSetting	Retrieves global DNS client settings such as the suffix search list.
Get-DNSClientServerAddress	Obtains the list of DNS servers that are used for each interface.
Register-DnsClient	Registers all of the IP addresses on the computer on the configured DNS server.
Set-DnsClient	Sets the interface-specific DNS client configurations on the computer.
Set-DnsClientGlobalSetting	Configures the global DNS client settings such as the suffix search list.
Set-DnsClientServerAddress	Configures the computer's network adapter with the IP addresses of the DNS server.

(continued from previous page)

Cmdlet	Purpose
Set-NetIPAddress	Sets information about the IP address configuration.
Set-NetIPv4Protocol	Sets information about the IPv4 protocol configuration. (Set-NetIP6Protocol returns the same information for the IPv6 protocol.)
Set-NetIPInterface	Modifies the IP interface properties.
Test-Connection	Runs connectivity tests that are similar to those that are used by ping.
Test-NetConnection	<ul style="list-style-type: none"> • Displays the following: • Results of a DNS lookup • Listing of IP interfaces • Option to test a TCP connection • IPsec rules • Confirmation of connection establishment
Resolve-Dnsname	Performs a DNS name query resolution for the specified name

IPv4 Troubleshooting Tools

Windows Server 2012 includes a number of command-line tools that can help you diagnose network problems. These tools were commonly used in earlier Windows Server editions.

Ipconfig

Ipconfig is a command-line tool that displays the current TCP/IP network configuration.

Additionally, you can use the **ipconfig** command to refresh DHCP and DNS settings. The following table describes the command-line options for **ipconfig**.

Use the following tools to troubleshoot IPv4:

- Ipconfig
- Ping
- Tracert
- Pathping
- Telnet
- Netstat
- Resource Monitor
- Windows Network Diagnostics
- Event Viewer

Command	Description
ipconfig /all	View detailed configuration information
ipconfig /release	Release the leased configuration back to the DHCP server
ipconfig /renew	Renew the leased configuration
ipconfig /displaydns	View the DNS resolver cache entries
ipconfig /flushdns	Purge the DNS resolve cache

Ping

Ping is a command-line tool that verifies IP-level connectivity to another TCP/IP computer. It sends ICMP echo request messages and displays the receipt of corresponding echo reply messages. Ping is the primary TCP/IP command that you use to troubleshoot connectivity; however, firewalls might block the ICMP messages.

Tracert

Tracert is a command-line tool that identifies the path taken to a destination computer by sending a series of ICMP echo requests. Tracert then displays the list of router interfaces between a source and a destination. This tool also determines which router has failed, and what the latency (or speed) is. These results might not be accurate if the router is busy, because the ICMP packets are assigned a low priority by the router.

Pathping

Pathping is a command-line tool that traces a route through the network in a manner similar to Tracert. However, Pathping provides more detailed statistics on the individual steps, or *hops*, through the network. Pathping can provide greater detail, because it sends 100 packets for each router, which enables it to establish trends.

Route

Route is a command-line tool that allows you to view and modify the local routing table. You can use this to verify the default gateway, which is listed as the route 0.0.0.0. In Windows Server 2012, you can also use Windows PowerShell cmdlets to view and modify the routing table. The cmdlets for viewing and modifying the local routing table include **Get-NetRoute**, **New-NetRoute**, and **Remove-NetRoute**.

Telnet

You can use the Telnet Client feature to verify whether a server port is listening. For example, the command **telnet 10.10.0.10 25** attempts to open a connection with the destination server, 10.10.0.10, on port 25, SMTP. If the port is active and listening, it returns a message to the Telnet client.

Netstat

Netstat is a command-line tool that enables you to view network connections and statistics. For example, the command **netstat -ab** returns all listening ports and the executable that is listening.

Resource Monitor

Resource Monitor is a graphical tool that allows you to monitor system resource utilization. You can use Resource Monitor to view TCP and UDP ports that are in use. You can also verify which applications are using specific ports and the amount of data that they are transferring on those ports.

Network Diagnostics

Use Windows Network Diagnostics to diagnose and correct networking problems. In the event of a Windows Server networking problem, the **Diagnose Connection Problems** option helps you diagnose and repair the problem. Windows Network Diagnostics returns a possible description of the problem and a potential remedy. However, the solution might require manual intervention from the user.

Event Viewer

Event logs are files that record significant events on a computer, such as when a process encounters an error. When these events occur, the Windows operating system records the event in an appropriate event log. You can use Event Viewer to read the event log. IP conflicts, which might prevent services from starting, are listed in the System event log.

The IPv4 Troubleshooting Process

The first step in troubleshooting a network problem is identifying the scope of the problem. The causes of a problem that affects a single user will most likely differ from a problem that affects all users. If a problem affects only a single user, then the problem is likely related to the configuration of that one computer. If a problem affects all users, then the problem is likely either a server configuration issue or a network configuration issue. If a problem affects only a group of users, then you need to determine the common denominator among that group of users.

After you identify the scope of the problem, use the following tools to troubleshoot network connectivity:

Step	Windows PowerShell	Command-line tool
Identify the network path between hosts	Test-NetConnection -TraceRoute	tracert
Verify the network configuration is correct	Get-NetIPAddress	ipconfig
See if the remote host responds	Test-NetConnection	ping
Test the service on a remote host	Use an application such as Internet Explorer	
See if the default gateway responds	Test-NetConnection	ping

To troubleshoot network communication problems, you need to understand the overall communication process. You can identify where the process is breaking down and preventing communication only when you understand how the overall communication process works. This requires that you understand the routing and firewall configuration on your network.

To help identify the routing path through your network, you can use the Windows PowerShell cmdlet **Test-NetConnection -TraceRoute**, or you can use the command-line tool **tracert**.

The following are some of the steps that you can use to identify that cause of network communication problems:

1. If you know what the correct network configuration for the host should be, then use one of the following to verify that it is configured correctly:
 - Windows PowerShell: **Get-NetIPAddress**
 - Command-line: **ipconfig**

If the address that is returned is on the 169.254.0.0/16 network, it indicates that the host failed to obtain an IP address from DHCP.

2. To see if the remote host responds, use one of the following:
 - Windows PowerShell: **Test-NetConnection**
 - Command-line: **ping**

When you use either method to return the DNS name of the remote host, you verify both name resolution and whether the host responds. Be aware that Windows Firewall on member servers and client computers often blocks ping attempts. When this happens, the lack of a ping response might not be an indicator that the remote host is not functional, but only that the ping is being blocked. If you can ping other remote hosts on the same network, this might mean that the problem is on the remote host.

3. You can use an application to test the service you are connecting to on the remote host. For example, use Windows Internet Explorer® to test connectivity to a web server. You can also use Telnet to connect to the port of the remote application.

4. To see if the default gateway responds, use one of the following:

- Windows PowerShell: **Test-NetConnection**
- Command-line: **ping**

Most routers respond to Test-NetConnection and ping requests. If you do not get a response when you ping the default gateway, then there is likely a configuration error on the client computer, such as the default gateway being configured incorrectly. It is also possible that the router is experiencing errors.



Note: You can force **ping** to use IPv4 instead of IPv6 by using the **-4** option.

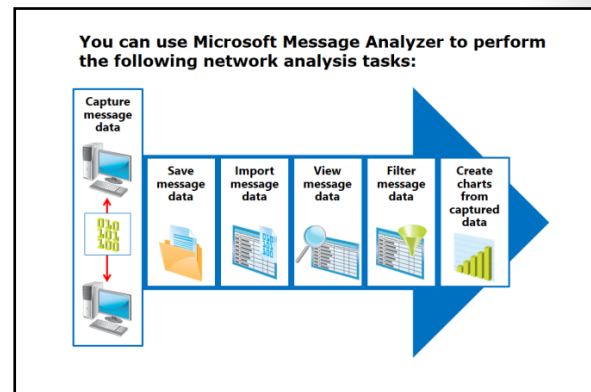
Question: What additional steps might you use to troubleshoot network connectivity problems?

What Is Microsoft Message Analyzer?

Microsoft Message Analyzer is a tool used to capture network traffic and then display and analyze information about that traffic. You can use Microsoft Message Analyzer to monitor live network traffic, or import, aggregate, and analyze data from log and trace files. Microsoft Message Analyzer replaces Network Monitor, which was last released as Network Monitor 3.4.

You can use Microsoft Message Analyzer to perform the following network analysis tasks:

- Capture message data
- Save message data
- Import message data
- View message data
- Filter message data



Microsoft Message Analyzer uses several built-in Trace Scenarios that you can access through the Microsoft Message Analyzer console. Trace Scenarios contain specific capture settings that enable you to quickly start a trace session and then capture the information you need for your troubleshooting task. These Trace Scenarios include predefined capture configuration for Windows Firewall troubleshooting, LAN and WAN monitoring, and Web Proxy troubleshooting. You can customize Trace Scenarios to remove items that do not require monitoring.

The Microsoft Message Analyzer console contains a Charts tab that creates charts from captured data. You can customize the parameters and data that will be included in the charts, including network transactions, operations, and the network protocol. Furthermore, you can define different types of chart views, such as Timeline Chart, Pie Chart, Grid View, or Bar Chart. Charts can help you understand incoming trace data by presenting complicated traffic information visually. Often, this feature is helpful when you need to perform mathematical calculations on the trace data, such as the number of retries required for a packet being sent between hosts.

Microsoft Message Analyzer introduces remote live monitoring, which is a feature that allows administrators to monitor the network from a remote host. Administrators can connect to both remote host network adapters and virtual machine network adapters in order to capture and analyze the network traffic data.

Microsoft Message Analyzer is capable of loading data from native Microsoft Message Analyzer files, event tracing log (.etl) files, Network Monitor capture files (.cap), comma-separated values (CSV) files, and several other formats. You can download Microsoft Message Analyzer free from Microsoft's website.



Reference Links: For more information about Microsoft Message Analyzer, see the Microsoft Message Analyzer Operating Guide at <http://go.microsoft.com/fwlink/?LinkID=331073>. To download Microsoft Message Analyzer, go to <http://go.microsoft.com/fwlink/?LinkID=331072>.

Demonstration: How to Capture and Analyze Network Traffic by Using Microsoft Message Analyzer

You can use Microsoft Message Analyzer to capture and view packets that are transmitted on a network. This allows you to view detailed information that you would not normally be able to see. This type of information can be useful for troubleshooting.

In this demonstration, you will see how to:

- Capture network traffic with Microsoft Message Analyzer.
- Analyze captured network traffic.
- Filter network traffic.

Demonstration Steps

Start a new Capture/Trace in Microsoft Message Analyzer

1. Sign in to **LON-SVR2** as **Adatum\Administrator** with a password of **Pa\$\$w0rd**.
2. Open a Windows PowerShell prompt and run the following command:

```
ipconfig /flushdns
```
3. From the Start screen open **Microsoft Message Analyzer**, choose **Do not update items**, and start a new **Capture/Trace** for using the **Firewall** trace scenario.

Capture packets from a ping request

1. In Microsoft Message Analyzer, start a packet capture.
2. At the Windows PowerShell prompt, run following cmdlet:

```
Test-NetConnection LON-DC1.adatum.com
```

3. In Microsoft Message Analyzer, stop the packet capture.

Analyze the captured network traffic

1. In Microsoft Message Analyzer, in the results pan, under the **Module** column, select the first **ICMP** packet group.
2. Expand the **ICMP** portion of the packet to view that it includes both **Echo Request and Echo Reply packets**. This is a **ping** request that was executed when running **Test-NetConnection** cmdlet.
3. View the source and destination IP addresses for each packet.

Filter the network traffic

1. In Microsoft Message Analyzer, enter the following filter criteria, and then apply the filter:
*DestinationAddress == 172.16.0.10
2. Verify that only packets that match the filter are displayed.
3. Close Microsoft Message Analyzer.

Lab: Implementing IPv4

Scenario

You have recently accepted a promotion to the server support team. One of your first assignments is configuring the infrastructure service for a new branch office.

After a security review, your manager has asked you to calculate new subnets for the branch office to support segmenting network traffic. You also need to troubleshoot a connectivity problem on a server in the branch office.

Objectives

After completing this lab, you will be able to:

- Identify appropriate subnets for a given set of requirements.
- Troubleshoot IPv4 connectivity issues.

Lab Setup

Estimated Time: 45 minutes

Virtual machines	20410C-LON-DC1 20410C-LON-RTR 20410C-LON-SVR2
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 **20410C-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in using the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
5. Repeat steps 2-4 for **20410C-LON-RTR**, and **20410C-LON-SVR2**.

Exercise 1: Identifying Appropriate Subnets

Scenario

The new branch office is configured with a single subnet. After a security review, all branch office network configurations are being modified to place servers on a separate subnet from the client computers. You need to calculate the new subnet mask and the default gateways for the subnets in your branch.

The current network for your branch office is 192.168.98.0/24. This network needs to be subdivided into three subnets as follows:

- One subnet with at least 100 IP addresses for clients

- One subnet with at least 10 IP addresses for servers
- One subnet with at least 40 IP addresses for future expansion

The main tasks for this exercise are as follows:

1. Calculate the bits required to support the hosts on each subnet.
2. Calculate subnet masks and network IDs.

► **Task 1: Calculate the bits required to support the hosts on each subnet**

1. How many bits are required to support 100 hosts on the client subnet?
2. How many bits are required to support 10 hosts on the server subnet?
3. How many bits are required to support 40 hosts on the future expansion subnet?
4. If all subnets are the same size, can they be accommodated?
5. Which feature allows a single network to be divided into subnets of varying sizes?
6. How many host bits will you use for each subnet? Use the simplest allocation possible, which is one large subnet and two equal-sized smaller subnets.

► **Task 2: Calculate subnet masks and network IDs**

1. Given the number of host bits allocated, what is the subnet mask that you will use for the client subnet? Calculate the subnet mask in binary and decimal.
 - The client subnet is using 7 bits for the host ID. Therefore, you can use 25 bits for the subnet mask.

Binary	Decimal

2. Given the number of host bits allocated, what is the subnet mask that you will use for the server subnet? Calculate the subnet mask in binary and decimal.
 - The server subnet is using 6 bits for the host ID. Therefore, you will use 26 bits for the subnet mask.

Binary	Decimal

3. Given the number of host bits allocated, what is the subnet mask that you can use for the future expansion subnet? Calculate the subnet mask in binary and decimal.
 - The future expansion subnet is using 6 bits for the host ID. Therefore, you will use 26 bits for the subnet mask.

Binary	Decimal

4. For the client subnet, define the network ID, first available host, last available host, and broadcast address. Assume that the client subnet is the first subnet allocated from the available address pool. Calculate the binary and decimal versions of each address.

Description	Binary	Decimal
Network ID		
First host		
Last host		
Broadcast		

5. For the server subnet, define the network ID, first available host, last available host, and broadcast address. Assume that the server subnet is the second subnet allocated from the available address pool. Calculate the binary and decimal versions of each address.

Description	Binary	Decimal
Network ID		
First host		
Last host		
Broadcast		

6. For the future allocation subnet, define the network ID, first available host, last available host, and broadcast address. Assume that the future allocation subnet is the third subnet allocated from the available address pool. Calculate the binary and decimal versions of each address.

Description	Binary	Decimal
Network ID		
First host		
Last host		
Broadcast		

Results: After completing this exercise, you should have identified the subnets required to meet the requirements of the lab scenario.

Exercise 2: Troubleshooting IPv4

Scenario

A server in the branch office is unable to communicate with the domain controller in the head office. You need to resolve the network connectivity problem.

The main tasks for this exercise are as follows:

1. Prepare for troubleshooting.
2. Troubleshoot IPv4 connectivity between LON-SVR2 and LON-DC1.

► Task 1: Prepare for troubleshooting

1. On LON-SVR2, open **Windows PowerShell**.
2. In the Windows PowerShell window, run the following cmdlet:

```
Test-NetConnection LON-DC1
```

3. Verify that you receive a reply that contains **PingSucceeded:True** from **LON-DC1**.
4. Open a **File Explorer** window, and browse to **\\LON-DC1\E\$\Labfiles\Mod05**.
5. From File Explorer, run the **Break2.ps1** script by using PowerShell.



Note: This script creates the problem that you will troubleshoot and repair in the next task.

6. Close File Explorer.

► Task 2: Troubleshoot IPv4 connectivity between LON-SVR2 and LON-DC1

1. Use your knowledge of IPv4 to troubleshoot and repair the connectivity problem between LON-SVR2 and LON-DC1. Consider using the following tools:
 - **Test-NetConnection**
 - **Test-NetConnection -TraceRoute**
 - **Get-NetRoute**
 - **New-NetRoute**
2. When you have repaired the problem, run the **Test-NetConnection LON-DC1** cmdlet from LON-SVR2 to confirm that the problem is resolved.

Results: After completing this lab, you should have resolved an IPv4 connectivity problem.

► Prepare for the next module

When you finish the lab, revert the virtual machines back 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 **20410C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20410C-LON-RTR** and **20410C-LON-SVR2**.

Lab Review Questions

Question: Why is variable-length subnetting required in this lab?

Question: Which Windows PowerShell cmdlet can you use to view the local routing table of a computer instead of using **route print**?

Module Review and Takeaways

Module Review Questions

Question: You have just started as a server administrator for a small organization with a single location. The organization is using the 131.107.88.0/24 address range for the internal network. Is this a concern?

Question: You are working for an organization that provides web hosting services to other organizations. You have a single /24 network from your ISP for the web hosts. You are almost out of IPv4 addresses and have asked ISP for an additional range of addresses. Ideally, you would like to supernet the existing network with the new network. Are there any specific requirements for supernetting?

Question: You have installed a new web-based application that runs on a non-standard port number. A colleague is testing access to the new web-based application, and indicates that he cannot connect to it. What are the most likely causes of his problem?

Best Practices

When implementing IPv4, use the following best practices:

- Allow for growth when planning IPv4 subnets. This ensures that you do not need to change your IPv4 configuration scheme.
- Define purposes for specific address ranges and subnets. This enables you to easily identify hosts based on their IP address, and to use firewalls to increase security.
- Use dynamic IPv4 addresses for clients. It is much easier to manage the IPv4 configuration for client computers by using DHCP than with manual configuration.
- Use static IPv4 addresses for servers. When servers have a static IPv4 address, it is easier to identify where services are located on the network.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
IP conflicts	
Multiple default gateways defined	
Incorrect IPv4 configuration	

Tools

Tool	Use	Where to find it
Microsoft Message Analyzer	Capture and analyze network traffic	Download from Microsoft website
Ipconfig	View network configuration	Command prompt
Ping	Verify network connectivity	Command prompt
Tracert	Verify network path between hosts	Command prompt
Pathping	Verify network path and reliability between hosts	Command prompt
Route	View and configure the local routing table	Command prompt
Telnet	Test connectivity to a specific port	Command prompt
Netstat	View network connectivity information	Command prompt
Resource monitor	View network connectivity information	Tools in Server Manager
Windows Network Diagnostics	Diagnose problem with a network connection	Properties of the network connection
Event Viewer	View network related system events	Tools in Server Manager

Module 6

Implementing Dynamic Host Configuration Protocol

Contents:

Module Overview	6-1
Lesson 1: Installing a DHCP Server Role	6-2
Lesson 2: Configuring DHCP Scopes	6-9
Lesson 3: Managing a DHCP Database	6-15
Lesson 4: Securing and Monitoring DHCP	6-19
Lab: Implementing DHCP	6-24
Module Review and Takeaways	6-29

Module Overview

Dynamic Host Configuration Protocol (DHCP) plays an important role in the Windows Server® 2012 infrastructure. It is the primary means of distributing important network configuration information to network clients, and it provides configuration information to other network-enabled services, including Windows® Deployment Services (Windows DS) and Network Access Protection (NAP). To support and troubleshoot a Windows Server–based network infrastructure, it is important that you understand how to deploy, configure, and troubleshoot the DHCP server role.

Objectives

After completing this module, you will be able to:

- Install the DHCP server role.
- Configure DHCP scopes.
- Manage a DHCP database.
- Secure and monitor the DHCP server role.

Lesson 1

Installing a DHCP Server Role

Using DHCP can help simplify client computer configuration. This lesson describes the benefits of DHCP, explains how the DHCP protocol works, and discusses how to control DHCP in a Windows Server 2012 network with Active Directory® Domain Services (AD DS).

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the benefits of using DHCP.
- Explain how DHCP allocates IP addresses to network clients.
- Explain how the DHCP lease generation process works.
- Explain how the DHCP lease renewal process works.
- Describe the purpose of a DHCP relay agent.
- Describe how DHCP interacts with DNS.
- Explain how a DHCP server role is authorized.
- Explain how to add and authorize the DHCP server role.

Benefits of Using DHCP

The DHCP protocol simplifies configuration of IP clients in a network environment. If you do not use DHCP, each time you add a client to a network, you need to configure it with information about the network on which you installed it, including the IP address, the network's subnet mask, and the default gateway for access to other networks.

When you need to manage many computers in a network, managing them manually can become a time-consuming process. Many corporations manage thousands of computer devices, including handhelds, desktop computers, and laptops. It is not feasible to manually manage the network IP configurations for organizations of this size.

With the DHCP server role, you can help to ensure that all clients have appropriate configuration information, which helps to eliminate human error during configuration. When key configuration information changes in the network, you can update it using the DHCP server role without having to change the information directly on each computer.

DHCP is also a key service for mobile users who change networks often. DHCP enables network administrators to offer complex network-configuration information to nontechnical users, without users having to deal with their network-configuration details.

DHCP version 6 (v6) stateful and stateless configurations are supported for configuring clients in an IPv6 environment. Stateful configuration occurs when the DHCPv6 server assigns the IPv6 address to the client, along with additional DHCP data. Stateless configuration occurs when the subnet router assigns the IPv6 address automatically, and the DHCPv6 server only assigns other IPv6 configuration settings.

DHCP reduces the complexity and amount of administrative work by using automatic IP configuration

Automatic IP Configuration	Manual IP Configuration
IP addresses are supplied automatically	IP addresses are entered manually
Correct configuration information is ensured	IP address could be entered incorrectly
Client configuration is updated automatically	Communication and network issues can result
A common source of network problems is eliminated	Frequent computer moves increase administrative effort

Clients can use the assigned DHCP address for a certain period of time known as a *lease*. You can set the lease time to optimize your overall IP address scheme. Clients attempt to automatically renew their lease after a period of time, usually after 50 percent of the lease period has passed. As long as there are IP addresses available, the DHCP continues to provide the renewals.

NAP

NAP is part of a toolset that can prevent full access to an intranet for computers that do not comply with system health requirements. NAP with DHCP helps isolate potentially malware-infected computers from the corporate network. DHCP NAP enables administrators to ensure that DHCP clients are compliant with internal security policies. For example, all network clients must be up-to-date and have a valid, up-to-date antivirus program installed before they are assigned an IP configuration that allows full access to an intranet.

Installing DHCP

You can install DHCP as a role on a Server Core installation of Windows Server 2012. A Server Core installation allows you to create a server with a reduced attack surface. To manage DHCP from the Server Core, you must install and configure the role from the command-line interface. You also can manage the DHCP role running on Server Core installation of Windows Server 2012 from a graphical user interface (GUI)-based console where the DHCP role is installed already.

Installing Windows Deployment Service and DHCP

Windows Deployment Service (WDS) uses the Preboot Execution Environment (PXE) for WDS clients that are receiving an operating system installation. These clients may not yet have an operating system, but if their network adapters are PXE-enabled, DHCP is used to assign them an IP address. The clients then use the IP address to communicate with a PXE server to start the operating system installation. You must have a working DHCP server with an active scope on the network for WDS to work properly.

You can install a WDS distribution point (DP) on a DHCP server, but special attention is needed to configure the User Datagram Protocol (UDP) ports, because both services use UDP port 67. The WDS server needs the following ports opened:

- UDP Port 67 – DHCP
- UDP Port 69 – TFTP
- UDP Port 4001 – PXE

You can avoid this issue if you deploy the WDS DP on a different server from the DHCP server. If this is not possible, you must configure the WDS service to listen on a different port by performing the following steps:

1. Modify the following registry key:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\WDS\Providers\WDSPXE
```

2. Set the registry value to:

```
UseDHCPPorts = 0
```

3. To set the new configuration, run the following command on the WDS server:

```
WDSUTIL /Set-Server /UseDHCPPorts:No /DHCPOption60:Yes
```

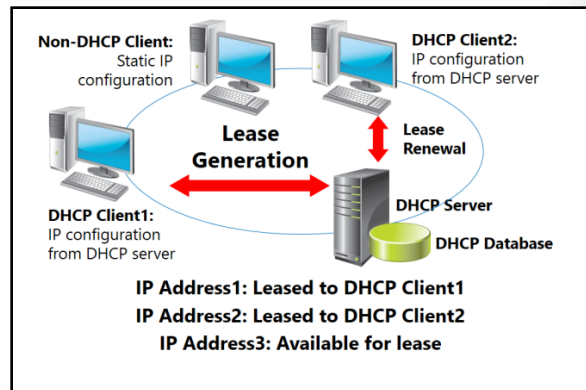
How DHCP Allocates IP Addresses

DHCP allocates IP addresses on a dynamic basis, otherwise known as a *lease*. Although you can set the lease duration to Unlimited, you typically set the duration for not more than a few hours or days. The default lease time for wired clients is eight days, and for wireless clients it is three days.

DHCP uses IP broadcasts to initiate communications. Therefore, DHCP servers are limited to communication within their IP subnet. This means that in many networks, there is a DHCP server for each IP subnet.

By default, all Microsoft operating systems are configured to obtain an IP address automatically. For a computer to be considered a DHCP client, it must be configured to obtain an IP address automatically. In a network where a DHCP server is installed, DHCP clients respond to DHCP broadcasts.

If a computer is configured with an IP address by an administrator, then that computer has a static IP address and is considered a non-DHCP client, and does not communicate with a DHCP server.

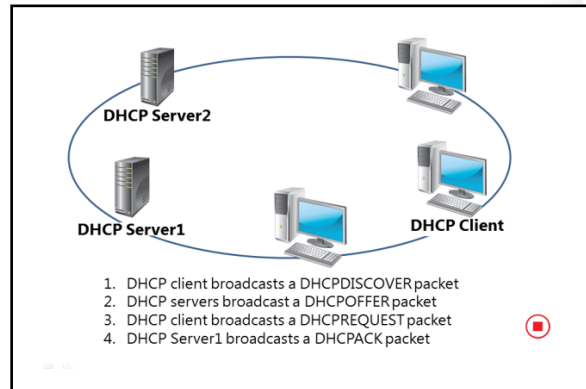


How DHCP Lease Generation Works

DHCP uses a four-step lease-generation process to assign an IP address to clients. Understanding how each step of this process works helps you troubleshoot problems when clients cannot obtain an IP address.

The following are the four steps of the DHCP lease-generation process:

1. The DHCP client broadcasts a DHCPDISCOVER packet to every computer in the subnet. The only computers that respond are computers that have either the DHCP server role or computers or routers that are running a DHCP relay agent. In the latter case, the DHCP relay agent forwards the message to the DHCP server with which it is configured.
2. A DHCP Server responds with a DHCPOFFER packet. This packet contains a potential address for the client.
3. The client receives the DHCPOFFER packet. It might receive packets from multiple servers. If it does, it usually selects the server that made the fastest response to its DHCPDISCOVER, which typically is the DHCP server closest to the client. The client then broadcasts a DHCPREQUEST that contains a server identifier. This informs the DHCP servers that receive the broadcast which server's DHCPOFFER the client has chosen to accept.



4. The DHCP servers receive the DHCPREQUEST. Servers that the client has not accepted use this message as the notification that the client declines that server's offer. The chosen server stores the IP address client information in the DHCP database and responds with a DHCPACK message. If the DHCP server cannot provide the address that was offered in the initial DHCPOFFER, the DHCP server sends a DHCPNAK message.



Additional Reading: For more information about DHCP technology in Windows Server 2012, see Dynamic Host Configuration Protocol (DHCP) Overview at <http://go.microsoft.com/fwlink/?LinkId=269709>.

How DHCP Lease Renewal Works

When the DHCP lease reaches 50 percent of the lease time, the client automatically attempts to renew the lease. This process occurs in the background. It is possible for a computer to have the same DHCP-assigned IP address for a long time if the computer is not restarted, as it will renegotiate the lease periodically.

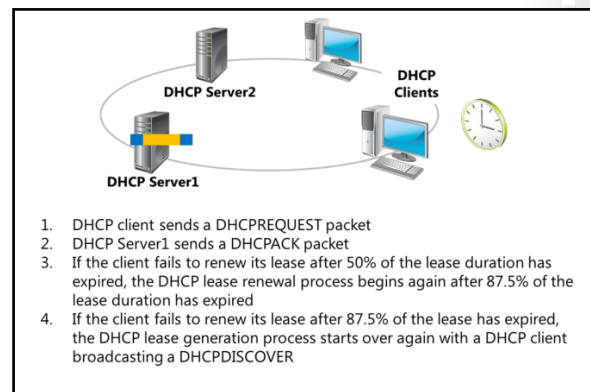
To attempt to renew the IP address lease, the client sends a unicast DHCPREQUEST message. The server that leased the IP address originally sends a DHCPACK message back to the client. This message contains any new parameters that have changed since the original lease was created. Note that these packets are not broadcast, because the client at this point has an IP address it can use for unicast communications.

If the DHCP client cannot contact the DHCP server, then the client waits until 87.5 percent of the lease time expires. If the renewal is unsuccessful, or in other words 100 percent of the lease time has expired, then the client computer attempts to contact the configured default gateway. If the gateway does not respond, the client considers itself to be on a new subnet and enters the Discovery phase, where it attempts to obtain an IP configuration from any DHCP server, as previously described.

Because client computers might be moved while they are turned off, for example a laptop computer that is plugged into a new subnet, client computers also attempt renewal during the startup process, or when the computer detects a network change. If renewal is successful, the lease period is reset.

DHCP Server Failover Protocol

The DHCP role on Windows Server 2012 supports a new feature named the *DHCP Server Failover protocol*. This protocol enables synchronization of lease information between multiple DHCP servers. It also increases DHCP service availability. If one DHCP server is not available, the other DHCP servers continue to service clients in the same subnet.



How DHCP Interacts with DNS

DHCP servers are primarily used to give client computers IP addresses dynamically. DNS servers are mainly used to find an IP address based on the given name or to find a name based on a given IP address. Starting with Windows 2000, DNS clients can register their records through the DNS dynamic update protocol.

Additionally, you can configure a DHCP server to register and update client names and IP addresses with a DNS server when those DHCP clients belong to that DNS zone. DHCP option code 81 returns a client's Fully Qualified Domain Name (FQDN) to the DHCP server, which can then dynamically update the individual client's resource record back to the DNS server by using the DNS dynamic update protocol.

- DHCP can:
 - Register client records into DNS zones
 - Use DNS dynamic update protocol
- To use secure DNS dynamic updates, add DHCP servers to the AD DS DnsUpdateProxy global group
- DHCP policies:
 - Automatically assign settings based on FQDN
 - Register workgroup computers with guest DNS suffix
 - Disable PTR registrations without disabling host record registration

DNS Dynamic Update Protocol

Depending on how you configure the DNS dynamic update function on the DNS server, using the DNS dynamic update protocol might not be secure. Instead, you can configure the secure DNS dynamic update functionality. The DNS server accepts updates only from clients that are authorized to make DNS dynamic updates to the objects they represent in AD DS. When using DHCP servers and DNS servers that are set for secure DNS dynamic updates, you can add the DHCP server's computer account to the AD DS DnsUpdateProxy global group. Membership in this group ensures that the DHCP server can perform secure DNS dynamic updates for a client's resource records.

DHCP Policies

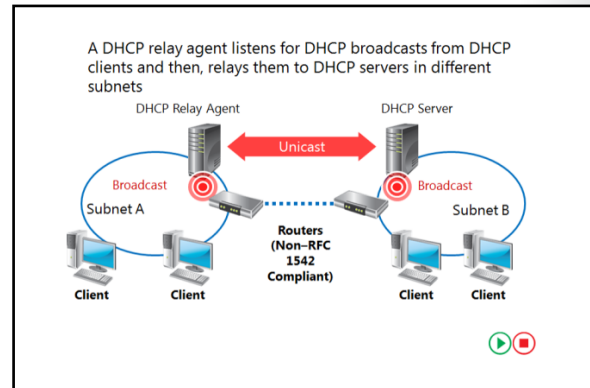
You can create DHCP policies in Windows Server 2012. Policy-based assignment allows the DHCP server to evaluate requests for IP addresses against policies that you define. The policies apply to a specific scope using a defined processing order and can be inherited from the server. When the request matches the conditions of a policy the DHCP server provides specific settings to the client. You can use DHCP policies to configure conditions based on the FQDN of the clients, and to register workgroup computers with a guest DNS suffix.

Previous to Windows Server 2012 R2, if you wanted to prevent a DNS reverse lookup record in DHCP, also known as pointer records registration (PTR), you had to disable both host and PTR record registration for DHCP clients. In Windows Server 2012 R2, you can allow a DHCP server to register a client's host record, but not the PTR record.

What Is a DHCP Relay Agent?

When initially attempting to get an IP address, DHCP clients use IP broadcasts to initiate communications. Because of this, DHCP servers and clients can only communicate within their IP subnet. This means that in many networks, there is a DHCP server for each IP subnet. If there are a large number of subnets, it might be expensive to deploy servers for every subnet. A single DHCP server might service collections of smaller subnets.

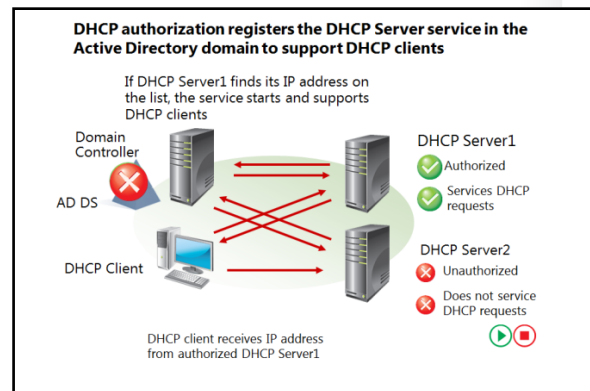
For the DHCP server to respond to a DHCP client request, it must be able to receive DHCP requests. You can enable this by configuring a DHCP relay agent on each subnet. A *DHCP relay agent* is a computer or router that listens for DHCP broadcasts from DHCP clients and then relays them to DHCP servers in different subnets.



With the DHCP relay agent, the DHCP broadcast packets can be relayed into another IP subnet across a router. Then, you can configure the DHCP relay agent in the subnet that requires IP addresses. Additionally, you can configure the agent with the IP address of the DHCP server. The agent can then capture the client broadcasts and forward them to the DHCP server in another subnet. You also can relay DHCP packets into other subnets using a router that is compatible with RFC 1542.

DHCP Server Authorization

DHCP allows a client computer to acquire configuration information about the network in which it starts. DHCP communication typically occurs before any authentication of the user or computer; and because the DHCP protocol is based on IP broadcasts, an incorrectly-configured DHCP server in a network can provide invalid information to clients. To avoid this, the server must be authorized. *DHCP authorization* is the process of registering the DHCP Server service in the Active Directory domain to support DHCP clients.



Active Directory Requirements

You must authorize the Windows Server 2012 DHCP server role in AD DS before it can begin leasing IP addresses. It is possible to have a single DHCP server providing IP addresses for subnets that contain multiple AD DS domains. Because of this, an Enterprise Administrator account must authorize the DHCP server.



Note: For authorization purposes, you must be an Enterprise Administrator in all domains, with the exception of the forest root domain. In the forest root domain, members of the Domain Admins group belong to the Enterprise Administrator group, which has adequate privilege to authorize a DHCP server.

Standalone DHCP Server Considerations

A standalone DHCP server is a computer that is running Windows Server 2012, that is not part of an AD DS domain, and that has the DHCP server role installed and configured. If the standalone DHCP server detects an authorized DHCP server in the domain, it does not lease IP addresses and then automatically shuts down.

Unauthorized DHCP Servers

Many network devices have built-in DHCP server software. As such, many routers can act as a DHCP server, but often these servers do not recognize DHCP-authorized servers, and might lease IP addresses to clients. In this situation you must perform an investigation to detect unauthorized DHCP servers, whether they are installed on devices or on non-Microsoft servers. Once you detect unauthorized DHCP servers, you should disable the DHCP service or functionality on them. You can find the IP address of the DHCP server by issuing the **ipconfig /all** command on the DHCP client computer.

Demonstration: Adding the DHCP Server Role

In this demonstration, you will see how to install and authorize the DHCP server role.

Demonstration Steps

Install the DHCP server role

1. Sign in to **LON-SVR1** as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Open **Server Manager** and install the **DHCP Server** role.
3. In the Add Role Wizard, accept all the default settings.
4. Close **Server Manager**.
5. Repeat steps 1 through 4 on LON-SVR2.

Authorize the DHCP Server

1. Switch to **LON-SVR1**.
2. Open the DHCP console.
3. Authorize the **lon-svr1.adatum.com** server in AD DS.
4. Repeat steps 1 through 3 on LON-SVR2, replacing the FQDN in step 3 as **LON-SVR2**.



Note: Leave all virtual machines in their current state for the next demonstration.

Lesson 2

Configuring DHCP Scopes

After you install the DHCP role on a server, you must configure the DHCP scopes. A DHCP scope is the primary method you can use to configure options for a group of IP addresses. A DHCP scope is based on an IP subnet, and can have settings specific to hardware or custom groups of clients. This lesson explains DHCP scopes and how to manage them.

Lesson Objectives

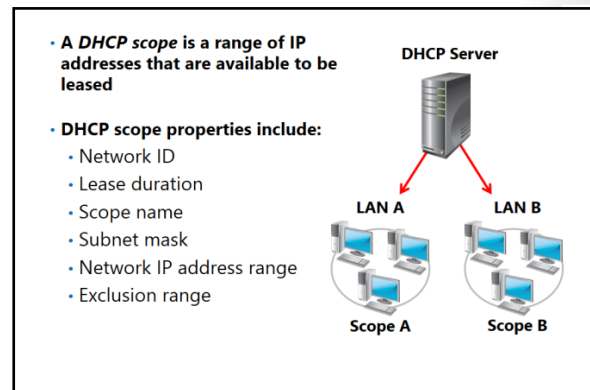
After completing this lesson, you will be able to:

- Describe the purpose of a DHCP scope.
- Describe a DHCP reservation.
- Describe the DHCP Options.
- Explain how to apply DHCP Options.
- Create and configure a DHCP scope.

What Are DHCP Scopes?

A *DHCP scope* is a range of IP addresses that are available for lease, and that are managed by a DHCP server. A DHCP scope typically is confined to the IP addresses in a given subnet.

For example, a DHCP scope for the network 192.168.1.0/24 (subnet mask of 255.255.255.0) supports a range from 192.168.1.1 through 192.168.1.254. When a computer or device in the 192.168.1.0/24 subnet requests an IP address, the scope that defined the range in this example allocates an address between 192.168.1.1 and 192.168.1.254.



Note: Remember that the DHCP server, if deployed to the same subnet, consumes an IPv4 address. This address should be excluded from the IPv4 address range.

To configure a scope, you must define the following properties:

- Name and description. This property identifies the scope.
- IP address range. This property lists the range of addresses that can be offered for lease, and usually lists the entire range of addresses for a given subnet.
- Subnet mask. This property is used by client computers to determine their location in the organization's network infrastructure.
- Exclusions. This property lists single addresses or blocks of addresses that fall within the IP address range, but that will not be offered for lease.
- Delay. This property is the amount of time to delay before making DHCPOFFER.

- Lease duration. This property lists the lease duration. Use shorter durations for scopes with limited IP addresses, and longer durations for more static networks.
- Options. You can configure many optional properties on a scope, but typically you configure:
 - Option 003 – Router (the default gateway for the subnet)
 - Option 006 – Domain Name System (DNS) servers
 - Option 015 – DNS suffix

IPv6 Scopes

You can configure the IPv6 scope options as a separate scope in the DHCP console's IPv6 node. The IPv6 node contains several different options that you can modify, and an enhanced lease mechanism.

When configuring a DHCPv6 scope, you must define the following properties:

- Name and description. This property identifies the scope.
- Prefix. The IPv6 address prefix is analogous to the IPv4 address range; in essence, it defines the network address.
- Exclusions. This property lists single addresses or blocks of addresses that fall within the IPv6 prefix but will not be offered for lease.
- Preferred lifetimes. This property defines how long leased addresses are valid.
- Options. Like IPv4, you can configure many options.


Windows PowerShell

In Windows Server 2012, Microsoft introduced several new Windows PowerShell cmdlets to configure and manage DHCP servers. Each cmdlet has parameters that need to be met, depending on actions to be taken. Many of the new cmdlets addressed scope creation and management, as shown in the following table.

Cmdlet name	Description
Add-DhcpServerv4Scope	Adds an IPv4 scope on the Dynamic Host Configuration Protocol (DHCP) server service.
Add-DhcpServerv6Scope	Adds an IPv6 scope to the Dynamic Host Configuration Protocol (DHCP) server service with the specified parameters.
Get-DhcpServerv4Scope	Returns the IPv4 scope configuration of the specified scopes.
Get-DhcpServerv4ScopeStatistics	Gets the IPv4 scope statistics corresponding to the IPv4 scope identifiers (IDs) specified for a Dynamic Host Configuration Protocol (DHCP) server service.
Get-DhcpServerv6Scope	Gets the scope information for the specified IPv6 prefixes on the Dynamic Host Configuration Protocol (DHCP) server service.
Get-DhcpServerv6ScopeStatistics	Gets the IPv6 prefix statistics that correspond to the IPv6 prefix specified for a Dynamic Host Configuration Protocol (DHCP) server service.
Remove-DhcpServerv4Scope	Deletes the specified IPv4 scopes from the Dynamic Host Configuration Protocol (DHCP) server service.

(continued from previous page)

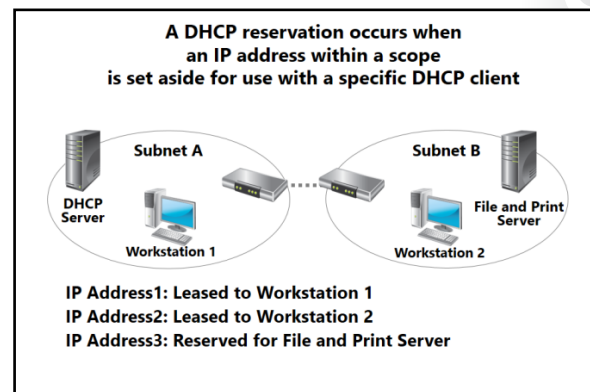
Cmdlet name	Description
Remove-DhcpServerv6Scope	Deletes the IPv6 Scopes from the Dynamic Host Configuration Protocol (DHCP) server service corresponding to the specified prefixes.
Set-DhcpServerv4Scope	Sets the properties of an existing IPv4 scope on the Dynamic Host Configuration Protocol (DHCP) server service.
Set-DhcpServerv6Scope	Modifies the properties of the IPv6 scope on the Dynamic Host Configuration Protocol (DHCP) server service.

 **Additional Reading:** For more information about DHCP Server cmdlets in Windows PowerShell, go to <http://go.microsoft.com/fwlink/?LinkID=331064>.

 **Additional Reading:** For additional Windows PowerShell cmdlets for DHCP added in Windows Server 2012 R2, go to <http://go.microsoft.com/fwlink/?LinkID=331065>.

What Is a DHCP Reservation?

As a best practice, you should consider providing network devices—such as network printers—with a predetermined IP address. Using a DHCP reservation, you can ensure that the IP addresses that you set aside from a configured scope are not assigned to another device. A DHCP reservation is a specific IP address from within a scope that is reserved permanently for lease to a specific DHCP client. A DHCP reservation also ensures that devices with reservations are guaranteed an IP address even if a scope is depleted of addresses. Configuring reservations enables you to centralize management of fixed IP addresses.



Configuring DHCP Reservations

To configure a reservation, you must know the device's network interface media access control (MAC) address or physical address. This address indicates to the DHCP server that the device should have a reservation. You can acquire a network interface's MAC address by using the **ipconfig /all** command. Typically, MAC addresses for network printers and other network devices are printed on the device. Most laptop computers also note this information on the bottom of their chassis.

What Are DHCP Options?

DHCP servers can configure more than just an IP address, they also provide information about network resources, such as DNS servers and the default gateway. DHCP options are values for common configuration data that apply to the server, scopes, reservations, and class options. You can apply DHCP options at the server, scope, user, and vendor levels. An option code identifies the DHCP options, and most option codes come from the RFC documentation found on the Internet Engineering Task Force (IETF) website.

DHCP options:

- Are values for common configuration data
- Apply to the server, scopes, reservations, and class options

Common scope options are:

- Router (Default Gateway)
- DNS Name
- DNS Servers
- WINS Servers

Common DHCP Options

The following table lists the common option codes that Windows-based DHCP clients request.

Option code	Name
1	Subnet mask
3	Router
6	DNS servers
15	DNS domain name
31	Perform router discovery
33	Static route
43	Vendor-specific information
47	NetBIOS scope ID
51	Lease time
58	Renewal (T1) time value
59	Rebinding (T2) time value
60	Preboot Execution (PXE) client
66	Boot server host name
67	Bootfile name
249	Classless static routes

PXE Boot options

Preboot Execution (PXE)-enabled network cards add the DHCP option 60 to their discover packets. Normally, DHCP clients send a DHCP option 67 packet and then DHCP servers return a DHCP 68 option offer. Because the ports used for DHCP are also used by the Windows Deployment Services (WDS) PXE server function, if you have DHCP and a PXE server deployed on the same machine, you must set DHCP to

make offers that also include the 60 option. A DHCP server then makes the DHCP 60 offer back to the client. You need to set DHCP Options 60 (PXE Client), 66 (Boot Server Host Name), and 67 (Bootfile Name). You can set options 66 and 67 in the Scope Options window in the DHCP console, but you must set the 60 option via the command line.

The following code sample lists the procedure:

```
C:\WINDOWS\system32>netsh
netsh>dhcp
netsh dhcp>server \\<server_machine_name>
netsh dhcp>add optiondef 60 PXEClient String 0 comment=PXE support
netsh dhcp>set optionvalue 60 STRING PXEClient
netsh dhcp>exit
```

After this code has run, a PXE server then sends back boot server and boot information to the PXE-enabled network client, which allows it to accept an operating system installation.

How Are DHCP Options Applied?

DHCP applies options to client computers. You need to understand these options when configuring DHCP, so you will know which level settings has priority when you are configuring different settings on multiple levels.

DHCP applies options in the following order:

1. Server level. A server-level option is assigned to all DHCP clients of the DHCP server.
2. Scope level. A scope-level option is assigned to all clients of a scope. Scope options override server options.
3. Class level. A class-level option is assigned to all clients that identify themselves as members of a class. Class options override both scope and server options.
4. Reserved client level. A reservation-level option is assigned to one DHCP client. Reserved client options apply to devices that have a DHCP reservation.

You can apply DHCP options at various levels:

- Server
- Scope
- Class
- Reserved client

Typically, you do not apply the class or reserved client options

If DHCP option settings are applied at each level and they conflict, then the option that is applied last overrides the previously-applied setting. For example, if the default gateway is configured at the scope level, and a different default gateway is applied for a reserved client, then the reserved client setting becomes the effective setting.

You can also configure address assignment policies at the server level or at the scope level. Address assignment policy contains a set of conditions that you define to lease different DHCP IP addresses and settings to different types of DHCP clients, such as computers, laptops, network printers, or IP phones. The conditions defined in these policies differentiate various types of clients, and include multiple criteria, such as MAC address or vendor information.

Demonstration: Creating and Configuring a DHCP Scope

You can create scopes using either the Microsoft Management Console (MMC) for the DHCP server role, or the Netsh network configuration command-line tool. The Netsh command-line tool allows you to manage scopes remotely if the DHCP server is running on a Server Core installation of Windows Server 2012. The Netsh command-line tool is also useful for scripting and automating server provisioning.

In Windows Server 2012, Microsoft introduced several new Windows PowerShell cmdlets to configure and manage DHCP servers.

In this demonstration, you will see how to configure scope and scope options by using both the DHCP console and the new Windows PowerShell cmdlets.

Demonstration Steps

Configure scope and scope options in DHCP

1. In DHCP, in the navigation pane, expand **lon-svr1.adatum.com**, expand and right-click **IPv4**, and then click **New Scope**.
2. Create a new scope with the following properties:
 - o Name: Branch Office
 - o IP Address Range: 172.16.0.100-172.16.0.200
 - o Length: **16**
 - o Subnet Mask: **255.255.0.0**
 - o Exclusions: 172.16.0.190-172.16.0.200
 - o Other settings: use default values
 - o Configure options Router 172.16.0.1
3. Use default settings for all other pages, and then activate the scope.

Configure scope and scope options in DHCP with Windows PowerShell

1. In Windows PowerShell, type the following cmdlets:

```
Add-DhcpServerv4Scope -Name "Branch Office 2" -StartRange 10.10.0.100 -EndRange 10.10.0.200 -SubnetMask 255.255.0.0
Add-DhcpServerv4ExclusionRange -ScopeID 10.10.0.0 -StartRange 10.10.0.190 -EndRange 10.10.0.200
Set-DhcpServerv4OptionValue -Router 10.10.0.1
Set-DhcpServerv4Scope -ScopeID 10.10.0.0 -State Active
```

2. In the **DHCP Manager** examine the scope just created.

Lesson 3

Managing a DHCP Database

The DHCP database stores information about the IP address leases. If there is a problem, it is important that you understand how to back up the database and resolve database issues. This lesson explains how to manage the database and its data.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the DHCP database.
- Explain how to back up and restore a DHCP database.
- Explain how to reconcile a DHCP database.
- Explain how to move a DHCP database.

What Is a DHCP Database?

The DHCP database is a dynamic database containing data that relates to scopes, address leases, and reservations. The database also contains the data file that stores both the DHCP configuration information and the lease data for clients that have leased an IP address from the DHCP server. By default, the DHCP database files are stored in the %systemroot%\System32\Dhcp folder.

DHCP Service Database Files

The following table describes some of the DHCP service database files.

The *DHCP database* is a dynamic database that contains configuration information such as:

- Scopes
- Reservations
- Address leases

Windows Server 2012 stores the DHCP database in the %Systemroot%\System32\Dhcp folder

The DHCP database files include:

- Dhcp.mdb
- Res*.log
- Dhcp.tmp
- J50.chk
- J50.log and J50*.log

File	Description
Dhcp.mdb	Dhcp.mdb is the DHCP server database file.
Dhcp.tmp	Dhcp.tmp is a temporary file that the DHCP database uses as a swap file during database index maintenance operations. Following a system failure, Dhcp.tmp sometimes remains in the <i>Systemroot\System32\Dhcp</i> directory.
J50.log and J50#####.log	J50.log and J50#####.log are logs of all database transactions. The DHCP database uses this log to recover data when necessary.
J50.chk	This is a checkpoint file.



Note: You should not remove or alter any of the DHCP service database files.

The DHCP server database is dynamic. It updates as DHCP clients are assigned, or as they release their TCP/IP configuration parameters. Because the DHCP database is not a distributed database like the

Windows Internet Name Service (WINS) server database, maintaining the DHCP server database is less complex.

By default, the DHCP database and related registry entries are backed up automatically at 60-minute intervals. You can change this default interval by changing the value of BackupInterval in the following registry key:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DHCP\Parameters
```


You can also back up a DHCP database manually at any time.

Backing Up and Restoring a DHCP Database

You can back up a DHCP database manually, or you can configure it to back up automatically. An automatic backup is called a synchronous backup. A manual backup is called an asynchronous backup.

Automatic (Synchronous) Backup

The default backup path for the DHCP backup is systemroot\System32\Dhcp\Backup. As a best practice, you can modify this path in the server properties to point to another volume.



Manual (Asynchronous) Backup

If you have an immediate need to create a backup, you can run the manual backup option in the DHCP console. This action requires either administrative-level permissions, or that the user account be a member of the DHCP administrators group.

What Is Backed Up?

When a synchronous or asynchronous backup occurs, the entire DHCP database is saved, including the following:

- All scopes
- Reservations
- Leases
- All options, including server options, scope options, reservation options, and class options
- All registry keys and other configuration settings (for example, audit log settings and folder location settings) that are set in DHCP server properties. These settings are stored in the following registry key:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DHCP\Parameters
```

To back up this key, open Registry Editor and save the specified key to a text file.



Note: The DNS dynamic update credentials (user name, domain, and password) that the DHCP server uses when registering DHCP client computers in DNS are not backed up with any backup method.

Restoring a Database

If you need to restore the database, use the Restore function in the DHCP server console. You will be prompted for the backup's location. Once you have selected the location, DHCP service stops, and the database is restored. To restore the database, the user account must either have administrative-level permissions, or be a member of the DHCP administrators group.

Backup Security

When the DHCP database file is backed up, it should be in a protected location that only the DHCP administrators can access. This ensures that any network information in the backup files remains protected.

Using Netsh

You also can use commands in the Netsh DHCP Server context to back up the database; this is useful for backing up the database to a remote location using a script file.

The following command is a script that you can use from the Netsh DHCP Server prompt to back up the DHCP data for all scopes:

```
export "c:\My Folder\Dhcp Configuration" all
```

To restore the DHCP database, use the following command:

```
import "c:\My Folder\Dhcp Configuration" all
```



Note: The Netsh DHCP Server context does not exist on server computers that do not have the DHCP server role installed.

Using Windows PowerShell

In Windows Server 2012, Microsoft introduced several new Windows PowerShell cmdlets you can use to configure and manage DHCP servers.

To back up the DHCP data for all scopes, use the following command:

```
Backup-DhcpServer -ComputerName 1on-svr1.adatum.com -Path C:\Windows\system32\dhcp\backup
```

To restore the DHCP database, use the following command:

```
Restore-DhcpServer -ComputerName 1on-svr1.adatum.com -Path  
C:\Windows\system32\dhcp\backup
```

The export operation exports the DHCP server service configuration and lease data, to a specified file.

To export, use the following command:

```
Export-DhcpServer -ComputerName 1on-svr1.adatum.com -File C:\exportdir\dhcpxport.xml
```

To import, use the following command:

```
Import-DhcpServer -ComputerName 1on-svr2.adatum.com -File C:\exports\dhcpxport.xml  
-BackupPath C:\dhcpbackup\
```

Reconciling a DHCP Database

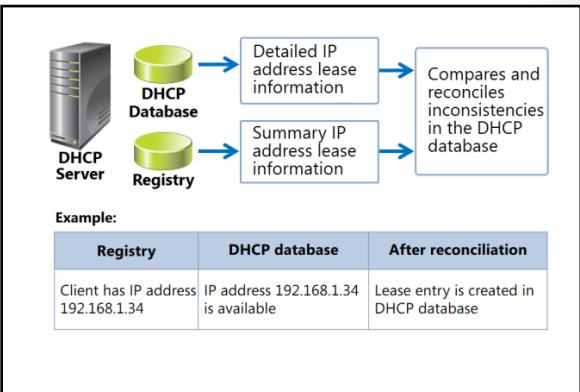
Reconciling scopes can fix inconsistencies that can affect client computers.

The DHCP Server service stores scope IP address-lease information in two forms:

- Detailed IP address lease information, which the DHCP database stores
- Summary IP address lease information, which the server's Registry stores

When you are reconciling scopes, the detail and summary entries are compared to find inconsistencies.

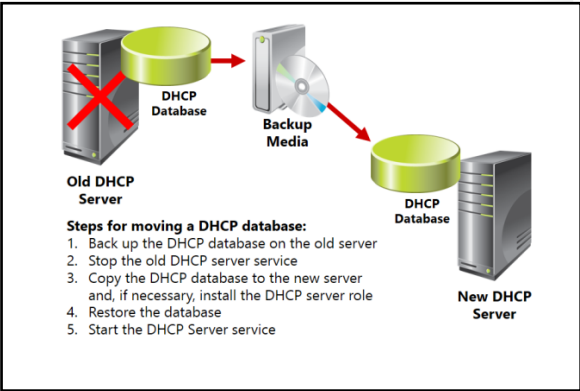
To correct and repair these inconsistencies, you must reconcile any scope inconsistencies. After you select and reconcile scope inconsistencies, the DHCP service either restores those IP addresses to the original owner, or creates a temporary reservation for those addresses. These reservations are valid for the lease time that is assigned to the scope. When the lease time expires, the addresses are then recovered for future use.



Moving a DHCP Database

In the event that you must move the DHCP server role to another server, as a best practice you should also move the DHCP database to the same server. This ensures that client leases are retained, and reduces the likelihood of client-configuration issues.

Initially, move the database by backing it up on to the old DHCP server. Then, shut down the DHCP service on the old DHCP server. Next, copy the DHCP database to the new server, where you can restore it using the normal database restore procedure.



Lesson 4

Securing and Monitoring DHCP

DHCP protocol has no built-in method for authenticating users. This means that if you do not take precautions, IP leases could be granted to devices and users who are unauthorized.

DHCP is a core service in many organization's network environments. If the DHCP service is not working properly, or if there is a situation that is causing problems with the DHCP server, it is important that you can identify the problem and determine potential causes to resolve the problem.

This lesson explains how to prevent unauthorized users from obtaining a lease, how to manage unauthorized DHCP servers, and how to configure DHCP servers so that a specific group can manage them.

Lesson Objectives

After completing this lesson, you will be able to:

- Explain how to prevent an unauthorized computer from obtaining a lease.
- Explain how to restrict unauthorized, non-Microsoft DHCP servers from leasing IP addresses.
- Explain how to delegate administration of the DHCP server role.
- Describe DHCP statistics.
- Describe DHCP audit logging.
- Identify common issues that are possible with DHCP.

Preventing an Unauthorized Computer from Obtaining a Lease

DHCP by itself can be difficult to secure—it is designed to work before the necessary information is in place for a client computer to authenticate with a domain controller. This is why you should take precautions to prevent unauthorized computers from obtaining a lease with DHCP.

Basic precautions that you should take to limit unauthorized access include:

- Ensuring that you reduce physical access. If users can access an active network connection to your network, their computers are likely to be able to obtain an IP address. If a network port is not being used, you should disconnect it physically from the switching infrastructure.
- Enabling audit logging on all DHCP servers. This can provide an historical view of activity, in addition to allowing you to trace when an unauthorized user obtained an IP address in the network. Make sure to review the audit logs at regular intervals.
- Requiring authenticated Layer 2 connections to the network: Most enterprise hardware switches now support Institute of Electrical and Electronics Engineers, Inc. (IEEE) 802.1X authentication. This allows for port-level user authentication. Secure wireless standards, such as Wi-Fi Protected Access (WPA) Enterprise and WPA2 Enterprise, also use 802.1X authentication.

To prevent an unauthorized computer from obtaining a lease:

- Ensure that unauthorized users do not have physical or wireless access to your network
- Enable audit logging for every DHCP server on your network
- Regularly check and monitor audit log files
- Use 802.1X-enabled LAN switches or wireless access points to access the network
- Configure NAP to validate that a client computer is compliant with system health requirements

- Implementing NAP. NAP enables administrators to validate that a client computer is compliant with system health requirements, such as running all the latest Windows operating system updates, or running an up-to-date antivirus client. If users who do not meet security requirements try to access the network, they receive an IP address configuration to access a remediation network where they can receive the necessary updates. The administrator can restrict access to the network by allowing only healthy computers access to the internal local area network (LAN).

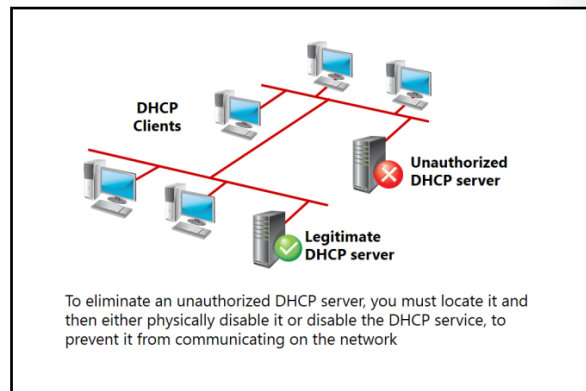
Restricting Unauthorized, Non-Microsoft DHCP Servers from Leasing IP Addresses

Many devices and network operating systems have multiple DHCP server implementations. Networks are almost never homogeneous in nature, so it is possible that at some point a DHCP server that does not check for Active Directory–authenticated servers will be enabled on the network. In this case, clients might obtain incorrect configuration data.

To eliminate an unauthorized DHCP server, you must first locate it. You must then prevent it from communicating on the network by disabling it physically, or by disabling the DHCP service.

If users complain that they do not have connectivity to the network, check the IP address of their DHCP server. Use the **ipconfig /all** command to check the IP address of the DHCP Server field. If the IP address is not the IP address of an authorized DHCP server, then there is probably an unauthorized server in the network.

You can use the DHCP Server Locator utility (Dhcploc.exe) to locate the DHCP servers that are active on a subnet. The DHCP Server Locator utility is found on the Windows Server 2008 installation DVD (\Support\Tools folder) or can be downloaded from the TechNet Gallery.



Delegating DHCP Administration

Ensure that only authorized persons can administer the DHCP server role. You can do this by performing either of the following tasks:

- Limit the membership of the DHCP Administrators group
- Assign users that require read-only access to DHCP membership of the DHCP Users group

The DHCP Administrators local group is used to restrict and grant access to only administering DHCP servers. The DHCP Administrators group is automatically created in AD DS when the DHCP server role is installed on a domain controller. It also is automatically created on a local computer when the DHCP server role is installed on domain members or on standalone servers. The groups have no members by default. Adding accounts to the membership of either group allows those accounts to administer the DHCP server.

To delegate who can administer the DHCP service:

- Limit the membership of the DHCP Administrators group
- Add users to the DHCP Users group if they need read-only access to the DHCP console

Account	Permissions
DHCP Administrators group	Can view and modify any data about the DHCP server
DHCP Users group	Has read-only DHCP console access to the server

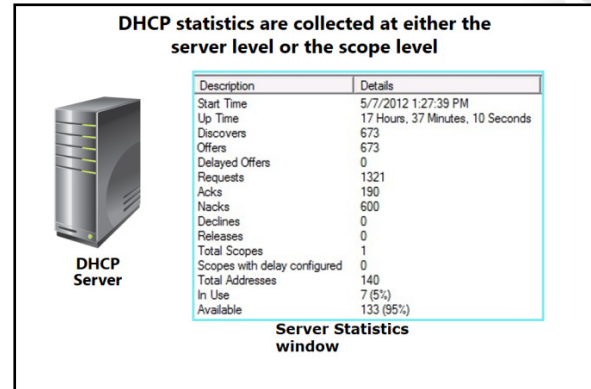
Permissions Required to Authorize and Administer DHCP

Only Enterprise administrators can authorize a DHCP service. If an administrator with lower credentials than an Enterprise administrator needs to authorize the domain the administrator should use Active Directory delegation. Any user in the DHCP Administrators group can manage the server's DHCP service. Any user in the DHCP Users group can have read-only access to the DHCP console.

What Are DHCP Statistics?

DHCP statistics provide information about DHCP activity and use. You can use this console to determine quickly whether there is a problem with the DHCP service or with the network's DHCP clients. An example in which statistics might be useful is if you notice an excessive amount of negative acknowledgement (NAK) packets, which might indicate that the server is not providing the correct data to clients.

You can configure the refresh rate for the statistics in the General tab of server's Properties dialog box.



DHCP Server Statistics

DHCP server statistics provide an overview of DHCP server usage. You can use this data to quickly understand the state of the DHCP server. Information such as number of offers, number of requests, total in-use addresses, and total available addresses can help to provide a picture of the server's health.

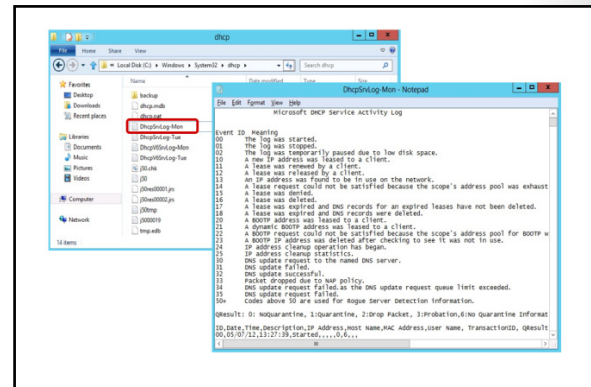
DHCP Scope Statistics

DHCP scope statistics provide far fewer details—such as total addresses in the scope, how many addresses are in use, and how many addresses are available. If you notice that there are a low number of addresses available in the server statistics, it might be that only one scope is near its depletion point. By using scope statistics, an administrator can quickly determine the status of the particular scope with respect to the addresses available.

What Is DHCP Audit Logging?

The DHCP audit log is a log that provides a traceable log of DHCP server activity. You can use this log to track lease requests, grants, and denials. This information allows you to troubleshoot DHCP server performance. The log files are stored in the %systemroot%\system32\dhcp folder by default. You can configure the log file settings in the server's Properties dialog box.

The DHCP audit log files are named based on the weekday that the file was created. For example, if audit logging is enabled on a Monday, the file name is DhcpSrvLog-Mon.log.



DHCP Audit Log Fields

The following table describes the fields in a DHCP audit log.

Field	Description
ID	A DHCP server event ID code
Date	The date on which the entry was logged on the DHCP server
Time	The time at which the entry was logged on the DHCP server
Description	A description of the DHCP server event
IP Address	The IP address of the DHCP client
Host Name	The host name of the DHCP client
MAC Address	The MAC address used by the client's network adapter hardware

Common Event ID Codes

Common event ID codes are written as follows:

- ID,Date,Time,Description,IP Address,Host Name,MAC Address

Common event ID codes include:

- 00,06/22/99,22:35:10,Started,,,,
- 56,06/22/99,22:35:10,Authorization failure, stopped servicing,,domain1.local,,
- 55,06/22/99,22:45:38,Authorized(servicing),,domain1.local

Discussion: Common DHCP Issues

The following table describes some common DHCP issues. Enter the possible solutions in the Solution column, and then discuss your answers with the class.

Common issues that can occur when you do not configure DHCP properly:

- Address conflicts
- Failure to obtain a DHCP address
- Address obtained from an incorrect scope
- DHCP database suffered data corruption or loss
- DHCP server has exhausted its IP address pool



10 minutes

Issue	Description	Example	Solution
Address conflicts	The same IP address is offered to two different clients.	An administrator deletes a lease. However, the client that had the lease is still operating as if the lease is valid. If the DHCP server does not verify the IP address, it might lease the IP address to another machine, causing an address conflict. This can also occur if two DHCP servers have overlapping scopes.	
Failure to obtain a DHCP address	The client does not receive a DHCP address and instead receives an Automatic Private IP Addressing (APIPA) self-assigned address.	If a client's network card driver is configured incorrectly, it might cause a failure to obtain a DHCP address. Additionally, the DHCP server or relay agent on the client's subnet might be not online. Another reason might be that the DHCP server has exhausted its scope, so the scope should be extended or modified.	
Address obtained from an incorrect scope	The client is obtaining an IP address from the wrong scope, causing it to experience communication problems.	If the client is connected to the wrong network, or if the DHCP relay agent is configured incorrectly, this error could occur.	
DHCP database suffers data corruption or loss	The DHCP database becomes unreadable or is lost due to a hardware failure.	A hardware failure can cause the database to become corrupted.	
DHCP server exhausts its IP address pool	The DHCP server's IP scopes have been depleted. Any new clients requesting an IP address are refused.	If all of the IP addresses that are assigned to a scope are leased, this error occurs.	

Lab: Implementing DHCP

Scenario

A. Datum Corporation has an IT office and data center in London, which supports the London location and other locations as well. A. Datum has recently deployed a Windows 2012 Server infrastructure with Windows 8® clients.

You have recently accepted a promotion to the server support team. One of your first assignments is to configure the infrastructure service for a new branch office. As part of this assignment, you need to configure a DHCP server that will provide IP addresses and configuration to client computers. Servers are configured with static IP addresses and do not use DHCP.

Objectives

After completing this lab, you should be able to:

- Implement DHCP.
- Implement a DHCP relay agent (optional).

Lab Setup

Estimated Time: 45 minutes

Virtual machines	20410C-LON-DC1 20410C-LON-SVR1 20410C-LON-RTR 20410C-LON-CL1 20410C-LON-CL2
User name	Adatum\Administrator
Password	Pa\$\$w0rd

For this lab, you will use the available virtual machine environment. Before beginning 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 Microsoft Hyper-V® Manager, click **20410C-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in using the following credentials:
 - User name: **Administrator**
 - Password: **Pa\$\$w0rd**
 - Domain: **Adatum**
5. Repeat steps 2 through 4 for 20410C-LON-SVR1 and 20410C-LON-CL1.
6. For the optional Exercise 2, you should repeat steps 2 through 4 for 20410C-LON-RTR and 20410C-LON-CL2.

Exercise 1: Implementing DHCP

Scenario

As part of configuring the infrastructure for the new branch office, you need to configure a DHCP server that will provide IP addresses and configuration to client computers. Servers are configured with static IP addresses and usually do not use DHCP for obtaining IP addresses.

One of the client computers in the branch office needs to access an accounting app in the head office. The network team uses firewalls based on IP addresses to restrict access to this app. The network team has requested that you assign a static IP address to this client computer. Rather than configuring a static IP address on the client computer manually, you decide to create a reservation in DHCP for the client computer.

The main tasks for this exercise are as follows:

1. Install the Dynamic Host Configuration Protocol (DHCP) server role.
2. Configure the DHCP scope and options.
3. Configure the client to use DHCP, and then test the configuration.
4. Configure a lease as a reservation.

► Task 1: Install the Dynamic Host Configuration Protocol (DHCP) server role

1. Sign in to LON-SVR1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Open **Server Manager**, and install the **DHCP Server** role.
3. In the Add Roles and Features Wizard, accept all defaults.

► Task 2: Configure the DHCP scope and options

1. In Server Manager, open the DHCP console.
2. Authorize the **lon-svr1.adatum.com** server in AD DS.
3. In DHCP, in the navigation pane, browse to **IPv4**, right-click **IPv4**, and then click **New Scope**.
4. Create a new scope with the following properties:
 - Name: **Branch Office**
 - IP Address Range: **172.16.0.100-172.16.0.200**
 - Length: **16**
 - Subnet Mask: **255.255.0.0**
 - Exclusions: **172.16.0.190-172.16.0.200**
 - Configure options **Router 172.16.0.1**
 - For all other settings use default values
5. Activate the scope.

► Task 3: Configure the client to use DHCP, and then test the configuration

1. Sign in to **20410C-LON-CL1** as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Reconfigure the Ethernet Connection using the following information:
 - **Configure Internet Protocol Version 4 (TCP/IPv4)**
 - **Obtain an IP address automatically**
 - **Obtain DNS server address automatically**

3. Open the Command Prompt window, and then initiate the DHCP process using the **ipconfig /renew** command.
4. To test the configuration, verify that LON-CL1 has received an IP address from the DHCP scope by typing **ipconfig /all** in the Command Prompt window.



Note: This command returns information such as IP address, subnet mask, and DHCP enabled status, which should be **Yes**.

► Task 4: Configure a lease as a reservation

1. To display the physical address of the network adapter, in the Command Prompt window, type **ipconfig /all**.
2. Switch to **LON-SVR1**.
3. Open the **DHCP** console.
4. In the DHCP console, in the navigation pane, browse to **Scope [172.16.0.0] Branch Office**, right-click **Reservations**, and then click **New Reservation**.
5. Create a new reservation for **LON-CL1** using the physical address of the **LON-CL1** network adapter, and the IP address **172.16.0.155**.
6. On LON-CL1, use the **ipconfig** command to renew and then verify the IP address.

Results: After completing this exercise, you should have implemented DHCP, configured DHCP scope and options, and configured a DHCP reservation.

► Prepare for the optional exercise

If you are going to complete the optional lab, revert the 20410C-LON-CL1 and 20410C-LON-SVR1 virtual machines by performing the following steps:

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20410C-LON-CL1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 1 through 3 for 20410C-LON-SVR1.
5. Start 20410C-LON-SVR1.

Exercise 2: Implementing a DHCP Relay Agent (Optional Exercise)

Scenario

To avoid configuring an additional DHCP server on the subnet, your manager has asked you to configure a DHCP relay agent for another subnet in your branch office.

The main tasks for this exercise are as follows:

1. Install a DHCP relay agent.
2. Configure a DHCP relay agent.
3. Test the DHCP relay agent with a client.


► **Task 1: Install a DHCP relay agent**

1. Sign in to LON-RTR as **Adatum\Administrator** with password **Pa\$\$w0rd**.
2. In Server Manager, open **Routing and Remote Access**.
3. Add the DHCP relay agent to the router.

► **Task 2: Configure a DHCP relay agent**

1. Open **Routing and Remote Access**.
2. Configure the DHCP relay agent by performing the following steps:
 - a. In the navigation pane, right-click **DHCP Relay Agent**, and then click **New Interface**.
 - b. In the **New Interface for DHCP Relay Agent** dialog box, click **Ethernet 2**, and then click **OK**.
 - c. In the **DHCP Relay Agent Properties – Ethernet 2 Properties** dialog box, click **OK**.
 - d. Right-click **DHCP Relay Agent**, and then click **Properties**.
 - e. In the **DHCP Relay Agent Properties** dialog box, in the **Server address** box, type **172.16.0.11**, click **Add**, and then click **OK**.
3. Close **Routing and Remote Access**.

► **Task 3: Test the DHCP relay agent with a client**

 **Note:** To test how a client receives an IP address from the DHCP relay agent in another subnet, you need to create another DHCP scope.

1. Sign in to LON-SVR1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Run Windows PowerShell as Administrator, and then type the following cmdlets:

```
Add-WindowsFeature -IncludeManagementTools dhcp
netsh dhcp add securitygroups
Restart-service dhcpserver
Add-DhcpServerInDC LON-SVR1 172.16.0.11
Add-DhcpServerv4Scope -Name "Branch Office 2" -StartRange 10.10.0.100 -EndRange
10.10.0.200 -SubnetMask 255.255.0.0
Add-DhcpServerv4ExclusionRange -ScopeID 10.10.0.0 -StartRange 10.10.0.190 -EndRange
10.10.0.200
Set-DhcpServerv4OptionValue -Router 10.10.0.1
Set-DhcpServerv4Scope -ScopeID 10.10.0.0 -State Active
```

3. To test the client, switch to **LON-CL2**.
4. Open the Network and Sharing Center window, and then configure the **Ethernet, Internet Protocol Version 4 (TCP/IPv4)** properties with the following settings:
 - **Obtain IP address automatically**
 - **Obtain DNS server address automatically**
5. Open the Command Prompt window.

6. In the Command Prompt window, at a command prompt, type the following command:

```
Ipconfig /renew
```

7. Verify that IP address and DNS server settings on LON-CL2 are obtained from DHCP Server scope installed on **LON-SVR1**.



Note: The IP address should be in the following range: **10.10.0.100/16** to **10.10.0.200/16**.

Results: After completing this exercise, you should have implemented a DHCP relay agent.

► Prepare for the next module

After you finish the lab, revert the virtual machines back 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 **20410C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for 20410C-LON-SVR1, 20410C-LON-RTR, and 20410C-LON-CL2.

Lab Review Questions

Question: For what is the DHCP scope used?

Question: How should you configure a computer to receive an IP address from the DHCP server?

Question: Why do you need MAC address for a DHCP server reservation?

Question: What information do you need to configure on a DHCP relay agent?

Module Review and Takeaways

Module Review Questions

Question: You have two subnets in your organization. You want to use DHCP to allocate addresses to client computers in both subnets, but you do not want to deploy two DHCP servers. What factors must you consider?

Question: Your organization has grown, and your IPv4 scope is almost out of addresses. What should you do?

Question: What information do you require to configure a DHCP reservation?

Question: Can you configure option 003 – Router as a Server-level DHCP scope option?

Best Practices

The following are some best practices you can follow:

- Design your IP addressing scheme carefully so that it accommodates the requirements of both your current and future IT infrastructure.
- Determine which devices need DHCP reservations, such as network printers, network scanners, or IP-based cameras.
- Secure your network from unauthorized DHCP servers.
- Configure the DHCP database on highly available disk drive configurations, such as a redundant array of independent disks (RAID)-5 or RAID-1, to provide DHCP service availability in case of a disk failure.
- Back up the DHCP database regularly. Test the restore procedure in an isolated, non-production environment.
- Monitor the system utilization of DHCP servers. Upgrade the DHCP server hardware if necessary to provide better service performance.

Tools

Tool	Use	Where to find it
DHCP	Graphical User Interface for managing DHCP Server	Server Manager
Windows PowerShell	Command-line interface for managing DHCP Server	Windows taskbar on the Desktop
Ipconfig.exe	Managing and troubleshooting client IP settings	Command line
Netsh.exe	Configuring both client and server-side IP settings, including those for DHCP server role	Command line
Regedit.exe	Editing and fine-tuning settings, including those for the DHCP server role	Windows interface or Command line

Module 7

Implementing DNS

Contents:

Module Overview	7-1
Lesson 1: Name Resolution for Windows Clients and Servers	7-2
Lesson 2: Installing a DNS Server	7-15
Lesson 3: Managing DNS Zones	7-21
Lab: Implementing DNS	7-25
Module Review and Takeaways	7-30

Module Overview

Name resolution is the process of software translating between names that users can read and understand, and numerical IP addresses, which are necessary for TCP/IP communications. Because of this, name resolution is one of the most important concepts of every network infrastructure. You can think about DNS as being like the Internet's phone book for computers. Client computers use the name resolution process when locating hosts on the Internet and when locating other hosts and services in an internal network. Domain Name System (DNS) is one of the most common technologies for name resolution. Active Directory® Domain Services (AD DS) depends heavily on DNS, as does Internet traffic. This module discusses some basic name resolution concepts, and installing and configuring a DNS Server service and its components.

Objectives

After completing this module, you will be able to:

- Describe name resolution for Windows® operating system clients and Windows Server® servers.
- Install and manage a DNS Server.
- Manage DNS zones.

Lesson 1

Name Resolution for Windows Clients and Servers

You can configure a computer to communicate over a network by using a name in place of an IP address. The computer then uses name resolution to find an IP address that corresponds to a name, such as a host name. This lesson focuses on different types of computer names, the methods used to resolve them, and how to troubleshoot problems with name resolution.

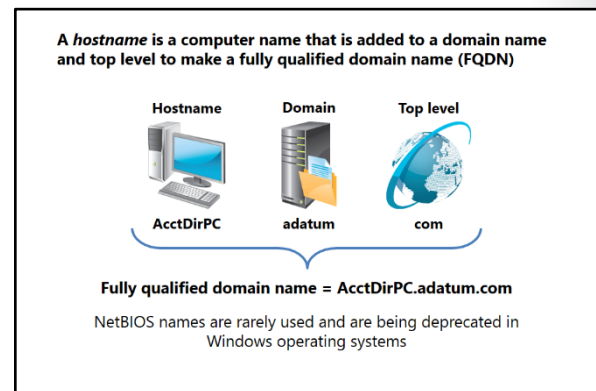
Lesson Objectives

After completing this lesson you will be able to:

- Describe computer names.
- Describe DNS.
- Describe DNS zones and records.
- Describe how Internet DNS names are resolved.
- Describe split DNS.
- Describe Link-local Multicast Name Resolution.
- Describe how a client resolves a name.
- Troubleshoot name resolution.

What Are the Computer Names Assigned to Computers?

The TCP/IP set of protocols identifies source and destination computers by their IP addresses. However, computer users are much better at using and remembering names than numbers. Because of this, administrators usually assign names to computers. Administrators then link these names to computer IP addresses in a name resolution system such as DNS. These names are in either in *host name* format, for example *dc1.contoso.com*, which is recognized by DNS, or in *NetBIOS name* format, for example *DC1*, which is recognized by Windows Internet Name Service (WINS).



Name Type

The type of name that an app uses, either host name or NetBIOS name, is determined by the application developer. If the application developer designs an application to request network services through Windows sockets, host names are used. If, on the other hand, the application developer designs an application to request services through NetBIOS, a NetBIOS name is used. Most current apps, including Internet apps, use Windows sockets—and thus use host names—to access network services.

Host Names

A *host name* is a user-friendly name that is associated with a computer's IP address to identify it as a TCP/IP host. The host name can be up to 255 characters long, and can contain alphabetic and numeric characters, periods, and hyphens.

You can use host names in various forms. The two most common forms are:

- An alias
- A fully qualified domain name (FQDN)

An alias is a single name that is associated with an IP address, such as *payroll*. You can combine an alias with a domain name to create an FQDN. An FQDN is structured for use on the Internet, and includes periods as separators. An example of an FQDN is *payroll.contoso.com*.

Creating Host Names

When you select host names, you should create host names that are intuitive and relatively easy-to-remember, yet still unique. The following lists some best practices to implement when creating host names:

- Select computer names that are easy for users to remember.
- Identify the owner of a computer in the computer name. For example, **john-doe-1** indicates that John Doe uses the computer.
- Select names that describe the purpose of the computer. For example, a file server named **past-accounts-1** indicates that the file server stores information related to past accounts.
- Do not use character case to convey the owner or purpose of a computer. DNS is not case-sensitive.
- Match the Active Directory domain name to the primary DNS suffix of the computer name.
- Use unique names for all computers in your organization. Do not assign the same computer name to different computers in different DNS domains.

What Is DNS?

DNS is a service that resolves FQDNs and other host names to IP addresses. All Windows Server operating systems include a DNS Server service.

When you use DNS, users on your network can locate network resources by typing in user-friendly names (for example, *www.microsoft.com*), which the computer then resolves to an IP address. The benefit is that IPv4 addresses may be difficult to remember (for example, 131.107.0.32), while a domain name typically is easier to remember. In addition, you can use host names that do not change while the underlying IP addresses can be changed to suit your organizational needs.

DNS can be used to:

- Resolve host names to IP addresses
- Locate domain controllers and global catalog servers
- Resolve IP addresses to host names
- Locate mail servers during email delivery

DNS uses a database of names and IP addresses, stored in a file or in AD DS, to provide this service. DNS client software performs queries on and updates to the DNS database. For example, within an organization, a user who is trying to locate a print server can use the DNS name *printserver.contoso.com*, and the DNS client software resolves the name to a printer's IP address, such as 172.16.23.55. Even if the printer's IP address changes, the user-friendly name can remain the same.

Originally, there was one file on the Internet that contained a list of all domain names and their corresponding IP addresses. This list quickly became too long to manage and distribute. DNS was developed to solve the problems associated with using a single Internet file. With the adoption of IPv6, DNS becomes even more important, because IPv6 addresses are even more complex than IPv4 addresses (for example, 2001:db8:4136:e38c:384f:3764:b59c:3d97).

DNS groups information about network resources into a hierarchical structure of domains.

The hierarchical structure of domains is an inverted tree structure beginning with a root domain at its apex, and descending into separate branches with common levels of parent domains, and descending downward even further into individual child domains.

As the Internet has grown, so has the number of domains from different countries. All countries in DNS have top-level country codes. The governing bodies in these countries can further create second-level domains that reflect categories such as .com, .org, and .net. For example, the United Kingdom (UK) has a top-level domain named .uk, and has further broken this down to the second-level for various activities. A commercial company in the UK may therefore have a FQDN of *companyname.com.uk*. This domain would not be the same as *companyname.com*, which is at an entirely different level.

The representation of the entire hierarchical domain structure as shown in the following illustration is known as a DNS namespace.




FIGURE 7.1: DNS NAMESPACE

The Internet uses a single DNS namespace with multiple root servers. To participate in the Internet DNS namespace, a domain name must be registered with a DNS registrar. This ensures that no two organizations attempt to use the same domain name.

If hosts that are located on the Internet do not need to resolve names in your domain, you can host a domain internally, without registering it. However, you must still ensure that the domain name is unique from Internet domain names, or connectivity to Internet resources might be affected. A common way to ensure uniqueness is to create an internal domain in the .local domain. The .local domain is reserved for internal use in much the same way that private IP addresses are reserved for internal use.

In addition to resolving host names to IP addresses, DNS can be used to:

- Locate domain controllers and global catalog servers. This is used when logging on to AD DS.
- Resolve IP addresses to host names. This is useful when a log file contains only the IP address of a host.
- Locate a mail server for email delivery. This is used for the delivery of all Internet email.

DNS Zones and Records

A *DNS zone* is the specific portion of a DNS namespace (such as adatum.com) that contains DNS records. A DNS zone is hosted on a DNS server that is responsible for responding to queries for records in a specific domain. For example, the DNS server that is responsible for resolving www.contoso.com to an IP address would contain the contoso.com zone.

You can store DNS zone content in a file or in the AD DS database. When the DNS server stores the zone in a file, that file is located in a local folder on the server. When the zone is not stored in AD DS, only one copy of the zone is a writable copy, and all the other copies are read-only.

The most commonly used types of zones in Windows Server DNS are forward lookup zones and reverse lookup zones.

A DNS zone is a specific portion of DNS namespace that contains DNS records

Zone types:

- Forward lookup zone
- Reverse lookup zone

Resource records in forward lookup zones include:

- A, MX, SRV, NS, SOA, and CNAME

Resource records in reverse lookup zones include:

- PTR

Forward Lookup Zones

Forward lookup zones resolve host names to IP addresses and host common resource records, including:

- Host (A) records
- Alias (CNAME) records
- Service (SRV) records
- Mail exchanger (MX) records
- Start of authority (SOA) records
- Name server (NS) records

The most common record type is the host (A) resource record.

Reverse Lookup Zones

Reverse lookup zones resolve IP addresses to domain names. A reverse lookup zone functions in the same manner as a forward lookup zone, but the IP address is part of the query and the host name is the returned information. Reverse lookup zones are not always configured, but you should configure them to reduce warning and error messages. Reverse lookup zones host SOA, NS, and pointer (PTR) resource records.

PTR Records

When you create host records in the DNS console, you also have the option to make a PTR record at the same time, if an appropriate reverse lookup zone exists. PTR records can be created automatically and added to a reverse lookup zone when an A record is created in a forward lookup zone. These PTR records are automatically deleted if the corresponding A resource record is deleted. You only need to manually create a PTR record once. Since it is not tied to an A resource record, it is not deleted if the A resource record is deleted. Client computers can create their PTR records when they dynamically update. A PTR record is in the format of IP Address, type of record (PTR) and hostname.

Many standard Internet protocols rely on reverse lookup zone lookup data to validate forward lookup zone information. For example, if the forward lookup indicates that training.contoso.com is resolved to 192.168.2.45, you can use a reverse lookup to confirm that 192.168.2.45 is associated with training.contoso.com.



Note: In Windows Server 2008 R2 and Windows Server 2012, you can also use DNSSec technology to perform similar type of verification. There are new enhancements to DNSSec in Windows Server 2012 R2 in encryption key management; however, these enhancements are beyond the scope of this lesson.

Many email servers use a reverse lookup as one way of reducing spam. By performing a reverse lookup, email servers try to detect open Simple Mail Transfer Protocol (SMTP) servers (open relays).

Having a reverse lookup zone is important if you have apps that rely on looking up hosts by their IP addresses. Many apps record this information in security or event logs. If you see suspicious activity from a particular IP address, you can look up the host name using the reverse lookup zone information.

Resource Records

The DNS zone file stores resource records. *Resource records* specify a resource type and the IP address to locate the resource. The most common resource record is a host (A) resource record. This is a simple record that resolves a host name to an IP address. The host can be a workstation, server, or another network device, such as a router.

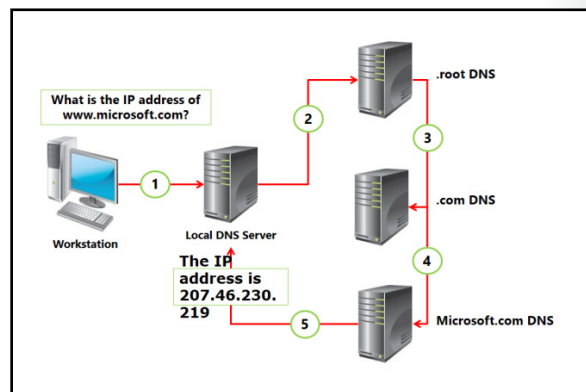
Resource records also help find resources for a particular domain. For instance, when a Microsoft Exchange Server needs to find the server that is responsible for delivering mail for another domain, it requests the mail exchanger (MX) resource record for that domain. This record points to the host (A) resource record of the host that is running the SMTP mail service.

Resource records also can contain custom attributes. MX records, for instance, have a Preference attribute, which is useful if an organization has multiple mail servers. The MX record tells the sending server which mail server the receiving organization prefers. SRV records also contain information about the port the service is listening to, and the protocol that you should use to communicate with the service.

How Internet DNS Names Are Resolved

When resolving DNS names on the Internet, an entire system of computers is used rather than just a single server. There are hundreds of servers on the Internet, called *root servers*, which manage the overall practice of DNS resolution. These servers are represented by 13 FQDNs. A list of these 13 servers is preloaded on each DNS server. When you register a domain name on the Internet, you are paying to become part of this system.

To see how these servers work together to resolve a DNS name, look at the following name resolution process for the name `www.microsoft.com`:



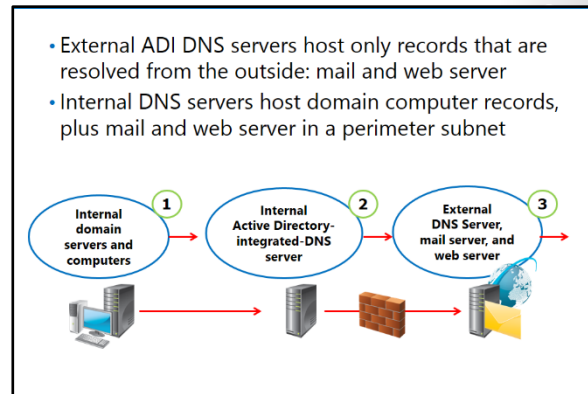
1. A workstation queries the local DNS server for the IP address `www.microsoft.com`.
2. If the local DNS server does not have the information, it queries a root DNS server for the location of the `.com` DNS servers.
3. The local DNS server queries a `.com` DNS server for the location of the `microsoft.com` DNS servers.
4. The local DNS server queries the `microsoft.com` DNS server for the IP address of `www.microsoft.com`.
5. The IP address of `www.microsoft.com` is returned to the workstation.

The name resolution process can be modified by caching or forwarding:

- **Caching.** After a local DNS server resolves a DNS name, it caches the results for the period of time defined by the time to live (TTL) value in the SOA record for the DNS zone. The default TTL is one hour. Subsequent resolution requests for the DNS name are given the cached information. Note that the TTL is not set by the caching server, but by the authoritative DNS server that resolved the name from its zone. When the TTL expires, the caching server must delete it. Subsequent requests for the same name would require a new name resolution request to the authoritative server.
- **Forwarding.** Instead of querying root servers, you can configure a DNS server to forward DNS requests to another DNS server. For example, requests for all Internet names can be forwarded to a DNS server at an Internet service provider (ISP).

What is Split DNS?

In Microsoft operating systems, DNS has two major functions: to resolve IP addresses to names (and vice versa), and to facilitate domain-level communications and authentication for AD DS. The ability to store service locator (SRV) records allows domain-joined clients to find domain controllers (DCs) for domain authentication and security while load balancing access to the various DC's using DNS round-robin functionality. However, Internet-level untrusted users from outside the firewall should never be able to access the SRV records and other sensitive AD DS information from the internal DNS servers. That data must remain separate and inaccessible from outside the firewall. At the same time, DNS records of servers and services hosting Internet level resources, such as web, mail and proxy servers, must remain accessible.



Split DNS, also known as *Split-brain DNS*, uses the same DNS domain name for both Internet and internal, domain-joined resources. However, the DNS server role is assigned to separate servers: one or more servers for the Internet, and the other server(s) for the AD DS domain. Deploying DNS in this manner requires extra steps to ensure the separation of sensitive information found on the AD DS domain side from the Internet side, and to ensure that outside the firewall access only goes to the DNS server deployed on the Internet side.

Because DNS is such a vital function for the AD DS, the DNS server role is usually included with domain controllers when they are deployed. This role can be integrated into AD DS so that DNS records are stored as Active Directory objects and attributes. The DNS zone type in this instance is referred to as Active Directory Integrated (ADI). ADI zones replace DNS zone transfers with AD DS replication and can ensure secure dynamic updates of client records to the zone. In a domain, using ADI DNS is considered a best practice.

With Split DNS, internal clients are only configured with the IP addresses of the ADI DNS servers, which are domain controllers. All client DNS dynamic updates are written to the servers. All DNS queries from internal clients go only to these DNS servers. If any resolutions for names are needed beyond the internal domain, such as for Internet web servers, the ADI DNS servers forward these requests to the Internet-facing DNS server. The Internet-facing DNS servers are normally deployed in the perimeter network between the firewalls. Although they have the same domain name as the ADI DNS servers, the Internet-facing DNS servers do not store the same data. All records in the Internet-facing DNS server zone are created manually. Normally the Internet-facing DNS server zone only contains records for itself and other servers that are located in the perimeter network and need to be accessed from the Internet.

When a query to the Internet-facing DNS server comes in from the Internet requesting a resolution on any domain-level resource, such as an SRV record, the Internet-facing DNS server rejects the query because it does not have any of the SRV records—these are only stored in the domain ADI DNS servers. Because it considers itself authoritative for the zone, the Internet-facing DNS server does not make an iterative query to the ADI DNS servers.

To further enhance security, you can set a firewall rule on the inside firewall, that is, the firewall between the internal and perimeter networks, to reject all DNS (UDP port 53) queries from the perimeter to the internal network, while still allowing DNS replies.



Note: When using DirectAccess for portable clients, be aware that when the client is deployed outside of the internal network it uses the Name Resolution Policy Table (NRPT) for continued access to internal resources. This sends DNS name queries for internal resources to the ADI DNS servers. With split DNS and DirectAccess clients, you need to add the Fully Qualified Domain Names (FQDN) of any Internet-level web servers kept in the perimeter network to the NRPT as a firewall exception rule.

What Is Link-local Multicast Name Resolution?

In Windows Server 2012, a new method for resolving names to IP addresses is Link-local Multicast Name Resolution (LLMNR). Because of various limitations that are beyond the scope of this lesson, LLMNR typically is used only on localized networks. Although LLMNR is able to resolve IPv4 addresses, it has been designed specifically for IPv6; therefore, if you want to use it, you must have IPv6 supported and enabled on your hosts.

LLMNR is commonly used in networks where:

- There are no DNS or NetBIOS services for name resolution.
- Implementation of these services is not practical for any reason.
- These services are not available.

For example, you might want to set up a temporary network for testing purposes without a server infrastructure.

LLMNR is supported on Windows Vista®, Windows Server 2008, and all newer Windows operating systems. It uses a simple system of request and reply messages to resolve computer names to IPv6 or IPv4 addresses. For a node to respond to a LLMNR request Network Discovery must be enabled, but Network Discovery is not needed just to make a request for name resolution.

To use LLMNR, you need to turn on the Network Discovery feature for all nodes on the local subnet. This feature is available in the Network and Sharing Center. Be aware that Network Discovery is usually disabled for any network that you designate as Public.

LLMNR is an additional method for name resolution that does not use DNS or WINS

- LLMNR is designed for IPv6
- Works only on Windows Vista, Windows Server 2008, and all newer Windows operating systems
- Network Discovery must be enabled
- Can be controlled via Group Policy

If you want to control the use of LLMNR on your network, you can configure it via Group Policy. To disable LLMNR via Group Policy, set the following Group Policy value:

Group Policy = Computer Configuration\Administrative Templates\Network\DNS Client\Turn off Multicast Name Resolution.

Set this value to Enabled if you do not want to use LLMNR, or to Disabled if you want to use LLMNR.

How a Client Resolves a Name

Windows operating systems support a number of different methods for resolving computer names, such as DNS, WINS, and the host name resolution process.

DNS

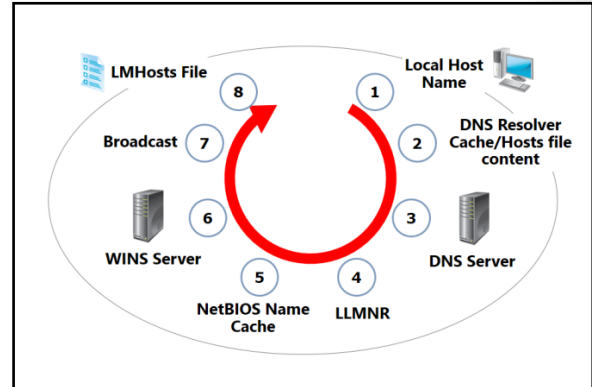
As previously discussed, DNS is the Microsoft standard for resolving host names to IP Addresses. For more information on DNS, refer back to second topic of this Lesson, *What is DNS*.

WINS

WINS provides a centralized database for registering dynamic mappings of a network's NetBIOS names. Windows operating systems retain support for WINS to provide backward compatibility.

You can resolve NetBIOS names by using:

- Broadcast messages. Broadcast messages, however, do not work well on large networks because routers do not propagate broadcasts.
- Lmhosts file on all computers. Using an Lmhosts file for NetBIOS name resolution is a high maintenance solution, because you must maintain the file manually on all computers.
- Hosts file on all computers. Similar to an Lmhosts file, you can also use a hosts file for NETBIOS name resolution. This file is also stored locally on each machine, and it is used for fixed mappings of names to IP addresses, on local network segment.



Note: The DNS server role in Windows Server 2008 R2 and Windows Server 2012 also provides a new zone type, the GlobalNames zone. You can use GlobalNames zone to resolve single-label names that are unique across an entire forest. This eliminates the need to use the NetBIOS-based WINS to provide support for single-label names.

Host Name Resolution Process

When an app specifies a host name and uses Windows sockets, TCP/IP uses the DNS resolver cache and DNS when attempting to resolve the host name. The hosts file is loaded into the DNS resolver cache. If NetBIOS over TCP/IP is enabled, TCP/IP also uses NetBIOS name resolution methods when resolving host names.

Windows operating systems resolve host names by performing the following tasks in this specific order:

1. Checks whether the host name is the same as the local host name.
2. Searches the DNS resolver cache. In the DNS client resolver cache, entries from hosts file are preloaded.

3. Sends a DNS request to its configured DNS servers.
4. Searches the network using LLMNR, if it is enabled.
5. Converts the host name to a NetBIOS name and checking the local NetBIOS name cache.
6. Contacts the host's configured WINS servers.
7. Broadcasts as many as three NetBIOS name query request messages on the subnet that is attached directly.
8. Searches the Lmhosts file.



Note: You can control the order used to resolve names. For example, if you disable NetBIOS over TCP/IP, none of the NetBIOS name resolution methods are attempted. Alternatively, you can modify the NetBIOS node type, which changes the order in which the NetBIOS name resolution methods are attempted.



Additional Reading: To learn more about LLMNR, please see <http://go.microsoft.com/fwlink/?LinkID=331077>.

Troubleshooting Name Resolution

Like most of other technologies, name resolution sometimes requires troubleshooting. Issues can occur when the DNS server, its zones, and its resource records are not configured properly. When resource records are causing issues, it can sometimes be difficult to identify the issue because configuration problems are not always obvious.

Windows Server 2012 R2 Cmdlets

Additionally, Windows Server 2012 R2 provides the following functionality:

- **Step-DnsServerSigningKeyRollover.** This cmdlet forces a Key Signing Key (KSK) rollover when waiting for a parent delegation signer (DS) update. If a server that is hosting a securely-delegated zone is unable to check to determine that the DS record in its parent is updated, this parameter allows you to force a rollover. The server expects the DS record to be manually updated in the parent.
- **Add-DnsServerTrustAnchor -Root.** The Root parameter set enables you to retrieve trust anchors from the URL specified in the **RootTrustAnchorsURL** property of the DNS server. This cmdlet has the following alias: **Retrieve-DnsServerRootTrustAnchor**.
- **RootTrustAnchorsURL.** The **Get-DnsServerSetting** and **Set-DnsServerSetting** cmdlets are extended to add a new output string of **RootTrustAnchorURL**.

Windows PowerShell® has extended functionality in Windows Server 2012 R2 with enhanced zone-level statistics that are accessible through the **Get-DnsServerStatistics** cmdlet.



Additional Reading: For more information on the parameters for the **Get-DnsServerStatistics** cmdlet, go to <http://go.microsoft.com/fwlink/?LinkID=331076>.

A new Windows PowerShell DNS module with numerous cmdlets was introduced with Windows Server 2012 R2

- More cmdlets were added with Windows Server 2012 R2

Command-line tools to troubleshoot configuration issues:

- Nslookup
- DNSCmd
- Dnslint
- Ipconfig

The troubleshooting process:

1. Identify client DNS server with nslookup or Resolve-DnsName
2. Communicate via ping
3. Use nslookup to verify records

The following table details the **ZoneTransferStatistics** cmdlet, which returns information about full and incremental zone transfers.

Functionality	Returns information about zone transfer requests:
RequestReceived	<ul style="list-style-type: none"> Received when the DNS server is a primary server for a zone
RequestSent	<ul style="list-style-type: none"> Sent when the DNS server is a secondary server for a zone
ResponseReceived	<ul style="list-style-type: none"> Received when the DNS server is a secondary server for a zone
SuccessReceived	<ul style="list-style-type: none"> Successful and received when the DNS server is a secondary server for a zone
SuccessSent	<ul style="list-style-type: none"> Successful and received when the DNS server is a primary server for a zone

The following table details the ZoneUpdateStatistics cmdlet.

Functionality	Dynamic update information:
DynamicUpdateReceived	<ul style="list-style-type: none"> Dynamic update requests that are received by the DNS server
DynamicUpdateRejected	<ul style="list-style-type: none"> Dynamic updates that are rejected by the DNS server

To get zone level statistics, type the following code at an elevated Windows PowerShell prompt:

```
PS C:\> $statistics = Get-DnsServerStatistics -ZoneName Adatum.com
$statistics.ZoneQueryStatistics
$statistics.ZoneTransferStatistics
$statistics.ZoneUpdateStatistics
```

Command-Line Tools and Commands for Troubleshooting

The command-line tools and commands that you use to troubleshoot these and other configuration issues are as follows:

- **Nslookup.** Use this tool to query DNS information. The tool is flexible and can provide valuable information about DNS server status. You also can use it to look up resource records and validate their configuration. Additionally, you can test zone transfers, security options, and MX record resolution.
- **DNSScmd.** Use this command-line tool to manage the DNS server role. This tool is useful in scripting batch files to help automate routine DNS management tasks or to perform simple unattended setup and configuration of new DNS servers on your network.
- **Dnslint.** Use this tool to diagnose common DNS issues. This tool diagnoses configuration issues in DNS quickly, and can generate a report in HTML format regarding the status of the domain that you are testing.



Reference Links: The **Dnslint** command can be downloaded from <http://go.microsoft.com/fwlink/?LinkId=286763>.

- **Ipconfig.** Use this command to view and modify IP configuration details that the computer uses. This tool includes additional command-line options that you can use to troubleshoot and support DNS clients. You can view the client local DNS cache using the command **ipconfig /displaydns**, and you can clear the local cache using **ipconfig /flushdns**. If you want to reregister a host in DNS, you can use **ipconfig /registerdns**.
- **Monitoring on DNS server.** Perform simple local queries and recursive queries from the DNS server Monitoring tab to test if the server can communicate with upstream servers. You also can schedule these tests for regular intervals. The DNS server Monitoring tab is available only in Windows Server 2008 and Windows Server 2012 in the DNS Server Name Properties dialog box.

In Windows Server 2012, there is a new set of Windows PowerShell cmdlets in that you can use for DNS client and server management. Some of the most commonly used cmdlets are as follows:

- **Clear-DNSClientCache.** This cmdlet clears the client cache, similar to **ipconfig /flushdns**.
- **Get-DNSClient.** This cmdlet displays the details of the network interfaces.
- **Get-DNSClientCache.** This cmdlet displays the content of the local DNS client cache.
- **Register-DNSClient.** This cmdlet registers all of the IP addresses on the computer onto the configured DNS server.
- **Resolve-DNSName.** This cmdlet performs a DNS name resolution for a specific name, similar to **nslookup**.
- **Set-DNSClient.** This cmdlet sets the interface-specific DNS client configurations on the computer.
- **Test-DNSServer.** This cmdlet tests that a specified computer is a functioning DNS server.

The Troubleshooting Process

When you troubleshoot name resolution, you must understand what name resolution methods the computer is using, and in what order the computer uses them. Be sure to clear the DNS resolver cache between resolution attempts.

If you cannot connect to a remote host and suspect a name resolution problem, you can troubleshoot the name resolution by performing the following steps:

1. Open an elevated command prompt, and then clear the DNS resolver cache by typing the following command at a command prompt:

```
ipconfig /flushdns
```

Alternatively, you can open Windows PowerShell and type the equivalent cmdlet at a Windows PowerShell prompt:

```
Clear-DNSClientCache
```

2. Attempt to ping the remote host by its IP address. This helps identify whether the issue is related to name resolution. If the ping succeeds by using the IP address but fails by using its host name, then the problem is related to name resolution.
3. Attempt to ping the remote host by using its host name. For accuracy, use the FQDN with a trailing period. For example, if you were working at Contoso, Ltd., you would enter the following command at a command prompt:

```
Ping LON-dc1.contoso.com
```


- If the ping is successful, the problem is probably not related to name resolution. If the ping is unsuccessful, edit the C:\windows\system32\drivers\etc\hosts text file, and then add the appropriate entry to the end of the file. For instance, in the previous Contoso, Ltd. example, you would add the following line and then save the file:

```
10.10.0.10LON-dc1.contoso.com
```

- Perform the Ping-by-host-name test once more. Name resolution should now be successful. Verify that the name was resolved correctly by examining the DNS resolver cache. To display the DNS resolver cache, at a command prompt type the following, or use the equivalent Windows PowerShell cmdlet:

```
Ipconfig /displaydns
```

- Remove the entry that you added to the hosts file, and then clear the resolver cache once more.
- At the command prompt, type the following command:

```
Nslookup.exe -d LON-dc1.contoso.com. > filename.txt
```

Examine the contents of the filename.txt file to identify the failed stage in name resolution.



Note: You should also know how to interpret the DNS resolver cache output so that you can identify whether the name resolution problem is with the client computer's configuration, the name server, or the configuration of records within the name server zone database. Interpreting the DNS resolver cache output is beyond the scope of this lesson.

Demonstration: Troubleshooting Name Resolution

In this demonstration, you will see how to use Windows PowerShell cmdlets and command-line tools to troubleshoot DNS.

Demonstration Steps

Use Windows PowerShell cmdlets to troubleshoot DNS

- Sign in to LON-DC1 and LON-CL1 as **Adatum\Administrator** with a password of **Pa\$\$w0rd**.
- On LON-CL1, open **Windows PowerShell**, run the following cmdlet, and then examine the results:

```
Get-DnsClientServerAddress
```

- In the **Network and Sharing Center**, record the **static TCP/IP** address properties, and then change the network interface to **automatic**.
- Switch back to **Windows PowerShell** and run the following cmdlet, and then make a note of the results:

```
Get-DnsClientServerAddress  
Clear-DnsClientCache
```

- Write the Interface Index value of the Ethernet interfaces' IPv4 row, here:
- Run the following cmdlet:

```
Resolve-DnsName lon-dc1
```

Note that the cmdlet issues the following error message: "A DNS server is not found."

7. Run the following cmdlets, where *X* is the Interface Index value that you wrote down in step 5:

```
Set-DnsClientServerAddress -InterfaceIndex X -ServerAddress 172.16.0.10
Get-DnsClientServerAddress
Resolve-DnsName lon-dc1
```

The error does not report back, and an address is returned.

8. Switch back to the **Network and Sharing Center** and enter the **static TCP/IP** you wrote down earlier.
9. In Windows PowerShell, use the following cmdlets:

```
Get-DnsClientCache
Clear-DnsClientCache
Get-DnsClientCache
Get-DnsClientGlobalSetting
Register-DnsClient
```

10. Close both the **Windows PowerShell** and the **Network and Sharing Center** windows.

Using Command-line tools to troubleshoot DNS

1. Run an elevated **Command Prompt** as **Administrator** and run **ipconfig /all**.
2. Run the **nslookup** command, and then search for the **LON-CL1** address. Close the nslookup command.
3. Switch to **LON-DC1** and open an elevated **Command Prompt** as **Administrator**.
4. Run the **dnscmd /?** Command, and note the options.
5. Run **ipconfig /displaydns** and note the output values displayed.
6. Run the command **ipconfig /flushdns** and then run **ipconfig /displaydns** again.
7. Run the **ping** command on **LON-CL1**.
8. Use the **ipconfig /displaydns** command to display the host record for LON-CL1.

Although the request packets are ignored, note that the command returned the **FQDN**, which proves that the name resolution was successful.

9. Close all open windows and then sign out from LON-CL1 and LON-DC1.

Lesson 2

Installing a DNS Server

To use a DNS Server, you must first install it. Installing the DNS Server service on a DNS server is a simple procedure. To manage your DNS Server service, it is important that you understand the DNS server components and their purpose. In this lesson, you will learn about DNS components and how to install and manage the DNS Server role.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the components of a DNS solution.
- Describe root hints.
- Describe DNS queries.
- Describe forwarding.
- Explain how DNS server caching works.
- Explain how to install the DNS server role.

What Are the Components of a DNS Solution?

The components of a DNS solution include DNS servers, DNS servers on the Internet, and DNS resolvers (or *DNS clients*).

DNS Server

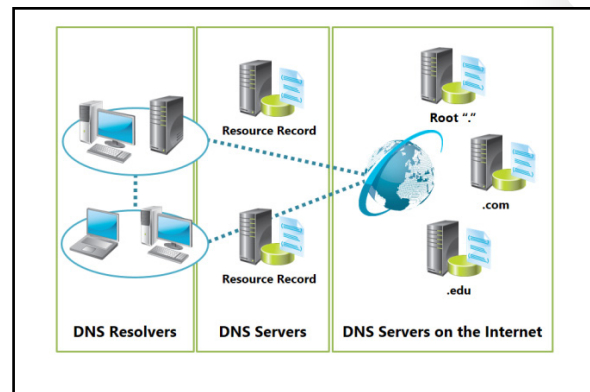
A DNS server answers recursive and iterative DNS queries. DNS servers also can host one or more zones of a particular domain. Zones contain different resource records. DNS servers also can cache lookups to save time for common queries.

DNS Servers on the Internet

DNS servers on the Internet are accessible publicly. These servers host information about public domains, such as common top-level domains (TLDs), for example, .com, .net, and .edu.

DNS Resolver

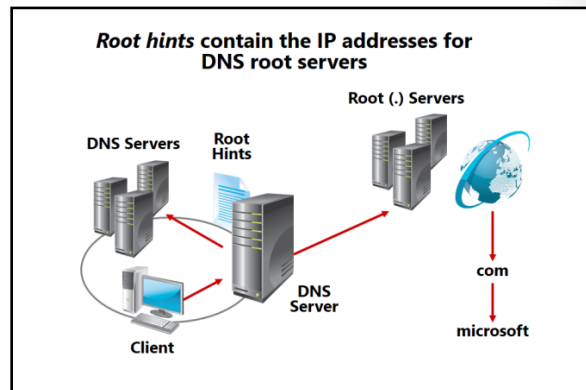
The DNS resolver generates and sends iterative or recursive queries to the DNS Server. A DNS resolver can be any computer that is performing a DNS lookup that requires interaction with the DNS server. DNS servers also can issue DNS requests to other DNS servers.



What Are Root Hints?

Root hints are a list of the 13 FQDNs on the Internet that your DNS server uses if it cannot resolve a DNS query by using its own zone data, a DNS forwarder, or its own cache. The root hints list the highest servers in the DNS hierarchy, and can provide the necessary information for a DNS server to perform an iterative query to the next lowest layer of the DNS namespace.

Root Servers are installed automatically when you install the DNS role. They are copied from the `cache.dns` file that is included in the DNS role setup files. You also can add root hints to a DNS server to support lookups for non-contiguous domains within a forest.



When a DNS server communicates with a root hint server, it uses only an iterative query. To configure a server to use only recursive queries to a forwarder, configure the forwarder on the DNS server properties. If you want to disable all iterative queries, deselect the **Use root hints if no forwarders are available** check box on the Forwarders tab. If you configure the server to use only a forwarder, and you disable root hints, it attempts to send a recursive query to its forwarding server; if the forwarding server does not answer this query, the first server responds that the host could not be found.

It is important to understand that recursion on a DNS server and recursive queries are not the same thing. Recursion on a DNS server means that the server uses its root hints to try to resolve a DNS query, whereas a recursive query is a query that is made to a DNS server in which the requester asks the server to assume the responsibility for providing a complete answer to the query.

What Are DNS Queries?

A *DNS query* is a name resolution query that is sent to a DNS Server. The DNS server then provides either an authoritative or a non-authoritative response to the client query.



Note: It is important to note that DNS servers also can act as DNS resolvers, and send DNS queries to other DNS servers.

Authoritative or Non-Authoritative Responses

The two types of responses are:

- **Authoritative.** An authoritative response is one in which the server returns an answer that it knows is correct, because the request is directed to the authoritative server that manages the domain. A DNS server is authoritative when it hosts a primary or secondary copy of a DNS zone.
- **Non-authoritative.** A non-authoritative response is one where the DNS server that contains the requested domain in its cache answers a query by using forwarders or root hints. Because the answer provided might not be accurate (because only the authoritative DNS server for the given domain can issue that information), it is called a non-authoritative response.

- Queries are recursive or iterative
- DNS clients and DNS servers initiate queries
- DNS servers are authoritative or non-authoritative for a namespace
- An authoritative DNS server for the namespace either:
 - Returns the requested IP address
 - Returns an authoritative "No, that name does not exist"
- A non-authoritative DNS server for the namespace either:
 - Checks its cache
 - Uses forwarders
 - Uses root hints



If the DNS server is authoritative for the query's namespace, the DNS server checks the zone and then does one of the following:

- Returns the requested address.
- Returns an authoritative "No, that name does not exist."



Note: An authoritative answer can be given only by the server with direct authority for the queried name.

If the local DNS server is non-authoritative for the query's namespace, the DNS server does one of the following:

- Checks its cache and returns a cached response.
- Forwards the unresolvable query to a specific server, called a *forwarder*.
- Uses well-known addresses of multiple root servers to find an authoritative DNS server to resolve the query. This process uses root hints.

Recursive Queries

In a recursive query, the requester asks the DNS server to obtain a fully resolved IP address of the requested resource, before returning the answer to the requestor. The DNS server may have to perform several queries to other DNS servers before it finds the answer. Recursive queries are generally made by a DNS client to a DNS server, or by a DNS server that is configured to pass unresolved queries to another DNS server, in the case of a DNS server configured to use a forwarder.

A recursive query has two possible results:

- The DNS server returns the IP address of the host requested.
- The DNS server cannot resolve an IP address.

For security reasons, it sometimes is necessary to disable recursive queries on a DNS server so that the DNS server in question does not attempt to forward its DNS requests to another server. This is useful when you do not want a particular DNS server to communicate outside its local network.

Iterative Queries

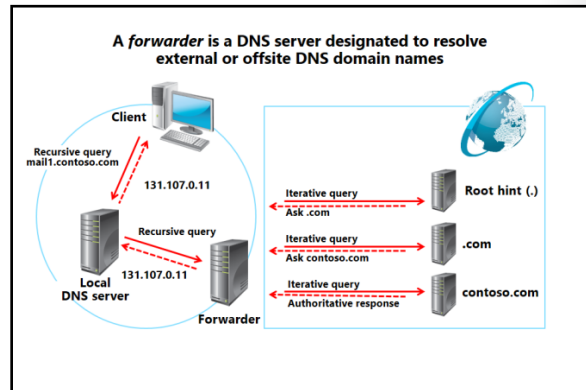
Iterative queries access domain name information that resides across the DNS system. You can use iterative queries to resolve names across many servers quickly and efficiently. When a DNS server receives a request that it cannot answer using its local information or its cached lookups, it makes the same request to another DNS server by using an iterative query. When a DNS server receives an iterative query, it might answer with either the IP address for the domain name (if known), or with a referral to the DNS servers that are responsible for the domain being queried. The DNS server continues this process until it locates a DNS server that is authoritative for the queried name, or until an error or time-out condition is met.

What Is Forwarding?

A *forwarder* is a network DNS server that forwards queries for external names to DNS servers outside of its network. You also can create and use conditional forwarders to forward queries according to specific domain names.

Once you designate a network DNS server as a forwarder, other DNS servers in the network forward the queries that they cannot resolve locally to that server. By using a forwarder, you can manage name resolution for names outside of your network, such as names on the Internet. This improves the efficiency of name resolution for your network's computers.

The forwarder must be able to communicate with the DNS server that is located on the Internet. This means either you configure it to forward requests to another DNS server, or you configure it to use root hints to communicate.



Best Practice: Use a central forwarding DNS server for Internet name resolution. This can improve security because you can isolate the forwarding DNS server in a perimeter network, which ensures that no server within the network is communicating directly to the Internet.

Conditional Forwarder

A *conditional forwarder* is a DNS server on a network that forwards DNS queries according to the query's DNS domain name. For example, you can configure a DNS server to forward all queries that it receives for names ending with *corp.contoso.com* to the IP address of a specific DNS server, or to the IP addresses of multiple DNS servers. This is useful when you have multiple DNS namespaces in a forest.

Conditional Forwarding in Windows Server 2008 R2 and Windows Server 2012

In Windows Server 2008 R2 and Windows Server 2012, the conditional forwarder configuration is in a node in the DNS console. You can replicate this information to other DNS servers through Active Directory-integrated DNS.




Best Practice: Use conditional forwarders if you have multiple internal namespaces. This results in faster name resolution.

How DNS Server Caching Works

DNS caching increases the performance of the organization's DNS system by decreasing the time it takes to provide DNS lookups.

When a DNS server resolves a DNS name successfully, it adds the name to its cache. Over time, this builds a cache of domain names and their associated IP addresses for most of the domains that the organization uses or accesses. The default time to keep a name in the cache is one hour. The zone owner can change this by modifying the SOA record for the appropriate DNS zone.



A caching-only server is the ideal type of DNS server to use as a forwarder. It does not host any DNS zone data; it only answers lookup requests for DNS clients.

In Windows Server 2012, you can access the content of the DNS server cache by selecting the Advanced view in the DNS Manager console. In this view, cached content is displayed as a node in DNS Manager. You can also delete single entries (or the entire cache) from the DNS server cache. Alternatively, you can use the Windows PowerShell **Get-DNSServerCache** cmdlet to view the cache content.

The DNS client cache is stored on the local computer by the DNS client service. To view client-side caching, at a command prompt run the **ipconfig /displaydns** command. This displays the local DNS client cache. If you need to clear the local cache, you can use the Windows PowerShell **Get-DNSClientCache** and **Clear-DNSClientCache** cmdlets, or the **ipconfig /flushdns** command.

To prevent DNS client caches from being overwritten, use the DNS Cache Locking feature available in Windows Server 2008 R2 and Windows Server 2012. When enabled, the cached records cannot be overwritten for the duration of the time-to-live (TTL) value. Cache locking provides improved security against cache poisoning attacks. This type of attack occurs when a false name resolution is provided by an attacker's DNS server. This false data is kept in the cache for as long as the attacker's DNS server has set the TTL value for that record, and therefore falsifies or *poisons* the cache.

How to Install the DNS Server Role

The DNS server role is not installed on Windows Server 2012 by default. Instead, you must add it in a role-based manner when you configure the server to perform the role. You install the DNS server role by using the Add Roles and Features Wizard in Server Manager.

You can also add the DNS server role when you promote your server to a domain controller. You do this from the domain controller Options page of the Active Directory Domain Services Installation Wizard.

Once you install the DNS server role, the DNS Manager snap-in becomes available to add to your administrative consoles. The snap-in is added automatically to the Server Manager console and to the DNS Manager console. You can run the DNS Manager from the Start screen by typing **dnsmgmt.msc**.

DNS server installation methods:

- Server Manager
- Active Directory Domain Services Installation Wizard

Tools available to manage DNS Server:

- DNS Manager snap-in
 - Server Manager
 - DNS Manager console (dnsmgmt.msc)
- DNSCmd command-line tool
- Windows Powershell
- Remote Server Administrative Tools

When you install the DNS server role, the **dnscmd.exe** command-line tool is also added. You can use the **DNSCmd** tool to script and automate DNS configuration. For help with this tool, at a command prompt, type **dnscmd.exe /?**

In Windows Server 2012, another method you can use to manage a DNS server is to use Windows PowerShell. It is recommended that you use Windows PowerShell cmdlets for command-line-based management of the DNS server. In addition, you can use the command-line tools **Nslookup**, **DNSCmd**, **Dnslint**, and **Ipconfig** in the Windows PowerShell environment.

To administer a remote DNS server, add the Remote Server Administrative Tools to your administrative workstation, which must be running a Windows Vista Service Pack 1 (SP1) or newer Windows operating system.

Demonstration: Installing the DNS Server Role

Many organizations now have or will want more than one DNS server on their network. You can install additional DNS servers by using the Server Manager console. If you want to enable your DNS server to resolve Internet names, you will most likely want to enable forwarding.

In this demonstration, you will see how to:

- Install a second DNS server.
- Create a forward lookup zone by using Windows PowerShell.
- Configure forwarding.

Demonstration Steps

Install a second DNS server

1. Sign in to LON-DC1 and LON-SVR1 as **Adatum\Administrator** with a password of **Pa\$\$w0rd**.
2. On LON-SVR1, open **Server Manager**.
3. Start the Add Roles and Features Wizard.
4. Add the DNS Server role.

Create a forward lookup zone by using Windows PowerShell

1. In Windows PowerShell, run the following cmdlet:

```
Add-DnsServerPrimaryZone -Name fabrikam.com -DynamicUpdate Secure -ReplicationScope Domain
```

2. Go to the DNS Console and verify that the **fabrikam.com** forward lookup zone was created with the appropriate settings.

Configure forwarding

- Configure the DNS Server with a forwarder on IP address **172.16.0.10**.



Note: Leave all virtual machines in their current state for the next demonstration.

Lesson 3

Managing DNS Zones

The DNS server hosts zone data in an Active Directory database or in the zone file. Also, the DNS server can host several types of zones. In this lesson, you will learn about DNS zone types and about Active Directory–integrated DNS zones.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe DNS zone types.
- Describe dynamic updates.
- Describe Active Directory–integrated zones.
- Explain how to create an Active Directory–integrated zone.

What Are DNS Zone Types?

There are four DNS zone types:

- Primary
- Secondary
- Stub
- Active Directory–integrated

Zones	Description
Primary	Read/write copy of a DNS database
Secondary	Read-only copy of a DNS database
Stub	Copy of a zone that contains only records used to locate name servers
Active Directory-integrated	Zone data is stored in AD DS rather than in zone files

Primary Zone

When the DNS server is both the host and the primary source for information about a zone, the zone is a *primary zone*. In addition, the DNS server stores the master copy of the zone data either in a local file or in AD DS. When the DNS server stores the zone data in a file, the primary zone file by default is named *zone_name.dns*, and is located on the server in the %windir%\System32\Dns folder. When the zone is not stored in AD DS, the primary zone server is the only DNS server that has a writable copy of the database.

Secondary Zone

When the DNS server is the host, but is the secondary source for zone information, the zone is a *secondary zone*. The zone information at this server must be obtained from another DNS server that also hosts the zone. This DNS server must have network access to the DNS server to receive updated zone information.

Because a secondary zone is a copy of a primary zone that another server hosts, the secondary zone cannot be stored in AD DS. Secondary zones can be useful if you are replicating data from non-Windows DNS zones.

Stub Zone

A *stub zone* is a replicated copy of a zone that contains only those resource records that are necessary to identify that zone's authoritative DNS servers. A stub zone resolves names between separate DNS namespaces, which might be necessary when a corporate merger requires that the DNS servers for two separate DNS namespaces resolve names for clients in both namespaces.

A stub zone consists of the following:

- The delegated zone's SOA resource record, NS resource records, and A resource records
- The IP address of one or more master servers used to update the stub zone

The master servers for a stub zone are one or more DNS servers that are authoritative for the child zone. Usually this is the DNS server that is hosting the primary zone for the delegated domain name.

Active Directory–Integrated Zone

If AD DS stores the zone data, DNS can use the multimaster replication model to replicate the primary zone data. This enables you to simultaneously edit zone data on more than one DNS server.

What Are Dynamic Updates?

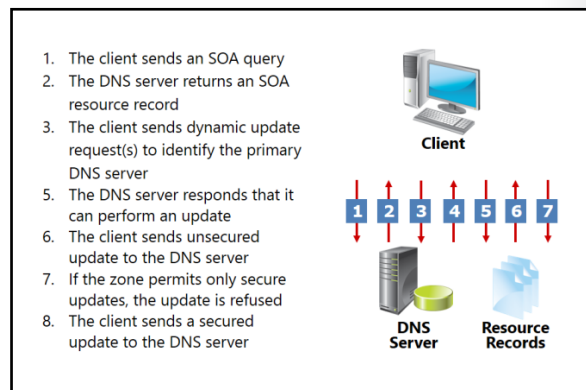
A *dynamic update* is an update to DNS in real time. Dynamic updates are important for DNS clients that change locations, because they can dynamically register and update their resource records without manual intervention.

The Dynamic Host Configuration Protocol (DHCP) client service performs the registration, regardless of whether the client's IP address is obtained from a DHCP server, or is fixed. The registration occurs during the following events:

- When the client starts and the DHCP client service is started
- When an IP address is configured, added, or changed on any network connection
- When an administrator executes the Windows PowerShell cmdlet **Register-DNSClient** or runs the **ipconfig /registerdns** at a command prompt

The process of dynamic updates is as follows:

1. The client identifies a name server and sends an update. If the name server hosts only a secondary zone, the name server refuses the client's update. If the zone is not an Active Directory–integrated zone, the client may have to do this several times.
2. If the zone supports dynamic updates, the client eventually reaches a DNS server that can write to the zone. This DNS server is one of the following:
 - a. The primary server for a standard, file-based zone
 - b. Any domain controller that is a name server for an Active Directory–integrated zone, which is, by default, considered primary because it is writable
3. If the zone is configured for secure dynamic updates, the DNS server refuses the change. The client then authenticates and resends the update.



In some configurations, you may not want clients to update their records even in a dynamic update zone. In this case, you can configure the DHCP server to register the records on the client's behalf. By default, a client registers that it is a (host/address) record, and the DHCP server registers the PTR (pointer/reverse lookup) record.

By default, Windows operating systems attempt to register their records with their DNS server. You can modify this behavior in the client IP configuration, or through Group Policy. Domain Controllers also register their SRV records (and their host records) in DNS. SRV records are registered automatically each time the NETLOGON service starts.

What Are Active Directory–Integrated Zones?

A DNS server can store zone data in the AD DS database provided that the DNS server is an AD DS domain controller. When the DNS server stores zone data in this way, this creates an Active Directory–integrated zone.

The benefits of an Active Directory–integrated zone are significant:

- **Multimaster updates.** Unlike standard primary zones, which can only be modified by a single primary server, Active Directory–integrated zones can be written to by any writable domain controller to which the zone is replicated. This builds redundancy into the DNS infrastructure. In addition, Multimaster updates are particularly important in organizations that use dynamic update zones and have locations that are distributed geographically. Clients can update their DNS records without having to connect to a potentially geographically-distant primary server.
- **Replication of DNS zone data by using AD DS replication.** One of the characteristics of Active Directory replication is attribute-level replication in which only changed attributes are replicated. An Active Directory–integrated zone can thus avoid replicating the entire zone file as in traditional DNS zone transfer models.
- **Secure dynamic updates.** An Active Directory–integrated zone can enforce secure dynamic updates.
- **Granular security.** As with other Active Directory objects, an Active Directory–integrated zone enables you to delegate administration of zones, domains, and resource records by modifying the access control list (ACL) on the zone.

Question: Can you think of any disadvantages to storing DNS information in AD DS?

Benefits of an Active Directory–integrated zone:

- Allows multimaster writes to zone
- Replicates DNS zone information by using AD DS replication
 - Leverages efficient replication topology
 - Uses efficient incremental updates for Active Directory replication processes
- Enables secure dynamic updates
- Delegates zones, domains, resource records for increased security

Examples of contoso.com zones include:

- hqdc01
- filesvr01
- desktop101

Demonstration: Creating an Active Directory–Integrated Zone

To create an Active Directory integrated zone, you must install a DNS server on a domain controller. All changes in an Active Directory–integrated zone replicate to the other DNS servers that are on domain controllers through the AD DS multimaster replication model.

In this demonstration, you will see how to:

- Promote a server as a domain controller.
- Create an Active Directory–integrated zone.
- Create a record.
- Verify replication to a second DNS server.

Demonstration Steps

Promote a server as a domain controller

1. Install the AD DS server role on LON-SVR1.
2. Start the Active Directory Domain Services Configuration Wizard.
3. Install the DNS Server service.

Create an Active Directory–integrated zone

1. On LON-DC1, open the DNS Manager console.
2. Start the New Zone Wizard.
3. Create a new Active Directory–integrated forward lookup zone named **Contoso.com** that allows only secure dynamic updates.
4. Review the records in the Contoso.com zone.

Create a record

- Create a **New Host** record in Contoso.com zone named **www**, and have it point to **172.16.0.100**.

Verify replication to a second DNS server

- Verify that new record is replicating to the **LON-SVR1** DNS server.

Lab: Implementing DNS

Scenario

Your manager has asked you to configure the domain controller in the branch office as a DNS server. You have also been asked to create some new host records to support a new app that is being installed. Finally, you need to configure forwarding on the DNS server in the branch office to support Internet name resolution.

Objectives

After completing this lab, you should be able to:

- Install and configure DNS.
- Create host records in DNS.
- Manage the DNS server cache.

Lab Setup

Estimated Time: 40 minutes

Virtual machines	20410C-LON-DC1 20410C-LON-SVR1 20410C-LON-CL1
User name	Adatum\Administrator
Password	Pa\$\$w0rd

For this lab, you will use the available virtual machine environment. Before beginning 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 **20410C-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in using the following credentials:
 - User name: **Administrator**
 - Password: **Pa\$\$w0rd**
 - Domain: **Adatum**
5. Repeat steps 2 to 4 for **20410C-LON-SVR1** and **20410C-LON-CL1**.

Exercise 1: Installing and Configuring DNS

Scenario

Contoso is a partner organization working closely with users in the new branch office. In order to support name resolution between A Datum's branch office and Contoso, you decide to enable DNS forwarding between the two DNS domains.

As part of configuring the infrastructure for the new branch office, you must configure a DNS server that provides name resolution for the branch office. This includes the forwarding for Contoso.com

The DNS server in the branch office will be a domain controller. The Active Directory integrated zones required to support logons will be replicated automatically to the branch office.

The main tasks for this exercise are as follows:

1. Configure LON-SVR1 as a domain controller without installing the Domain Name System (DNS) server role.
2. Create and configure Contoso.com zone on LON-DC1.
3. Review configuration settings on the existing DNS server to confirm root hints.
4. Add the DNS server role for the branch office on the domain controller.
5. Verify replication of the Adatum.com Active Directory®–integrated zone.
6. Use Windows PowerShell commands to test non-local resolution.
7. Configure Internet name resolution to forward to the head office.
8. Use Windows PowerShell to confirm name resolution.

► **Task 1: Configure LON-SVR1 as a domain controller without installing the Domain Name System (DNS) server role**

1. Use **Add roles and features** in Server Manager to add the **Active Directory Domain Services** role to LON-SVR1.
2. Start the Add Roles and Features Wizard to promote LON-SVR1 to domain controller.
3. Choose to add LON-SVR1 as an additional domain controller in **Adatum.com** domain.
4. Do not install the DNS server.

► **Task 2: Create and configure Contoso.com zone on LON-DC1**

1. On the LON-DC1 virtual machine, open the DNS Manager console.
2. Create new Forward Lookup Zone with the following parameters:
 - Zone name: **contoso.com**
 - Zone type: **Primary Zone**
 - Store the zone in Active Directory: **No**

► **Task 3: Review configuration settings on the existing DNS server to confirm root hints**

1. In DNS Manager on LON-DC1, open the **Properties** dialog box for LON-DC1.
2. Review root hints and forwarder configuration.

► **Task 4: Add the DNS server role for the branch office on the domain controller**

- Use Server Manager to add the DNS Server role to LON-SVR1.

► **Task 5: Verify replication of the Adatum.com Active Directory®–integrated zone**

1. On LON-SVR1, open the DNS Manager console.
2. Expand **Forward Lookup Zones** and verify that both the **Adatum.com** and **_msdcs.Adatum.com** zones are replicated.
3. If you do not see these zones, open Active Directory Sites and Services, force replication between **LON-DC1** and **LON-SVR1**, and then repeat steps 1 and 2.

► **Task 6: Use Windows PowerShell commands to test non-local resolution**

1. On LON-SVR1, make **127.0.0.1** the **preferred DNS server** for LON-SVR1 using the following **Windows PowerShell** cmdlet, where *X* is the Interface Index number, which you can find in the Get-DnsClient cmdlet:

```
Set-DnsClientServerAddress -InterfaceIndex X -ServerAddress 0.0.0.0
```

2. Open a Windows PowerShell window on **LON-SVR1**, and try to resolve **www.contoso.com** by using the **Resolve-DNSName** cmdlet.

You receive a negative reply (this is expected).

► **Task 7: Configure Internet name resolution to forward to the head office**

1. Type the following cmdlet, and then press Enter:

```
Set-DnsServerForwarder -IPAddress '172.16.0.10' -PassThru
```

2. Type the following cmdlet, and then press Enter:

```
Restart-Computer
```

► **Task 8: Use Windows PowerShell to confirm name resolution**

1. Sign in to **LON-SVR1** as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. In Windows PowerShell, type the following cmdlet, and then press Enter:

```
nslookupwww.contoso.com
```

You should get a reply and an IP address.

Results: After completing this exercise, you should have installed and configured DNS on 20410C-LON-SVR1.

Exercise 2: Creating Host Records in DNS

Scenario

Several new web-based apps are being implemented in the A. Datum head office. Each app requires that you configure a host record in DNS. You have been asked to create the new host records for these apps.

The main tasks for this exercise are as follows:

1. Configure a client to use LON-SVR1 as a DNS server.
2. Create several host records for web apps in the Adatum.com domain.
3. Verify replication of new records to LON-SVR1.
4. Use the ping command to locate new records from LON-CL1.

► **Task 1: Configure a client to use LON-SVR1 as a DNS server**

1. Sign in to LON-CL1 as **Adatum\Administrator** using the password **Pa\$\$w0rd**.
2. Open **Control Panel**.

3. Open the **Properties** dialog box for the **Ethernet** adapter.
4. Configure the **preferred DNS server** to be **172.16.0.11**.

► **Task 2: Create several host records for web apps in the Adatum.com domain**

1. On LON-DC1, open **DNS Manager**.
2. Navigate to the **Adatum.com** forward lookup zone.
3. Create a new record named **www** with the IP address **172.16.0.200**.
4. Create a new record named **ftp** with IP address **172.16.0.201**.

► **Task 3: Verify replication of new records to LON-SVR1**

1. On LON-SVR1, open **DNS Manager**.
2. Navigate to the **Adatum.com** forward lookup zone.
3. Ensure that records **www** and **ftp** display.



Note: If the **www** and **ftp** resource records do not display within several minutes, refresh the Adatum.com zone.

► **Task 4: Use the ping command to locate new records from LON-CL1**

1. On LON-CL1, open a Command Prompt window.
2. Ping **www.adatum.com**. Ensure that ping resolves this name to **172.16.0.200**.
3. Ping **ftp.adatum.com** and ensure that it resolves to **172.16.0.201**.

Results: After completing this exercise, you should have configured DNS records.

Exercise 3: Managing the DNS Server Cache

Scenario

After you changed some host records in zones configured on LON-DC1, you noticed that clients that use LON-SVR1 as their DNS server were still receiving old IP addresses during the name resolving process. You want to determine which component is caching this data.

The main tasks for this exercise are as follows:

1. Use the ping command to locate an Internet record from LON-CL1.
2. Update an Internet record to point to the LON-DC1 IP address.
3. Examine the content of the DNS cache.
4. Clear the cache, and retry the ping command.

► **Task 1: Use the ping command to locate an Internet record from LON-CL1**

1. On LON-CL1, in the Command Prompt window, use **ping** to locate **www.contoso.com**.
2. Ensure that the name resolves to an IP address, and then document the IP address.

► **Task 2: Update an Internet record to point to the LON-DC1 IP address**

1. On LON-DC1, open the **DNS Manager** console.
2. Navigate to the **contoso.com** forward lookup zone.
3. Change the IP address for the record **www** to **172.16.0.10**.
4. From LON-CL1, ping **www.contoso.com**.

Note that this record is still resolved with the old IP.

► **Task 3: Examine the content of the DNS cache**

1. On LON-SVR1, in the **DNS Manager** console, enable **Advanced View**.
2. Browse the content of the **Cached Lookups** container for the **com** namespace and note the IP address for **www** record.
3. On LON-CL1, at a command prompt, type **ipconfig /displaydns**.
4. Examine the cached content and notice the IP address for the **www** record.

► **Task 4: Clear the cache, and retry the ping command**

1. Clear the cache on the LON-SVR1 DNS server, by using the **Clear-DNSServerCache** cmdlet.
2. Retry the ping to **www.contoso.com** on LON-CL1.
The result still returns the old IP address.
3. Clear the client resolver cache on LON-CL1 by typing **ipconfig /flushdns** in the Command Prompt window.
4. On **LON-CL1**, retry ping to **www.contoso.com**. (The result should work.)

Results: After completing this exercise, you should have examined the DNS server cache.

► **Prepare for the next module**

After you finish the lab, revert the virtual machines to their initial state.

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20410C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20410C-LON-SVR1** and **20410C-LON-CL1**.

Lab Review Questions

Question: Can you install the DNS server role on a server that is not a domain controller? If yes, are there any limitations?

Question: What is the most common way to carry out Internet name resolution on a local DNS?

Question: How can you browse the content of the DNS resolver cache on a DNS server?

Module Review and Takeaways

Module Review Questions

Question: You are troubleshooting DNS name resolution from a client computer. What must you remember to do before each test?

Question: You are deploying DNS servers into an Active Directory domain, and your customer requires that the infrastructure is resistant to single points of failure. What must you consider when planning the DNS configuration?

Question: What benefits do you realize by using forwarders?

Best Practices

When implementing DNS, use the following best practices:

- Always use host names instead of NetBIOS names.
- Use forwarders rather than root hints.
- Be aware of potential caching issues when troubleshooting name resolution.
- Use Active Directory–integrated zones instead of primary and secondary zones.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Clients sometimes cache invalid DNS records.	
DNS Server performs slowly.	

Tools

Tool	Use	Where to find it
DNS Manager console	Manage DNS server role	Administrative Tools
Nslookup	Troubleshoot DNS	Command-line tool
Ipconfig	Troubleshoot DNS	Command-line tool
Windows PowerShell cmdlets	Manage and troubleshoot DNS	Windows PowerShell

Module 8

Implementing IPv6

Contents:

Module Overview	8-1
Lesson 1: Overview of IPv6	8-2
Lesson 2: IPv6 Addressing	8-7
Lesson 3: Coexistence with IPv4	8-13
Lesson 4: IPv6 Transition Technologies	8-17
Lab: Implementing IPv6	8-22
Module Review and Takeaways	8-27

Module Overview

IPv6 is a technology that helps the Internet support a growing user base and an increasingly large number of IP-enabled devices. IPv4 has been the underlying Internet protocol for almost thirty years. Its robustness, scalability, and limited feature set is now challenged by the growing need for new IP addresses. This is due in large part to the rapid growth of new network-aware devices.

This module discusses the features and benefits of IPv6, how IPv6 affects IPv4 networks, and how to integrate IPv6 into IPv4 networks by using various transition technologies.

Objectives

After completing this module, you will be able to:

- Describe the features and benefits of IPv6.
- Describe IPv6 addressing.
- Describe IPv6 coexistence with IPv4.
- Describe IPv6 transition technologies.

Lesson 1

Overview of IPv6

IPv6 has been included with Windows® client operating systems and servers beginning with Windows Server® 2008. The use of IPv6 is becoming more common on corporate networks and on the Internet.

It is important for you to understand how this technology affects current networks, and how to integrate IPv6 into those networks. This lesson discusses the benefits of IPv6, and how it differs from IPv4.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the benefits of IPv6.
- Describe the differences between IPv4 and IPv6.
- Describe the IPv6 address format.

Benefits of IPv6

IPv6 support is included in Windows Server 2012 and Windows 8. The following list of benefits describes why IPv6 is being implemented.

Larger Address Space

The IPv6 address space is 128-bit, which is much larger than the 32-bit address space in IPv4. A 32-bit address space has 2^{32} or 4,294,967,296 possible addresses; a 128-bit address space has 2^{128} or 340,282,366,920,938,463,463,374,607,431,768,211,456 (or 3.4×10^{38} or 340 undecillion) possible addresses. As the Internet continues to grow, IPv6 provides for the required larger address space.

Benefits of IPv6 include:

- Larger address space
- Hierarchical addressing and routing infrastructure
- Stateless and stateful address configuration
- Required support for Internet Protocol security (IPsec)
- End-to-end communication
- Required support for Quality of Service (QoS)
- Improved support for single-subnet environments
- Extensibility

Hierarchical Addressing and Routing Infrastructure

The public IPv6 address space is allocated more efficiently than it is for IPv4. IPv4 addresses are not all allocated in geographical blocks, but IPv6 public addresses are. This means that even though there are many more addresses, Internet routers can process data much more efficiently because of address optimization.

Stateless and Stateful Address Configuration

IPv6 has auto-configure capability without Dynamic Host Configuration Protocol (DHCP), and it can discover router information so that hosts can access the Internet. This is referred to as a *stateless* address configuration. A *stateful* address configuration is when you use the DHCPv6 protocol. This provides network administrators with flexibility in how IPv6 addresses and configuration information is distributed to clients.

Required Support for IPsec

The IPv6 standards require support for the Authentication Header (AH) and encapsulating security payload (ESP) headers that are defined by Internet Protocol security (IPsec). Although support for specific IPsec authentication methods and cryptographic algorithms are not specified, IPsec is defined from the start as the way to protect IPv6 packets. This guarantees the availability of IPsec on all IPv6 hosts. IPsec support was not required for IPv4 hosts, but was commonly implemented.

End-to-End Communication

One of the design goals for IPv6 is to provide sufficient address space so that you do not have to use translation mechanisms such as Network Address Translation (NAT). This simplifies communication because IPv6 hosts can communicate directly with each other over the Internet. This also simplifies support for apps such as video conferencing and other peer-to-peer apps. However, many organizations may choose to continue using translation mechanisms as a security measure.

Required Support for Quality of Service

An IPv6 packet contains a Quality of Service (QoS) field that specifies how fast the packet should be processed. This enables IPv6 packet traffic to be assigned a priority. For example, when you are streaming video traffic, it is critical that the packets arrive in a timely manner. You can set the QoS field to ensure that network devices recognize that the packet delivery is time sensitive. Support for QoS was optional for IPv4 hosts.

Improved Support for Single-Subnet Environments

All IPv6 hosts are configured automatically with a link-local address that allows the host to communicate on the local subnet. However, like Automatic Private IP Addressing (APIPA), which was optionally implemented in IPv4 environments, computers are not configured automatically with a default gateway or Domain Name System (DNS) server.

Extensibility

IPv6 has been designed so that developers can extend it with much fewer constraints than IPv4. As a network administrator, you will not be extending IPv6, but programs that you purchase may take advantage of this to enhance IPv6 functionality.

Differences Between IPv4 and IPv6

When the IPv4 address space was designed, it was unimaginable that it could ever be exhausted. However, because of changes in technology and an allocation practice that did not anticipate the growth of number of Internet hosts, it was clear by 1992 that a replacement would be necessary.

When the IPv6 address space was designed, the addresses were made 128 bits long so that the address space could be subdivided into hierarchical routing domains that reflect modern-day Internet topology. With 128 bits, there are enough bits to create multiple levels of hierarchy, and there is flexibility for designing hierarchical addressing and routing. These features are currently lacking on the IPv4-based Internet.

Feature	IPv4	IPv6
Fragmentation	Performed by routers and sending host	Performed only by sending host
Address Resolution	Broadcast ARP Request frames	Multicast Neighbor Solicitation messages
Manage multicast group membership	IGMP	Multicast listener discovery
Router Discovery	ICMP Router Discovery (optional)	ICMPv6 Router Solicitation and Router Advertisement (required)
DNS host records	A records	AAAA records
DNS reverse lookup zones	IN-ADDR.ARPA	IP6.ARPA
Minimum packet size	576 bytes	1280 bytes

Hexadecimal Numbering System (Base 16)

In the hexadecimal numbering system, some letters represent numbers; this is because there must be 16 unique symbols for each position. Because 10 symbols (0 through 9) already exist, there must be six new symbols for the hexadecimal system; hence, the letters A through F are used. The hexadecimal number 10 is equal to the decimal number 16.



Note: You can use the Calculator app included with Windows Server 2012 to convert between binary, decimal, and hexadecimal numbers.

To convert an IPv6 binary address that is 128 bits long, you break it into 8 blocks of 16 bits. You then convert each of these 8 blocks of 16 bits into 4 hexadecimal characters. For each of the blocks, you evaluate 4 bits at a time. You should number each section of 4 binary numbers 1, 2, 4, and 8, starting from the right and moving left. That is:

- The first bit [001**0**] is assigned the value of 1.
- The second bit [00**1**0] is assigned the value of 2.
- The third bit [0**0**10] is assigned the value of 4.
- The fourth bit [**0**010] is assigned the value of 8.

To calculate the hexadecimal value for this section of 4 bits, add up the value of each bit that is set to 1. In the example of 0010, the only bit that is set to 1 is the bit assigned the value 2. The rest are set to 0. Therefore, the hexadecimal value of this section of 4 bits is 2.

Converting from Binary to Hexadecimal

The following table describes converting 8 bits of binary into hexadecimal for the binary number [0010][1111]:

Binary	0010	1111
Values of each binary position	8421	8421
Adding values where the bit is 1	0+0+2+0=2	8+4+2+1=15 or hexadecimal F

The following example is a single IPv6 address in binary form. Note that the binary representation of the IP address is quite long. The following two lines of binary numbers represent one IP address:

```
00100000000000010000110110111000000000000000000010111100111011
000000101010101000000000111111111111110001010001001110001011010
```

The 128-bit address is now divided along 16-bit boundaries (8 blocks of 16 bits):

```
0010000000000001 0000110110111000 0000000000000000 0010111100111011
0000001010101010 0000000011111111 1111111000101000 1001110001011010
```

Each block is further broken into sections of 4 bits. The following table shows the binary and corresponding hexadecimal values for each section of 4 bits:

Binary	Hexadecimal
[0010][0000][0000][0001]	[2][0][0][1]
[0000][1101][1011][1000]	[0][D][B][8]
[0000][0000][0000][0000]	[0][0][0][0]
[0010][1111][0011][1011]	[2][F][3][B]
[0000][0010][1010][1010]	[0][2][A][A]
[0000][0000][1111][1111]	[0][0][F][F]
[1111][1110][0010][1000]	[F][E][2][8]
[1001][1100][0101][1010]	[9][C][5][A]

Each 16-bit block is expressed as 4 hexadecimal characters, and is then delimited with colons. The result is as follows:

2001:0DB8:0000:2F3B:02AA:00FF:FE28:9C5A

You can simplify IPv6 representation further by removing the leading zeros within each 16-bit block. However, each block must have at least a single digit. With leading zero suppression, the address representation becomes the following:

2001:DB8:0:2F3B:2AA:FF:FE28:9C5A

Compressing Zeros

When multiple contiguous zero blocks occur, you can compress these and represent them in the address as a double colon (::); this further simplifies the IPV6 notation. The computer recognizes "::" and substitutes it with the number of blocks necessary to make the appropriate IPv6 address.

In the following example, the address is expressed using zero compression:

2001:DB8::2F3B:2AA:FF:FE28:9C5A

To determine how many 0 bits are represented by the "::", you count the number of blocks in the compressed address, subtract this number from 8, and then multiply the result by 16. Using the previous example, there are 7 blocks. Subtract 7 from 8, and then multiply the result (1) by 16. Thus, there are 16 bits or 16 zeros in the address where the double colon is located.

You can use zero compression only once in a given address. If you use it twice or more, then there is no way to show how many 0 bits are represented by each instance of the double colon (::).

To convert an address into binary, use the reverse of the method described previously:

1. Add in 0s using zero compression.
2. Add leading 0s.
3. Convert each bit that is set to 1 (one) into its binary equivalent.

Lesson 2

IPv6 Addressing

An essential part of working with IPv6 is understanding the different address types and when they are used. This helps you understand the overall communication process between IPv6 hosts and perform troubleshooting. You also need to understand the processes available for configuring a host with an IPv6 address to ensure that hosts are configured properly.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the structure of IPv6 addresses.
- Describe the structure of global unicast addresses.
- Describe unique local unicast addresses.
- Describe link-local unicast addresses and zone IDs.
- Describe address autoconfiguration for IPv6.
- Explain how to configure IPv6 client settings on a network host.

IPv6 Address Structure

Each IPv6 address is 128 bits long. The prefix is the part of the address that indicates the bits that have fixed values, or that are the subnet prefix's bits. This is equivalent to the network ID for IPv4 addresses.

Prefixes for IPv6 subnets, routes, and address ranges are expressed in the same way as IPv4 Classless Interdomain Routing (CIDR) notations. An IPv6 prefix is written in address/prefix-length notation. For example, 2001:DB8::/48 and 2001:DB8:0:2F3B::/64 are IPv6 address prefixes.

- The number of network bits is defined by the prefix
- Each host has 64-bits allocated to the interface identifier

Type of address	IPv4 address	IPv6 address
Unspecified	0.0.0.0	::
Loopback	127.0.0.1	::1
Autoconfigured	169.254.0.0/16	FE80::/64
Broadcast	255.255.255.255	Uses multicasts instead
Multicast	224.0.0.0/4	FF00::/8



Note: IPv6 uses prefixes instead of a subnet mask.

When a unicast IPv6 address is assigned to a host, the prefix is 64 bits long. The remaining 64 bits are allocated to the interface identifier, which uniquely identifies the host on that network. The interface identifier can be either randomly generated, assigned by DHCPv6, or based on the media access control (MAC) address of the network. By default, the host bits are generated randomly unless assigned by DHCPv6.



Note: The routes on an IPv6 router have varying prefix sizes that are determined by the size of the network.

IPv6 Equivalents to IPv4 Special Addresses

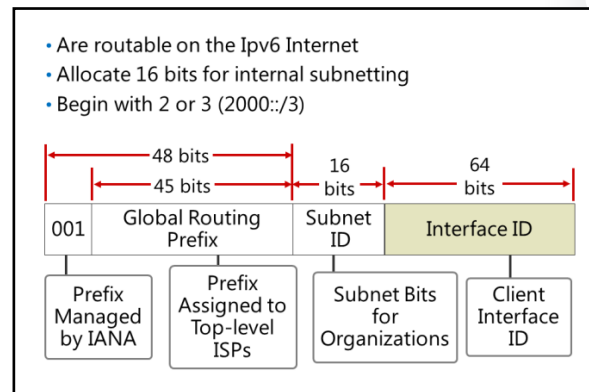
The following table shows IPv6 equivalents to some common IPv4 addresses.

	IPv4 Address	IPv6 Address
Unspecified address	0.0.0.0	::
Loopback address	127.0.0.1	::1
Autoconfigured addresses	169.254.0.0/16	FE80::/64
Broadcast address	255.255.255.255	Uses multicasts instead
Multicast addresses	224.0.0.0/4	FF00::/8

Global Unicast Addresses

Global unicast addresses are equivalent to public IPv4 addresses that are available from an Internet Service Provider (ISP). They are routable and reachable globally on the IPv6 portion of the Internet. Unlike the limited number of internet-addressable IPv4 addresses that remain, there are many global unicast addresses available for use.

The global unicast address space is designed to allow each ISP customer to obtain a large number of IPv6 addresses. The first 48 bits are used to identify the customer site. The next 16 bits are allocated for the customer to perform subnetting within their own network.



Note: The network 2001:0db8::/32 is reserved for documentation and is not routable.

The structure of a global unicast address is as follows:


- Fixed portion set to 001. The three high-order bits are set to 001. The address prefix for currently assigned global addresses is 2000::/3. Therefore, all global unicast addresses begin with either 2 or 3.
- Global routing prefix. This field identifies the global routing prefix for a specific organization's site. The combination of the three fixed bits and the 45-bit global routing prefix is used to create a 48-bit site prefix, which is assigned to an organization's individual site. Once the assignment occurs, routers on the IPv6 Internet then forward IPv6 traffic that matches the 48-bit prefix to the routers of the organization's site.
- Subnet ID. The Subnet ID is used within an organization's site to identify subnets. This field is 16 bits long. The organization's site can use these 16 bits within its site to create 65,536 subnets, or multiple levels of addressing hierarchy, and an efficient routing infrastructure.
- Interface ID. The Interface ID identifies the interface on a specific subnet within the site. This field is 64 bits long. This is either randomly generated or assigned by DHCPv6. In the past, the Interface ID was based on the MAC address of the network interface card to which the address was bound.

Unique Local Unicast Addresses

Unique local unicast addresses are the IPv6 equivalent of IPv4 private addresses. These addresses are routable within an organization, but not on the Internet.

IPv4 private IP addresses were a relatively small part of the overall IPv4 address space, and many companies used the same address space. This caused problems when separate organizations tried to communicate directly. It also caused problems when merging the networks of two organizations—possibly following a merger or a buyout.

- Are equivalent to IPv4 private addresses
- Require the organization ID to be randomly generated
- Allocates 16 bits for internal subnetting



To avoid the duplication problems experienced with IPv4 private addresses, the IPv6 unique local address structure allocates 40-bits to an organization identifier. The 40-bit organization identifier is randomly generated. The likelihood of two randomly generated 40-bit identifiers being the same is very small. This ensures that each organization has a unique address space.

The first 7 bits of the organization identifier have the fixed binary value of 1111110. All unique local addresses have the address prefix of FC00::/7. The Local (L) flag (the 8th bit) is set to 1 to indicate a local address. An L flag value set to 0 has not yet been defined. Therefore, unique local addresses with the L flag set to 1 have the address prefix of FD00::/8.


Link-Local Unicast Addresses

All IPv6 hosts have a link-local address that is used for communication only on the local subnet. The link-local address is generated automatically, and is non-routable. In this way, link-local addresses are similar to IPv4 APIPA addresses. However, a link-local address is an essential part of IPv6 communication.

Link-local addresses are used for communication in many scenarios where IPv4 would have used broadcasts. For example, link-local addresses are used when communicating with a DHCPv6 server. Link-local addresses are also used for neighbor discovery, which is the IPv6 equivalent of ARP in IPv4.

- Are automatically generated on all IPv6 hosts
- Are similar to IPv4 Automatic Private IP Addressing (APIPA) addresses
- Are sometimes used in place of broadcast messages
- Include a zone ID that identifies the interface

Examples: fe80::2b0:d0ff:fee9:4143%3
fe80::94bd:21cf:4080:e612%2



The prefix for link-local addresses is always FE80::/64. The final 64 bits are the interface identifier.

Zone ID

Regardless of the number of network interfaces in the host, each IPv6 host has a single link-local address. If the host has multiple network interfaces, the same link-local address is reused on each interface. To make it possible for hosts to identify link-local communication on each unique network interface, a zone ID is added to the link-local address.

A zone ID is used in the following format:

Address%zone_ID

Each sending host determines the zone ID that it will associate with each interface. There is no negotiation of zone ID between hosts. For example, on the same network, host A might use 3 for the zone ID on its interface, and host B might use 6 for the zone ID on its interface.

Each interface in a Windows-based host is assigned a unique interface index, which is an integer. In addition to physical network cards, interfaces also include loopback and tunnel interfaces. Windows-based IPv6 hosts use the interface index of an interface as the zone ID for that interface.

In the following example, the interface ID for the network interface is 3.

```
fe80::2b0:d0ff:fee9:4143%3
```

Autoconfiguring IPv6 Addresses

In most cases, you will use autoconfiguration to provide IPv6 hosts with an IPv6 address. Unlike IPv4 which uses primarily DHCP servers to provide addressing information, IPv6 also uses routers as part of the autoconfiguration process. The routers can provide the network address and a default gateway to clients in Router Advertisement messages.

Types of Autoconfiguration

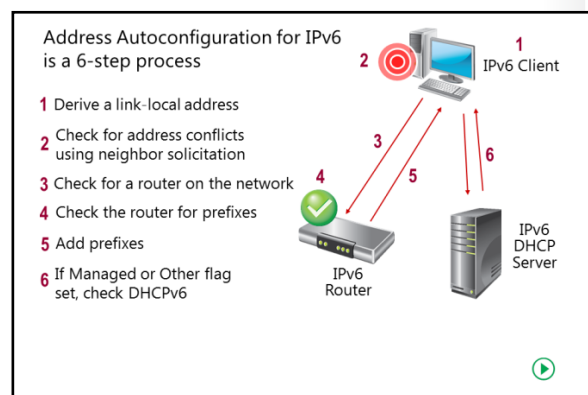
Types of autoconfiguration include:

- Stateless. With stateless autoconfiguration, address configuration is based on the receipt of Router Advertisement messages only. Stateless autoconfiguration includes a router prefix, but does not include additional configuration options such as DNS servers.
- Stateful. With stateful autoconfiguration, address configuration is based on the use of a stateful address configuration protocol such as DHCPv6 to obtain addresses and other configuration options. A host uses stateful address configuration when:
 - It receives instructions to do so in router advertisement messages.
 - There are no routers present on the local link.
- Both. With both, configuration is based on both receipt of router advertisement messages, and on DHCPv6.

Stateful Configuration

With stateful configuration, organizations can control how IPv6 addresses are assigned using DHCPv6. If there are any specific scope options that you need to configure—such as the IPv6 addresses of DNS servers—then a DHCPv6 server is necessary.

When IPv6 attempts to communicate with a DHCPv6 server, it uses multicast IPv6 addresses. This is different from IPv4, which uses broadcast IPv4 addresses.



Autoconfigured Address States

During autoconfiguration, the IPv6 address of a host goes through several states that define the life cycle of the IPv6 address. Autoconfigured addresses are in one or more of the following states:

- **Tentative.** In the tentative state, verification is occurring to determine if the address is unique. Duplicate address detection performs verification. When an address is in the tentative state, a node cannot receive unicast traffic.
- **Valid.** In the valid state, the address has been verified as unique, and can send and receive unicast traffic.
- **Preferred.** In the preferred state, the address enables a node to send and receive unicast traffic to and from it.
- **Deprecated.** In a deprecated state, the address is valid, but its use is discouraged for new communication.
- **Invalid.** In the invalid state, the address no longer allows a node to send or receive unicast traffic.

The Autoconfiguration Process

The process of IPv6 autoconfiguration follows these six general steps:

1. The client derives a link local address.
2. The client checks for address conflicts using neighbor solicitation and verifies that the link local address is unique.
3. The client checks for routers on the network.
4. The client checks to see which prefixes are configured on the router.
5. The client adds the prefixes locally.
6. If the Managed or Other flag is set then the client checks for DHCPv6 to obtain other configuration information.

Demonstration: Configuring IPv6 Client Settings

In most cases, IPv6 is configured dynamically by using DHCPv6 or router advertisements. However, you can also configure IPv6 manually with a static IPv6 address. The process for configuring IPv6 is similar to the process for configuring IPv4.

In this demonstration, you will see how to:

- View IPv6 configuration by using **IPconfig** and **Get-NetIPAddress**.
- Configure IPv6 on a domain controller and a server.
- Verify IPv6 communication is functional.

Demonstration Steps

View IPv6 configuration by using IPconfig and Get-NetIPAddress

1. Sign in to LON-DC1 and LON-SVR1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On LON-DC1, open a **Windows PowerShell®** prompt.
3. Use the **ipconfig** command to view the link-local IPv6 address on the Ethernet.
4. Use the **Get-NetIPAddress** cmdlet to view the network configuration and link-local IPv6 address on the **Ethernet** adapter.

Configure IPv6 on LON-DC1

1. On LON-DC1, use Server Manager to open the **Local Server Properties** dialog box, and then click **Ethernet**.
2. Open the **Internet Protocol Version 6 (TCP/IPv6) Properties** dialog box, and enter the following information:
 - **Use the following IPv6 address**
 - IPv6 address: **FD00:AAAA:BBBB:CCCC::A**
 - Subnet prefix length: **64**
 - **Use the following DNS server addresses**
 - Preferred DNS server: **::1**

Configure IPv6 on LON-SVR1

1. On LON-DC1, use Server Manager to open the **Local Server Properties** dialog box, and then click **Ethernet**.
2. Open the **Internet Protocol Version 6 (TCP/IPv6) Properties** dialog box, and enter the following:
 - **Use the following IPv6 address**
 - IPv6 address: **FD00:AAAA:BBBB:CCCC::15**
 - Subnet prefix length: **64**
 - **Use the following DNS server addresses**
 - Preferred DNS server: **FD00:AAAA:BBBB:CCCC::A**

Verify that IPv6 communication is functional

1. On LON-SVR1, open a **Windows PowerShell** prompt.
2. Use **ipconfig** to view the IPv6 address for the **Ethernet** adapter.
3. Use **ping -6** to test IPv6 communication with **LON-DC1**.
4. Use **ping -4** to test IPv4 communication with **LON-DC1**.
5. Use **Test-NetConnection** cmdlet to test IPv6 communication with IPv6 address of **LON-DC1 FD00:AAAA:BBBB:CCCC::**.

Lesson 3

Coexistence with IPv4

From its inception, IPv6 was designed for long-term coexistence with IPv4; in most cases your network will use both IPv4 and IPv6 for many years. Consequently, you need to understand how they coexist.

This lesson provides an overview of the technologies that support the two IP protocols' coexistence. This lesson also describes the different node types and IP stack implementations of IPv6. Finally, this lesson explains how DNS resolves names to IPv6 addresses and the various types of IPv6 transition technologies.

Lesson Objectives

After completing this lesson, you will be able to:


- Describe IP node types.
- Describe methods to provide coexistence for IPv4 and IPv6.
- Configure DNS to support IPv6.
- Explain IPv6 over IPv4 tunneling.

What Are Node Types?

When planning an IPv6 network, you should know what types of nodes or hosts are on the network. Describing the nodes in a specific way helps to define their capabilities on the network. Understanding the capabilities of each type of node is important if you use tunneling, because certain kinds of tunnels require specific node types. The descriptions for the various types of nodes are as follows:

- IPv4-only node. This is a node that implements only IPv4 (and has only IPv4 addresses) and does not support IPv6.
- IPv6-only node. This is a node that implements only IPv6 (and has only IPv6 addresses) and does not support IPv4. This node is able to communicate only with IPv6 nodes and program, and is not common today. However, it might become more prevalent as smaller devices (such as cellular phones and handheld computers), use the IPv6 protocol exclusively.
- IPv6/IPv4 node. This is a node that implements both IPv4 and IPv6. Windows Server 2008 and newer Windows Server operating systems, and Windows Vista and newer Windows client operating systems, use IPv4 and IPv6 by default.
- IPv4 node. This is a node that implements IPv4. It can be an IPv4-only node or an IPv6/IPv4 node.
- IPv6 node. This is a node that implements IPv6. It can be an IPv6-only node or an IPv6/IPv4 node.

Coexistence occurs when the largest number of nodes (IPv4 or IPv6 nodes) can communicate using an IPv4 infrastructure, an IPv6 infrastructure, or an infrastructure that is a combination of IPv4 and IPv6. You will achieve true migration when all IPv4 nodes are converted to IPv6-only nodes. However, for the foreseeable future, you can achieve practical migration when as many IPv4-only nodes as possible are converted to IPv6/IPv4 nodes. IPv4-only nodes can communicate with IPv6-only nodes only when you are using an IPv4-to-IPv6 proxy or translation gateway.



IPv4 and IPv6 Coexistence

Rather than replacing IPv4, most organizations add IPv6 to their existing IPv4 network. Starting with Windows Server 2008 and Windows Vista, Windows operating systems support the simultaneous use of IPv4 and IPv6 through a dual IP layer architecture. The Windows Server 2003 operating system uses less efficient dual-stack architecture.

Dual IP Layer Architecture

A dual IP layer architecture was implemented beginning with Windows Vista, and continuing through Windows Server 2012 and Windows 8.

This architecture contains both IPv4 and IPv6 Internet layers with a single implementation of transport layer protocols such as TCP and User Datagram Protocol (UDP). Dual stack makes it easier to migrate to IPv6, and there are fewer files to maintain to provide IPv6 connectivity. IPv6 is also available without adding any new protocols in the network-card configuration.

Windows Server 2012 uses a dual IP layer architecture that supports IPv4 and IPv6 in a single protocol stack

DNS records required for coexistence are:

- Host (A) resource records for IPv4 nodes
- IPv6 host (AAAA) resource records
- Reverse lookup pointer (PTR) resource records for IPv4 and IPv6 nodes

Dual Stack Architecture

Dual stack architecture contains both IPv4 and IPv6 Internet layers, and has separate protocol stacks that contain separate implementations of transport layer protocols, such as TCP and UDP. Tcpip6.sys, the IPv6 protocol driver in Windows Server 2003, contains a separate implementation of TCP and UDP.

DNS Infrastructure Requirements

Just as DNS is used as a supporting service on an IPv4 network, it is also used on an IPv6 network. When you add IPv6 to the network, you need to ensure that you add the records that are necessary to support IPv6 name-to-address and address-to-name resolution. The DNS records that are required for coexistence are:

- Host (A) resource records for IPv4 nodes
- IPv6 host (AAAA) resource records for IPv6 nodes
- Reverse lookup pointer (PTR) resource records for IPv4 and IPv6 nodes



Note: In most cases, the IPv6 host (AAAA) resource records that IPv6 nodes require are registered in DNS dynamically.

When a name can be resolved to both an IPv4 and IPv6 address, both addresses are returned to the client. The client then chooses which address to use based on prefix policies. In these prefix policies each prefix has a precedence level assigned to it. A higher precedence is preferred over a lower precedence.

The following table lists typical prefix policies for Windows Server 2012.

Prefix	Precedence	Label	Description
::1/128	50	0	IPv6 loopback
::/0	40	1	Default gateway
::ffff:0:0/96	10	4	IPv4 compatible address
2002::/16	7	2	6to4
2001::/32	5	5	Teredo
FC00::/7	3	13	Unique local
::/96	1	3	IPv4 compatible address (deprecated)
fec0::/10	1	11	Site local (deprecated)
3ffe::/16	1	12	6Bone (deprecated)



Note: You can view the prefix policies in Windows Server 2012 by using the Windows PowerShell **Get-NetPrefixPolicy** cmdlet.



Additional Reading: For more information about prefix policies, see Source and Destination Address Selection for IPv6 at <http://go.microsoft.com/fwlink/?LinkId=269711>.

Demonstration: Configuring DNS to Support IPv6

Similar to IPv4 nodes, IPv6 nodes use dynamic DNS automatically created host records. You can also manually create host records for IPv6 addresses. An IPv6 host (AAAA) resource record is a unique record type and different from an IPv4 host (A) resource record.

In this demonstration, you will see how to:

- Configure an IPv6 host (AAAA) resource record for an IPv6 address.
- Verify name resolution for an IPv6 host (AAAA) resource record.

Demonstration Steps

Configure an IPv6 host (AAAA) resource record

1. On LON-DC1, in Server Manager, open the **DNS** tool and browse to the **Adatum.com forward lookup zone**.
2. In DNS Manager, verify that IPv6 addresses have been registered dynamically for LON-DC1 and LON-SVR1.
3. Create a new host record in Adatum.com with the following settings:
 - Name: **WebApp**
 - IP address: **FD00:AAAA:BBBB:CCCC::A**

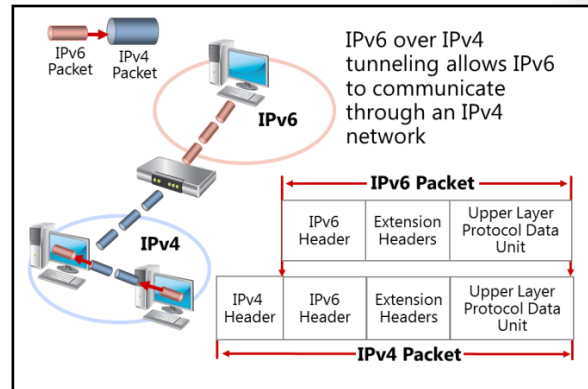
Verify name resolution for an IPv6 host (AAAA) resource record

1. On LON-SVR1, if necessary, open a **Windows PowerShell** prompt.
2. Use **ping** to test communication with **WebApp.adatum.com**.
3. Use the **Test-NetConnection** cmdlet to test communication with **WebApp.adatum.com**.

What Is IPv6 Over IPv4 Tunneling?

IPv6 over IPv4 tunneling is the encapsulation of IPv6 packets with an IPv4 header so that IPv6 packets can be sent over an IPv4-only infrastructure. Within the IPv4 header:

- The **IPv4 Protocol** field is set to **41** to indicate an encapsulated IPv6 packet.
- The **Source** and **Destination** fields are set to IPv4 addresses of the tunnel endpoints. You can configure tunnel endpoints manually as part of the tunnel interface, or they can be derived automatically.



Unlike tunneling for the Point-to-Point Tunneling Protocol (PPTP) and Layer Two Tunneling Protocol (L2TP), there is no exchange of messages for tunnel setup, maintenance, or termination. Additionally, IPv6 over IPv4 tunneling does not provide security for tunneled IPv6 packets. This means that when you use IPv6 tunneling, it does not need to establish a protected connection first.

You can manually configure IPv6 over IPv4 tunneling, or use automated technologies such as ISATAP, 6to4, or Teredo that implement IPv6 over IPv4 tunneling.

Lesson 4

IPv6 Transition Technologies

Transitioning from IPv4 to IPv6 requires coexistence between the two protocols. Too many programs, apps and services rely on IPv4 for it to be removed quickly. However, there are several technologies that aid transition by allowing communication between IPv4-only and IPv6-only hosts. There are also technologies that allow IPv6 communication over IPv4 networks.

This lesson provides information about Intra-Site Automatic Tunnel Addressing Protocol (ISATAP), 6to4, and Teredo, which help provide connectivity between IPv4 and IPv6 technology. This lesson also addresses PortProxy, which provides compatibility for apps.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe ISATAP.
- Describe 6to4.
- Describe Teredo.
- Describe PortProxy.
- Describe the transition process from IPv4 to IPv6.

What Is ISATAP?

ISATAP is an address-assignment technology that you can use to provide unicast IPv6 connectivity between IPv6/IPv4 hosts over an IPv4 intranet. IPv6 packets are tunneled in IPv4 packets for transmission over the network. Communication can occur directly between two ISATAP hosts on an IPv4 network, or communication can go through an ISATAP router if one network has only IPv6-only hosts.

ISATAP hosts do not require any manual configuration, and can create ISATAP addresses using standard address autoconfiguration mechanisms. Although the ISATAP component is enabled by default, it only assigns ISATAP-based addresses if it can resolve the ISATAP name on your network.

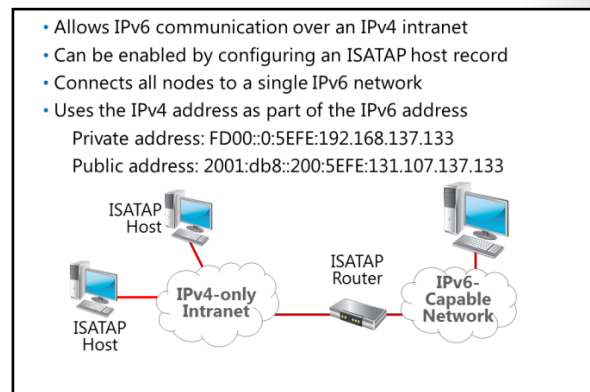
An ISATAP address that is based on a private IPv4 address is formatted like the following example:

[64-bit unicast prefix]:0:5EFE:w.x.y.z

An ISATAP address that is based on a public IPv4 address is formatted like the following example:

[64-bit unicast prefix]:200:5EFE:w.x.y.z

For example, FD00::5EFE:192.168.137.133 is an example of a private IPv4 address, and 2001:db8::200:5EFE:131.107.137.133 is an example of a public IPv4 address.



What Is an ISATAP Router?


If there are no IPv6-only hosts, then the ISATAP router advertises the IPv6 prefix that is used by ISATAP clients. The ISATAP interface on client computers is configured to use this prefix. When programs use the ISATAP interface to deliver data, the IPv6 packet is encapsulated in an IPv4 packet for delivery to the IPv4 address of the destination ISATAP host.

If there are IPv6-only hosts, then the ISATAP router also unpacks IPv6 packets. ISATAP hosts send packets to the IPv4 address of the ISATAP router. The ISATAP router then unpacks the IPv6 packets and sends them on to the IPv6-only network.

How to Enable ISATAP Tunneling


You can initiate ISATAP tunneling in many ways, but the simplest way is to configure an ISATAP host record in DNS that resolves to the IPv4 address of the ISATAP router. Windows hosts that can resolve this name automatically begin using the specified ISATAP router. By using this method, you can configure ISATAP for several computers simultaneously.

You can also define ISATAP name resolution in a hosts file, but this is not recommended because it is difficult to manage.

 **Note:** By default, DNS servers on Windows Server 2008 or newer Windows Server operating systems have a global query block list that prevents ISATAP resolution, even if the host record is created and properly configured. You need to remove ISATAP from the global query block list in DNS if you are using an ISATAP host record to configure ISATAP clients.

Other ways you can configure hosts with an ISATAP router are:

- Use the Windows PowerShell cmdlet **Set-NetIsatapConfiguration -Router x.x.x.x**.
- Use **Netsh Interface IPv6 ISATAP Set Router x.x.x.x**.
- Configure the **ISATAP Router Name** Group Policy setting.

 **Note:** All ISATAP nodes are connected to a single IPv6 subnet. This means that all ISATAP nodes are part of the same Active Directory® Domain Services (AD DS) site, which may not be desirable.

As such, you should use ISATAP only for limited testing. For intranet-wide deployment, you should instead deploy native IPv6 support.

What Is 6to4?


6to4 is a technology that you use to provide unicast IPv6 connectivity over the IPv4 Internet. You can use 6to4 to provide IPv6 connectivity between two IPv6 sites or between an IPv6 host and an IPv6 site. However, 6to4 is not suitable for scenarios that require NAT.

A 6to4 router provides a site with IPv6 connectivity over the IPv4 Internet. The 6to4 router has a public IPv4 address that is configured on the external interface, and a 6to4 IPv6 address that is configured on the internal interface. To

- Provides IPv6 connectivity over the IPv4 Internet
- Works between sites or from host to site
- Is not suitable for scenarios using NAT
- Uses the following network address format:
2002:WWXX:YYZZ:Subnet_ID::/64

To enable Windows Server 2012 as a 6to4 router:

- Enable ICS
- Use Windows PowerShell



configure client computers, the internal interface advertises the 6to4 network. Any client computer that begins to use the 6to4 network address is a 6to4 host. The 6to4 hosts in the site send 6to4 packets to the 6to4 router for delivery to other sites over the IPv4 Internet.

The IPv6 network address that is used for 6to4 is based on the IPv4 address of the external interface on an IPv6 router. The format of the IPv6 is 2002:WWXX:YZZZ:Subnet_ID:Interface_ID, where WWXX:YZZZ is the colon-hexadecimal representation of w.x.y.z, a public IPv4 address.

When a single host on the IPv4 Internet participates in 6to4, it is configured as a host/router. A 6to4 host/router does not perform routing for other hosts, but does generate its own IPv6 network used for 6to4.


Enabling 6to4 Router Functionality in Windows Operating Systems

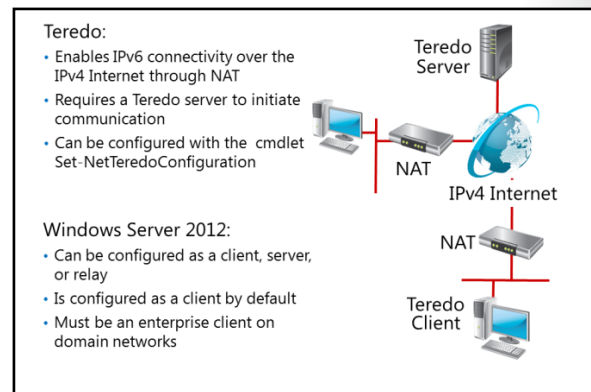
In most cases, you use existing network infrastructure components to act as a 6to4 router. However, you can configure Windows Server 2012 as a 6to4 router in the following ways:

- Enable Internet Connection Sharing. When you enable Internet Connection Sharing, Windows Server 2012 is configured automatically as a 6to4 router.
- Use Windows PowerShell. You can use the **Set-Net6to4Configuration** cmdlet to configure 6to4.


What Is Teredo?

Teredo is similar to 6to4 in that it allows you to tunnel IPv6 packets over the IPv4 Internet. However, Teredo functions correctly even when NAT is used for Internet connectivity. Teredo is required because many organizations use private IP addresses, which require NAT to access the Internet. If a NAT device can be configured as a 6to4 router, then Teredo is not required.

 **Note:** Teredo is only used if native IPv6, 6to4, or ISATAP do not provide connectivity.



IPv6 communication between two Teredo clients over the IPv4 Internet requires a Teredo server that is hosted on the IPv4 Internet. The Teredo server facilitates communication between the two Teredo clients by acting as a known central point for initiating communication. Typically, hosts behind a NAT device are allowed to initiate outbound communication, but are not allowed to accept inbound communication. To work around this problem, both Teredo clients initiate communication with the Teredo server. After connection is initiated with the Teredo server, and after the NAT device has allowed outbound communication, any further communication occurs directly between the two Teredo clients.

 **Note:** Several public Teredo servers are available for use on the Internet. Windows operating systems use the Microsoft-provided Teredo server at teredo.ipv6.microsoft.com by default.

Teredo can also facilitate communication with IPv6-only hosts on the IPv6 Internet by using a Teredo relay. The Teredo relay forwards packets from a Teredo client to the IPv6 Internet.

You can configure Windows Server 2012 as a Teredo client, Teredo relay, or Teredo server. To configure Teredo use the Windows PowerShell cmdlet **Set-NetTeredoConfiguration**. The default configuration for Teredo is as a client. When configured as a client, Teredo is disabled when attached to a domain network. To enable Teredo on a domain network, you must configure it as an enterprise client.

Teredo Address Structure

A Teredo address is a 128-bit IPv6 address, but it uses a different structure than typical unicast IPv6 addresses. The structure is as follows:

- 2001::/32 (32 bits). This is the Teredo-specific prefix that is used by all Teredo addresses.
- Teredo server IPv4 address (32 bits). This identifies the Teredo server.
- Options (16 bits). There are a number of options that describe the communication configuration, such as whether the client is behind NAT.
- Obscured external port (16 bits). This is the external port used for communication by the NAT device for this communication. It is obscured to prevent the NAT device from translating it.
- Obscured external IP address (32 bits). This is the external IP address of the NAT device. It is obscured to prevent the NAT device from translating it.

What Is PortProxy?

Application developers use specific network application programmer interfaces (APIs) to access network resources when they are writing apps. Modern APIs are able to use either IPv4 or IPv6, and leave the responsibility of choosing the IP version to the operating system. However, some older apps use APIs that can only use IPv4.

You use the PortProxy service to allow programs or apps that do not support IPv6 to communicate with IPv6 hosts. You enable PortProxy on the server where the program or app is running. Incoming IPv6 packets for the program or app are translated to IPv4, and then passed on to the program or app.

You can also use PortProxy as a proxy between IPv4-only and IPv6-only hosts. To do this, you must configure DNS to resolve the name of the remote host as the address of the PortProxy computer. For example, an IPv4-only host would resolve the name of an IPv6-only host as the IPv4 address of the PortProxy computer. Packets would then be sent to the PortProxy computer, which would then proxy them to the IPv6-only computer.

PortProxy has the following limitations:

- It is limited to TCP connections only. It cannot be used for programs or apps that use UDP.
- It cannot change address information that is embedded in the data portion of the packet. If the program or app (such as File Transfer Protocol (FTP)) embeds address information in the data portion, then it will not work.

You can configure PortProxy on Windows Server 2012 by using **netsh interface portproxy**. However, it is generally preferred to use a tunneling technology instead of PortProxy.

Use PortProxy to:

- Provide IPv6-only hosts with access to IPv4-only applications
- Provide access between IPv4-only and IPv6-only hosts

Limitations of PortProxy:

- Only TCP applications
- Cannot change embedded address information

Process for Transitioning to IPv6

The industry-wide migration from IPv4 to IPv6 is expected to take considerable time. This was taken into consideration when designing IPv6 and as a result, the transition plan for IPv6 is a multistep process that allows for extended coexistence.

To achieve the goal of an IPv6-only environment, use the following general guidelines:

- Upgrade your programs and apps to be independent of either IPv6 or IPv4. For example, you can change apps to use new Windows Sockets APIs so that name resolution, socket creation, and other functions are independent regardless of whether you are using IPv4 or IPv6.
- Upgrade routing infrastructure for native IPv6 routing. You must upgrade routers to support both native IPv6 routing and IPv6 routing protocols.
- Upgrade devices to support IPv6. The majority of current networking hardware supports IPv6, but many other types of devices do not. You need to verify that all network attached devices—such as printers and scanners—also support IPv6.
- Update the DNS infrastructure to support IPv6 address and pointer (PTR) resource records. You might have to upgrade the DNS infrastructure to support the new IPv6 host address (AAAA) resource records (required) and pointer (PTR) resource records in the IP6.ARPA reverse domain, but this is optional. Additionally, ensure that the DNS servers support both DNS traffic over IPv6, and DNS dynamic update for IPv6 host address (AAAA) resource records so that IPv6 hosts can register their names and IPv6 addresses automatically.
- Upgrade hosts to IPv6/IPv4 nodes. You must upgrade hosts to use both IPv4 and IPv6. This allows hosts to access both IPv4 and IPv6 resources during the migration process.

To transition from IPv4 to IPv6 you must:

- Update applications to support IPv6
- Update routing infrastructure to support IPv6
- Update devices to support IPv6
- Update DNS with records for IPv6
- Upgrade hosts to IPv4/IPv6 nodes

Most organizations will probably add IPv6 to an existing IPv4 environment and continue to have coexistence for an extended time. Many legacy programs, apps, and devices that do not support IPv6 are still in existence, and coexistence is much simpler than using transition technologies such as ISATAP. You should only remove IPv4 after resources that depend on it are either removed or updated to use IPv6.

IPv6 is enabled by default for Windows Vista and newer Windows client operating systems, and Windows Server 2008 and newer Windows Server operating systems. As a best practice, you should not disable IPv6 unless there is a technical reason to do so. Some features in Windows operating systems rely on IPv6.

Lab: Implementing IPv6

Scenario

The IT manager at A. Datum has been briefed by several program and app vendors about newly added support for IPv6 in their products. A. Datum does not have IPv6 support in place at this time. The IT manager would like you to configure a test lab that uses IPv6. As part of the test lab configuration, you also need to configure ISATAP to allow communication between an IPv4 network and an IPv6 network. This is the layout of the completed test environment.

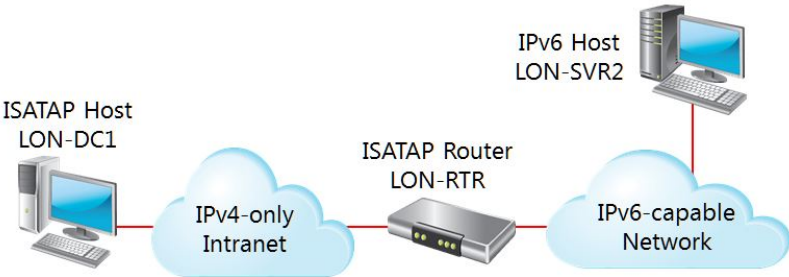


FIGURE 8.1: END RESULT OF THE LAB

Objectives

After completing this lab, you will be able to:

- Configure an IPv6 network.
- Configure an ISATAP router.

Lab Setup

Estimated time: 30 minutes

Virtual machines	20410C-LON-DC1 20410C-LON-RTR 20410C-LON-SVR2
User name	Adatum\Administrator
Password	Pa\$\$w0rd

For this lab, you will use the available virtual machine environment. Before beginning 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 **20410C-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in using the following credentials:
 - User name: **Administrator**
 - Password: **Pa\$\$w0rd**
 - Domain: **Adatum**
5. Repeat steps 2 to 4 for **20410C-LON-RTR** and **20410C-LON-SVR2**.

Exercise 1: Configuring an IPv6 Network

Scenario

For the first step in configuring the test lab, you need to configure LON-DC1 as an IPv4-only node, and LON-SVR2 as an IPv6-only node. You also need to configure LON-RTR to support IPv6 routing by adding a network to an interface on the IPv6 network, and by enabling router advertisements. The router advertisements allow the IPv6 clients on the IPv6 network to obtain the correct IPv6 network automatically through stateless configuration.

The main tasks for this exercise are as follows:

1. Verify IPv4 routing.
2. Disable IPv6 on LON-DC1.
3. Disable IPv4 on LON-SVR2.
4. Configure an IPv6 network on LON-RTR.
5. Verify IPv6 on LON-SVR2.

► Task 1: Verify IPv4 routing

1. On LON-SVR2, open a **Windows PowerShell** prompt.
2. Ping **LON-DC1** to verify that IPv4 is routing through LON-RTR.
3. Use the **ipconfig** command to verify that LON-SVR2 has only a link-local IPv6 address that cannot be routed.
4. Use the **Get-NetIPAddress** cmdlet to view the network configuration and link-local IPv6 address on the **Ethernet** adapter.

► Task 2: Disable IPv6 on LON-DC1

1. On LON-DC1, in Server Manager, select **Local Server**, and then select **172.16.0.10, IPv6 enabled** to open the Network Connections window.
2. In the Network Connections window, open the **Ethernet Properties** dialog box, and then disable IPv6 for **Ethernet** adapter to make LON-DC1 an IPv4-only host.

► Task 3: Disable IPv4 on LON-SVR2

1. On LON-SVR2, in Server Manager, select **Local Server**, and then select **10.10.0.11, IPv6 enabled** to open the Network Connections window.
2. In Network Connections window, open the **Ethernet Properties** dialog box, and then disable IPv4 for **Ethernet**, to make LON-SVR2 an IPv6-only host.

► Task 4: Configure an IPv6 network on LON-RTR

1. On LON-RTR, open **Windows PowerShell**.
2. Configure a network address that will be used on the IPv6 network by using the following Windows PowerShell **New-NetRoute** cmdlet to add an IPv6 network on **Ethernet 2** to the local routing table:

```
New-NetRoute -InterfaceAlias "Ethernet 2" -DestinationPrefix 2001:db8:0:1::/64
-Publish Yes
```

3. Allow clients to obtain the IPv6 network access automatically from LON-RTR by using the following **Set-NetIPInterface** cmdlet to enable router advertisements on Ethernet 2:

```
Set-NetIPInterface -InterfaceAlias "Ethernet 2" -AddressFamily IPv6 -Advertising Enabled
```

4. Use **ipconfig** to verify that Ethernet 2 has an IPv6 address on the **2001:db8:0:1::/64** network. This address is used for communication on the IPv6-only network.

► Task 5: Verify IPv6 on LON-SVR2

- On LON-SVR2, use **ipconfig** to verify that the Ethernet has an IPv6 address on the 2001:db8:0:1::/64 network.

The network address was obtained from the router through stateless configuration.

Results: After completing the exercise, you will have configured an IPv6-only network.

Exercise 2: Configuring an ISATAP Router

Scenario

After configuring the infrastructure for an IPv4-only network and an IPv6-only network, you need to configure LON-RTR as an ISATAP router to support communication between the IPv4-only nodes and the IPv6-only nodes.

To configure LON-RTR as an ISATAP router, you need to enable the IPv4 interface as the ISATAP router. Then you configure an IPv6 network on the ISATAP interface and enable advertising of the network route that includes that network. ISATAP clients will obtain the IPv6 network automatically from the advertisements.

To enable ISATAP automatically on clients, you need to create an ISATAP host record in DNS. Clients that can resolve this name automatically become ISATAP clients. To allow clients to resolve this name, you must remove ISATAP from the global query block list on the DNS server.

The main tasks for this exercise are as follows:

1. Add an ISATAP host record to DNS.
2. Enable the ISATAP router on LON-RTR.
3. Remove ISATAP from the Global Query Block List.
4. Enable LON-DC1 as an ISATAP client.
5. Test connectivity.

► Task 1: Add an ISATAP host record to DNS

1. On LON-DC1, in Server Manager, open the **DNS** tool.
2. Add an **ISATAP** host record in the Adatum.com domain that resolves to **172.16.0.1**. ISATAP clients resolve this host name to find the ISATAP router.

► Task 2: Enable the ISATAP router on LON-RTR

1. On LON-RTR, configure the IP address of the **Ethernet** adapter as the ISATAP router. Use the following **Set-NetIsatapConfiguration** cmdlet to enable ISATAP:

```
Set-NetIsatapConfiguration -Router 172.16.0.1
```

2. Use the following **Get-NetIPAddress** cmdlet to identify the interface index of the ISATAP interface with 172.16.0.1 in the link-local address.

```
Get-NetIPAddress | Format-Table InterfaceAlias,InterfaceIndex,IPv6Address
```

Record the Interface index here:	
----------------------------------	--

3. Use the **Get-NetIPInterface** cmdlet to verify the following on the ISATAP interface:
 - Forwarding is enabled
 - Advertising is disabled

```
Get-NetIPInterface -InterfaceIndex IndexYouRecorded -PolicyStore ActiveStore |  
Format-List
```

4. The ISATAP interface for an ISATAP router must have forwarding enabled and advertising enabled. Use the following **Set-NetIPInterface** cmdlet to enable router advertisements on the ISATAP interface:

```
Set-NetIPInterface -InterfaceIndex IndexYouRecorded -Advertising Enabled
```

5. Create a new IPv6 network that will be used for the ISATAP network. Use the following **New-NetRoute** cmdlet to configure a network route for the ISATAP interface:

```
New-NetRoute -InterfaceIndex IndexYouRecorded -DestinationPrefix 2001:db8:0:2::/64  
-Publish Yes
```

6. Use the following **Get-NetIPAddress** cmdlet to verify that the ISATAP interface has an IPv6 address on the 2001:db8:0:2::/64 network:

```
Get-NetIPAddress -InterfaceIndex IndexYouRecorded
```

► Task 3: Remove ISATAP from the Global Query Block List

1. On LON-DC1, in Windows PowerShell, run the following cmdlet:

```
dnscmd /config /globalqueryblocklist wpad
```

2. Restart the **DNS** service.
3. Ping **isatap.adatum.com** to verify it can be resolved.
The name should resolve and you should receive four replies from 172.16.0.1.

► Task 4: Enable LON-DC1 as an ISATAP client

1. On LON-DC1, use the following **Set-NetIsatapConfiguration** cmdlet to enable ISATAP:

```
Set-NetIsatapConfiguration -State Enabled
```

2. Use ipconfig to verify that the Tunnel adapter for ISATAP has an IPv6 address on the 2001:db8:0:2/64 network.
Notice that this address includes the IPv4 address of LON-DC1.

► Task 5: Test connectivity

1. On LON-SVR2, use the following ping command to test connectivity to the ISATAP address for LON-DC1:

```
ping 2001:db8:0:2:0:5efe:172.16.0.10
```

2. Use the Server Manager to modify the properties of TCP/IPv6 on the Ethernet, and add **2001:db8:0:2:0:5efe:172.16.0.10** as the preferred DNS server.
3. Use the ping command to test connectivity to **LON-DC1**.



Note: A ping from LON-DC1 to LON-SVR2 does not respond because the firewall configuration on LON-SVR2 blocks ping requests.

Results: After completing this exercise, you will have configured an ISATAP router on LON-RTR to allow communication between an IPv6-only network and an IPv4-only network.

► Prepare for the next module

After 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 **20410C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20410C-LON-SVR1**, **20410C-LON-RTR** and **20410C-LON-SVR2**.

Lab Review Questions

Question: Did you configure IPv6 statically or dynamically in this lab?

Question: Why did you not need to configure LON-DC1 with the IPv4 address of the ISATAP router?

Module Review and Takeaways

Review Questions

Question: What is the main difference between 6to4 and Teredo?

Question: How can you provide a DNS server to an IPv6 host dynamically?

Question: Your organization is planning to implement IPv6 internally. After some research, you have identified unique local IPv6 addresses as the correct type of IPv6 addresses to use for private networking. To use unique local IPv6 addresses, you must select a 40-bit identifier that is part of the network. A colleague suggests using all zeros for the 40 bits. Why is this not a good idea?

Question: How many IPv6 addresses should an IPv6 node be configured with?

Best Practices

Use the following best practices when implementing IPv6:

- Do not disable IPv6 on Windows 8 or Windows Server 2012.
- Enable coexistence of IPv4 and IPv6 in your organization rather than using transition technologies.
- Use unique local IPv6 addresses on your internal network.
- Use Teredo to implement IPv6 connectivity over the IPv4 Internet.

Module 9

Implementing Local Storage

Contents:

Module Overview	9-1
Lesson 1: Overview of Storage	9-2
Lesson 2: Managing Disks and Volumes	9-12
Lesson 3: Implementing Storage Spaces	9-23
Lab: Implementing Local Storage	9-29
Module Review and Takeaways	9-34

Module Overview

Storage is one of the key components that you must consider when planning and deploying a Windows Server® 2012 operating system. Most organizations require a great deal of storage because users work regularly with apps that create new files that require storage in a central location. When users keep their files for longer periods of time, storage demands increase. Every time a user logs on to a server, an audit trail is created in an event log; this also uses storage. Even as files are created, copied, and moved, storage is required.

This module introduces you to different storage technologies. It discusses how to implement the storage solutions in Windows Server 2012, and how to use Storage Spaces, a new feature that you can use to combine disks into pools that are then managed automatically.

Objectives

After completing this module you will be able to:

- Describe various storage technologies.
- Explain how to manage disks and volumes.
- Explain how to implement Storage Spaces.

Lesson 1

Overview of Storage

When you plan a server deployment, one of the key components that you require is storage. There are various types of storage that you can utilize, such as locally-attached storage, storage that is remotely accessed via Ethernet, or storage connected with optical fiber. You should be aware of each solution's benefits and limitations.

As you prepare to deploy storage for your environment, you need to make some important decisions. This lesson addresses questions to consider, such as:

- Does the storage need to be fast?
- Does the storage need to be highly available?
- How much storage does your deployment actually require?
- How much resilience do you need to add to the initial storage requirement to ensure that your investment remains secure in the future?

Lesson Objectives

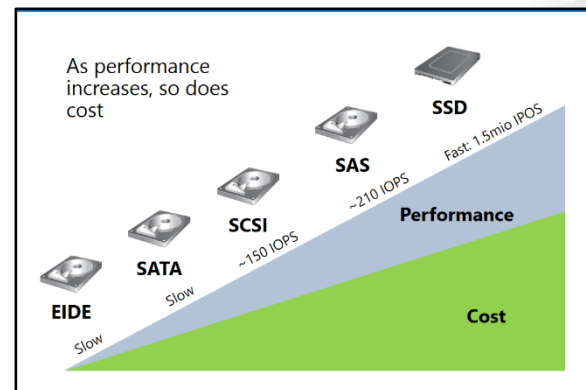
After completing this lesson, you will be able to:

- Describe disk types and performance.
- Describe direct-attached storage.
- Describe network-attached storage.
- Describe a storage area network (SAN).
- Describe Redundant Array of Independent Disks (RAID).
- Describe RAID levels.
- Describe the new file and storage service features in Windows Server 2012 and Windows Server 2012 R2.

Disk Types and Performance

There are various types of disks available that you can use to provide storage to server and client systems. The speed of disks is measured in Input Outputs per second (IOPS). The most common types of disks are:


- *Enhanced Integrated Drive Electronics (EIDE)*. EIDE is based on standards that were created in 1986. The Integrated Drive Electronics (IDE) interface supports both the Advanced Technology Attachment 2 (ATA-2) and Advanced Technology Attachment Packet Interface (ATAPI) standards. *Enhanced* refers to the ATA-2 (Fast ATA) standard. Due to the addressing standards of this technology, there is a 128 gigabyte (GB) limitation on storage using EIDE. Also, the speed of an EIDE drive is limited to a maximum of 133 megabytes (MB) per second. EIDE drives are almost never used on servers today.



- **Serial Advanced Technology Attachment (SATA).** Introduced in 2003, SATA is a computer bus interface, or channel, for connecting the motherboard or device adapters to mass storage devices such as hard disk drives and optical drives. SATA was designed to replace EIDE. It is able to use the same low-level commands as EIDE, but SATA host adapters and devices communicate via a high-speed serial cable over two pairs of conductors. It can operate at speeds of 1.5, 3.0, and 6.0 GB per second, depending on the SATA revision (1, 2 or 3 respectively). SATA disks are generally low-cost disks that provide mass storage. Because SATA drives are less expensive than other drive options, but also provide less performance, organizations might choose to deploy SATA drives when they require large amounts of storage but not high performance. SATA disks are also less reliable compared to serial attached SCSI (SAS) disks.

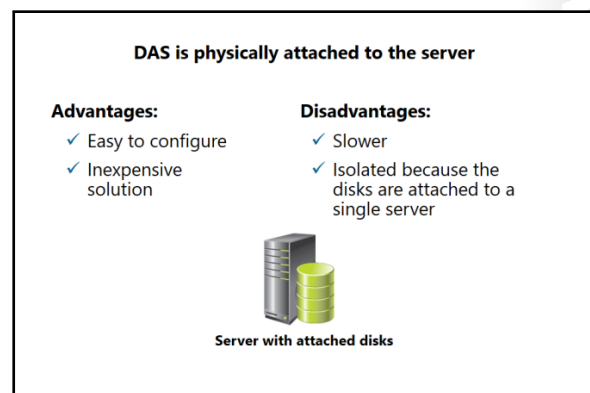
A variation on the SATA interface is eSATA, which is designed to enable high-speed access to externally-attached SATA drives.

- **Small computer system interface (SCSI).** SCSI is a set of standards for physically connecting and transferring data between computers and peripheral devices. SCSI was originally introduced in 1978 and became a standard in 1986. Similar to EIDE, SCSI was designed to run over parallel cables; however, recently the usage has been expanded to run over other mediums. The 1986 parallel specification of SCSI had initial speed transfers of 5 MB per second. The more recent 2003 implementation, Ultra 640 SCSI, also known as Ultra 5, can transfer data at speeds of 640 MB per second. SCSI disks provide higher performance than SATA disks, but are also more expensive.
- **SAS.** SAS is a further implementation of the SCSI standard. SAS depends on a point-to-point serial protocol that replaces the parallel SCSI bus technology, and uses the standard SCSI command set. SAS offers backward-compatibility with second generation SATA drives. SAS drives are reliable and made for 24 hours a day, seven days a week (24/7) operation in data centers. With up to 15,000 rotations per minute (RPM), these disks are also the fastest traditional hard disks.
- **Solid State Drives (SSDs).** SSDs are data storage devices that use solid-state memory to store data rather than using the spinning disks and movable read/write heads that are used in other disks. SSDs use microchips to store the data and do not contain any moving parts. SSDs provide fast disk access, use less power, and are less susceptible to failure from being dropped than traditional hard disks (such as SAS drives), but are also much more expensive per GB of storage. SSDs typically use a SATA interface, so you can usually replace hard disk drives with SSDs without any modifications.

 **Note:** Fibre Channel, FireWire, or USB-attached disks are also available storage options. They define either the transport bus or the disk type. For example, USB-attached disks use mostly with SATA or SSD drives to store data.

What Is Direct Attached Storage?

Almost all servers provide some built-in storage. This type of storage is referred to as *direct attached storage* (DAS). DAS can include disks that are physically located inside the server or connect directly with an external array, or disks that connect to the server with a USB cable or an alternative method. Because DAS storage is physically connected to the server, if the server suffers a power failure, the storage is unavailable. DAS comes in various disk types such as SATA, SAS or SSD, which affect the speed and the performance of the storage, and has both advantages and disadvantages.



Advantages of Using DAS

A typical DAS system is made up of a data storage device that includes a number of hard disk drives that connect directly to a computer through a host bus adapter (HBA). Between the DAS and the computer, there are no network devices such as hubs, switches, or routers. Instead, the storage is connected directly to the server that utilizes it, making DAS the easiest storage system to deploy and maintain.

DAS is also usually the least expensive storage available today, and is widely available in various speeds and sizes to accommodate various installations. In addition to being inexpensive, DAS is very easy to configure. In most instances, you would simply plug in the device, ensure that the running Windows® operating system recognizes it, and then use Disk Management to configure the disks.

Disadvantages of Using DAS

Storing data locally on DAS makes data centralization more difficult because the data is located on multiple servers. This can make it more complex to back up the data and, for users, more difficult to locate the data they want to find. Furthermore, if any one device that has DAS connected to it suffers a power outage, the storage on that computer becomes unavailable.

DAS also has drawbacks in its access methodologies. Due to the way reads and writes are handled by the server operating system, DAS can be slower than other storage technologies. Another drawback is that DAS shares the processing power and server memory of the server to which it is connected. This means that, on very busy servers, disk access might become slow when the operating system is overloaded.

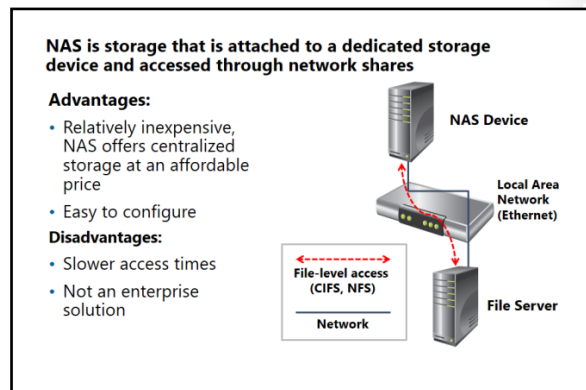
What Is Network Attached Storage?

Network attached storage (NAS) is storage that is connected to a dedicated storage device and then accessed over the network. NAS is different from DAS in that the storage is not directly attached to each individual server, but rather is accessible across the network to many servers. NAS has two distinct solutions: a low-end appliance (NAS only), and an enterprise-class NAS that integrates with SAN.

Each NAS device has a dedicated operating system that solely controls the access to the data on the device, which reduces the overhead associated with sharing the storage device with other server services. An example of NAS software is Windows Storage Server, a feature of Windows Server 2012.

NAS devices typically provide file-level access to the storage. This means that the data on the storage is accessible only as files, and you must use protocols like Common Internet Files System (CIFS), Server Message Block (SMB), or Network File System (NFS) to access the files.

To enable NAS storage, you need a storage device. Frequently, these devices do not have any server interfaces such as keyboards, mice, and monitors. To configure the device, you need to provide a network configuration and then access the device across the network. You can then create network shares on the device by using the name of the NAS and the share created. These shares are then accessible to users on the network.



Advantages of Using NAS

NAS is an ideal choice for organizations that are looking for a simple and cost-effective way to achieve fast data access for multiple clients at the file level. Users of NAS benefit from performance and productivity gains because the processing power of the NAS device is dedicated solely to the distribution of the files.

NAS also fits nicely into the market as a mid-priced solution. It is not expensive, but it suits more needs than DAS in the following ways:

- NAS storage is usually much larger than DAS.
- NAS offers a single location for all critical files, rather than dispersing them on various servers or devices with DAS.
- NAS offers centralized storage at an affordable price.
- NAS units are accessible from any operating system. They often have multi-protocol support and can serve up data via CIFS and NFS simultaneously. For example, Windows and Linux hosts can simultaneously access a NAS unit.

NAS can also be considered a Plug and Play solution that is easy to install, deploy, and manage, with or without IT staff onsite.

Disadvantages of Using NAS

NAS is slower than SAN technologies. NAS is frequently accessed via Ethernet protocols. Because of this, it relies heavily on the network supporting the NAS solution. For this reason, NAS is commonly used as a file sharing/storage solution and cannot (and should not) be used with data-intensive applications such as Microsoft® Exchange Server and Microsoft SQL Server®.

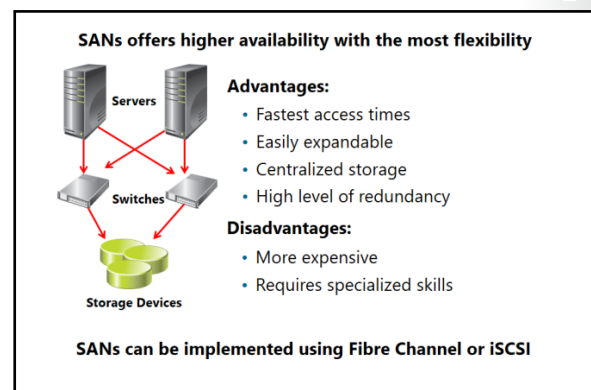
NAS is affordable for small to mid-size businesses, but provides less performance and may be less reliable than a SAN. For this reason, most large enterprises use SANs rather than NAS.

 **Additional Reading:** For more information about Windows Storage Server 2012 R2, see Windows Server 2012 R2 at <http://go.microsoft.com/fwlink/?LinkID=199647>.

What Is a SAN?

The third type of storage is a storage area network (SAN). A SAN is a high-speed network that connects computer systems or host servers to high-performance storage subsystems. A SAN usually includes various components such as host bus-adapters (HBAs), special switches to help route traffic, and storage disk arrays with logical unit numbers (LUNs) for storage.

A SAN enables multiple servers to access a pool of storage in which any server can potentially access any storage unit. Because a SAN uses a network, you can use a SAN to connect many different devices and hosts and provide access to any connected device from anywhere.



SANs provide block level access. This means that, rather than accessing the content on the disks as files by using a file access protocol, SANs write blocks of data directly to the disks using protocols such as Fibre Channel over Ethernet (FCOE) or Internet Small Computer System Interface (iSCSI).

Today, most SAN solutions offer SAN and NAS together. The backend head units, disks, and technologies are identical; the access method is the only thing that changes. Enterprises often provision block storage from the SAN to the servers using FCOE or iSCSI, whereas NAS services are made available via CIFS and NFS.

Advantages of Using SAN

SAN technologies read and write at block levels, making data access much faster. For example, with most DAS and NAS solutions, if you write a file of 8 GB, the entire file has to be read/written and its checksum calculated. With a SAN, the file is written to the disk based on the block size for which the SAN is set up. This speed is accomplished by using fiber channel and block level writing, instead of having to read/write an entire file by using a checksum.

SANs also provide:

- Centralization of storage into a single pool, which enables storage resources and server resources to grow independently. They also enable storage to be dynamically assigned from the pool when it is required. Storage on a given server can be increased or decreased as needed without complex reconfiguring or re-cabling of devices.
- Common infrastructure for attaching storage, which enables a single common management model for configuration and deployment.
- Storage devices that are inherently shared by multiple systems.
- Data transfer directly from device to device without server intervention.
- A high level of redundancy. Most SANs are deployed with multiple network devices and paths through the network. As well, the storage device contains redundant components such as power supplies and hard disks.

Disadvantages of Using SAN

The main drawback to SAN technology is that due to the complexities in the configuration, SAN often requires management tools and expert skills. It is also considerably more expensive than DAS or NAS—an entry level SAN can often cost as much as a fully-loaded server with a DAS or an NAS device, and that is without any SAN disks or configuration.

To manage a SAN, you often use command-line tools. You must have a firm understanding of the underlying technology, including the LUN setup, the Fibre Channel network, the block sizing, and other factors. In addition, each storage vendor often implements SANs using different tools and features. Because of this, organizations often have dedicated personnel whose only job is to manage the SAN deployment.



Note: You can implement SANs using a variety of technologies. The most common options are Fibre Channel and iSCSI.

What Is RAID?

RAID is a technology that you can use to configure storage systems to provide high reliability and (potentially) high performance. RAID implements storage systems by combining multiple disks into a single logical unit called a *RAID array*. Depending on the configuration, a RAID array can withstand the failure of one or more of the physical hard disks contained in the array, and/or provide higher performance than is available by using a single disk.

RAID provides an important component—redundancy—that you can use when planning and deploying Windows Server 2012 servers. In most organizations, it is important that the servers are available all of the time. Most servers provide highly redundant components such as redundant power supplies and redundant network adapters. The goal of this redundancy is to ensure that the server remains available even when a single component on the server fails. By implementing RAID, you can provide the same level of redundancy for the storage system.

RAID:

- Combines multiple disks into a single logical unit to provide fault tolerance and performance
- Provides fault tolerance by using:
 - Disk mirroring
 - Parity information
- Can provide performance benefits by spreading disk I/O across multiple disks
- Can be configured using several different levels
- Should not replace server backups

How RAID Works

RAID enables fault tolerance by using additional disks to ensure that the disk subsystem can continue to function even if one or more disks in the subsystem fail. RAID uses two options for enabling fault tolerance:

- **Disk mirroring.** With disk mirroring, all of the information that is written to one disk is also written to another disk. If one of the disks fails, the other disk is still available.
- **Parity information.** Parity information is used in the event of a disk failure to calculate the information that was stored on a disk. If you use this option, the server or RAID controller calculates the parity information for each block of data that is written to the disks, and then stores this information on another disk or across multiple disks. If one of the disks in the RAID array fails, the server can use the data that is still available on the functional disks along with the parity information to recreate the data that was stored on the failed disk.

RAID subsystems can also provide potentially better performance than single disks by distributing disk reads and writes across multiple disks. For example, when implementing disk striping, the server can read information from all hard disks in the stripe set. When combined with multiple disk controllers, this can provide significant improvements in disk performance.



Note: Although RAID can provide a greater level of tolerance for disk failure, you should not use RAID to replace traditional backups. If a server has a power surge or catastrophic failure and all of the disks fail, then you would still need to rely on standard backups.

Hardware RAID vs. Software RAID

Implement hardware RAID by installing a RAID controller in the server and then configure it by using the RAID controller configuration tool. With this implementation, the RAID configuration is hidden from the operating system, but the RAID arrays are exposed to the operating system as single disks. The only configuration you need to perform in the operating system is to create volumes on the disks.

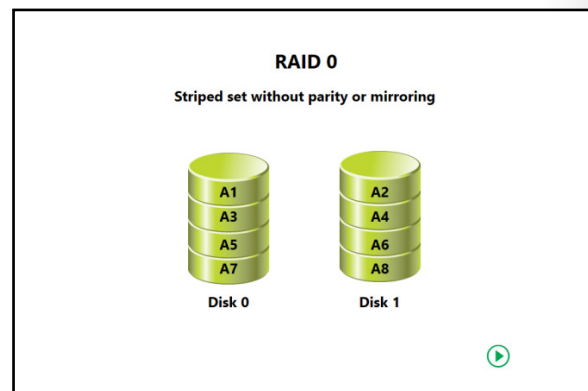
Software RAID is implemented by exposing all of the disks that are available on the server to the operating system, and then configuring RAID from within the operating system. Windows Server 2012 supports the use of software RAID, and you can use Disk Management to configure several different levels of RAID.

When choosing to implement hardware or software RAID, consider the following:

- Hardware RAID requires disk controllers that are RAID-capable. Most disk controllers shipped with new servers have this functionality.
- To configure hardware RAID, you need to access the disk controller management program. Normally, you can access this during the server boot process or by using a web page that runs management software.
- Implementing disk mirroring for the disk containing the system and boot volume with software RAID can require additional configuration when a disk fails. Because the RAID configuration is managed by the operating system, you must configure one of the disks in the mirror as the boot disk. If that disk fails, you may need to modify the boot configuration for the server to start the server. This is not an issue with hardware RAID, because the disk controller accesses the available disk and exposes it to the operating system.
- In older servers, you may get better performance with software RAID when using parity, because the server processor can calculate parity more quickly than the disk controller can. This is no longer an issue with newer servers, where you may get better performance on the server because you can offload the parity calculations to the disk controller.

RAID Levels

When implementing RAID, you need to decide what level of RAID to implement. The following table lists the features for each different RAID level.



Level	Description	Performance	Space utilization	Redundancy	Comments
RAID 0	Striped set without parity or mirroring Data is written sequentially to each disk	High read and write performance	All space on the disks is available	A single disk failure results in the loss of all data	Use only in situations where you require high performance and can tolerate data loss
RAID 1	Mirrored set without parity or striping Data is written to both disks simultaneously	Good performance	Can only use the amount of space that is available on the smallest disk	Can tolerate a single disk failure	Frequently used for system and boot volumes with hardware RAID
RAID 2	Data is written in bits to each disk with parity written to separate disk or disks	Extremely high performance	Uses one or more disks for parity	Can tolerate a single disk failure	Requires that all disks be synchronized Not currently used
RAID 3	Data is written in bytes to each disk with parity written to separate disk or disks	Very high performance	Uses one disk for parity	Can tolerate a single disk failure	Requires that all disks be synchronized Rarely used
RAID 4	Data is written in blocks to each disk with parity written to a dedicated disk	Good read performance, poor write performance	Uses one disk for parity	Can tolerate a single disk failure	Rarely used
RAID 5	Striped set with distributed parity Data is written in blocks to each disk with parity spread across all disks	Good read performance, poor write performance	Uses the equivalent of one disk for parity	Can tolerate a single disk failure	Commonly used for data storage where performance is not critical, but maximizing disk usage is important
RAID 6	Striped set with dual distributed parity Data is written in blocks to each disk with double parity written across all disks	Good read performance, poor write performance	Uses the equivalent of two disks for parity	Can tolerate two disk failures	Commonly used for data storage where performance is not critical but maximizing disk usage and availability are important

(continued from previous page)

Level	Description	Performance	Space utilization	Redundancy	Comments
RAID 0+1	Striped sets in a mirrored set A set of drives is striped, and then the strip set is mirrored	Very good read and write performance	Only half the disk space is available due to mirroring	Can tolerate the failure of two or more disks as long as all failed disks are in the same striped set	Not commonly used
RAID 1+0 (or 10)	Mirrored set in a stripe set Several drives are mirrored to a second set of drives, and then one drive from each mirror is striped	Very good read and write performance	Only half the disk space is available due to mirroring	Can tolerate the failure of two or more disks as long as both disks in a mirror do not fail	Frequently used in scenarios where performance and redundancy are critical, and the cost of the required additional disks is acceptable
RAID 5+0 (or 50)	Striped set with distributed parity in a stripe set Drives are striped with RAID 5, and then striped without parity	Good read performance, better write performance than RAID 5	The equivalent of at least two disks is used for parity	Provides better fault tolerance than a single RAID level	This level is recommended for applications that require high fault tolerance, capacity, and random positioning performance Requires at least six drives



Note: The most common RAID levels are RAID 1 (also known as mirroring), RAID 5 (also known as striped set with distributed parity), and RAID 1+0 (also known as mirrored set in a stripe set).

Question: Should all disks be configured with the same amount of fault tolerance?

Windows Server 2012 and Windows Server 2012 R2 Storage Features

Windows Server 2012 and Windows Server 2012 R2 include some important enhancements to the File and Storage Services server role. The new features include:

- **Storage Spaces.** Storage Spaces is a storage virtualization feature that you can use to add multiple physical disks of any type and size to a storage pool, and then create highly available virtual disks from the storage pool. With Storage Spaces, you can implement and manage a storage infrastructure that provides a high level of performance and redundancy without implementing any special storage infrastructure.
- **Data deduplication.** Data deduplication optimizes volume storage by finding redundant data on a volume, and then ensuring that the data is stored only once on the volume. It does this by storing the data in a single location, and then providing a reference to this single location in place of other redundant copies of the data. Data is segmented into 32 KB to 218 KB chunks, so data deduplication can optimize not only redundant files, but also portions of files that are redundant on the volume.
- **iSCSI Target Server.** Windows Server 2012 includes the iSCSI Target Server role to provide block storage to other servers and applications. iSCSI enables you to deploy a highly available SAN infrastructure using a standard network infrastructure. Windows Server 2012 R2 provides enhancements to the iSCSI Target Server role by supporting the creation of larger virtual disks that use the .vhdx format, optimizing disk caching, and increasing the number of sessions per server.
- **Management enhancements.** Windows Server 2012 provides a single management console for the File and Storage Services server role. You can use this console to manage all the file and storage components on a local or a remote server. Windows Server 2012 also provides new Windows PowerShell commands you can use to manage disks and storage.
- **Work Folders.** Work Folders enable users to access work files on computers and devices that are not members of an AD DS domain. You can synchronize the Work Folder contents from corporate file servers to the devices, so that users can work with the files easily. Administrators can maintain control over corporate data by setting permissions and device management policies to manage how users can use Work Folders.
- **Distributed File System (DFS) enhancements.** Windows Server 2012 R2 provides several new features for DFS, including the following:
 - A Windows PowerShell module for managing DFS
 - A database cloning feature for initial synchronization
 - A database corruption recovery feature
 - An option to disable cross-file remote differential compression (RDC)

Windows Server 2012 and Windows Server 2012 R2 provide several file and storage services enhancements including:

- Storage Spaces
- Data deduplication
- iSCSI Target Server
- Management enhancements
- Work Folders
- DFS enhancements

If you disable cross-file RDC, the network bandwidth used for replication increases; however, this decreases the processor load on file servers.



Note: Storage Spaces and storage pools are covered later in this lesson. The other storage enhancements are covered in Course 20411C: *Administering Windows Server 2012* and Course 20412C: *Configuring Advanced Windows Server 2012 Services*.

Lesson 2

Managing Disks and Volumes

Identifying which storage technology that you want to deploy is the first critical step in making sure that your environment is prepared for data storage requirements. This, however, is only the first step. There are other steps that you need to take to prepare for data storage requirements.

For example, once you have identified the best storage solution, or have chosen a combination of storage solutions, you need to figure out the best way to manage that storage. Ask yourself the following questions:

- What disks are you going to allocate to a storage pool?
- Are the type of file systems going to be the same for all disks?

This lesson addresses these and similar questions, including why it is important to manage disks and what tools you need to manage disks.

Lesson Objectives

After completing this lesson, you will be able to:

- Explain how to select a partition table format.
- Describe the difference between basic and dynamic disk types.
- Explain how to select a file system.
- Describe a resilient file system.
- Describe mount points and links.
- Explain how to create mount points and links.
- Describe the process of extending and shrinking volumes.

Selecting a Partition Table Format

A partition table format, or *partition style*, refers to the method that an operating system such as Windows Server 2012 uses to organize partitions or volumes on a disk. For Windows operating systems, you can decide between master boot record (MBR) and globally unique identifier (GUID) partition table (GPT).

MBR

The *MBR partition table format* is the standard partitioning scheme that has been used on hard disks since the first personal computers came out in the 1980s. The MBR partition table format has the following characteristics:

- A partition supports a maximum of four primary partitions per drive.
- A partition can have maximum of 2 terabytes (TB) (2.19×10^{12} bytes).

MBR

- Standard Partition table format since early 1980s
- Supports a maximum of 4 primary partitions per drive
- Can partition a disk up to 2 TB

GPT

- GPT is the successor of MBR partition table format
- Supports a maximum of 128 partitions per drive
- Can partition a disk up to 18 EB

✓ Use MBR for disks smaller than 2 TB

✓ Use GPT for disks larger than 2 TB

- If you initialize a disk larger than 2 TB using MBR, the disks are only able to store volumes up to 2 TB and the rest of the storage is not used. You must convert the disk to GPT if you want to use all of its space.



Note: You can use the MBR partition table format for disk drives that never surpass 2 TB in size. This provides you with a bit more space, because GPT requires more disk space than MBR.

GPT

The GPT was introduced with Windows Server 2003 and Windows XP 64-bit Edition to overcome the limitations of MBR, and to address larger disks. GPT has the following characteristics:

- GPT is the successor of MBR partition table format.
- GPT supports a maximum of 128 partitions per drive.
- A partition can have up to 18 exabytes (EB).
- A hard disk can have up to 8 zettabytes (ZB), with 512 kilobytes (KB) logical block addressing (LBA).
- To boot from a GPT partition table, your BIOS must support GPT.



Note: If your hard disk is larger than 2 TB, you must use the GPT partition table format.



Additional Reading: For frequently asked questions about the GUID partitioning table disk architecture, see <http://go.microsoft.com/fwlink/?LinkID=266748>.

Selecting a Disk Type

When selecting a type of disk for use in Windows Server 2012, you can choose between basic disks and dynamic disks.

Basic Disk

Basic storage uses normal partition tables that are used by all versions of the Windows operating system. A disk that is initialized for basic storage is called a *basic disk*. A basic disk contains basic partitions, such as primary partitions and extended partitions. You can subdivide extended partitions into logical drives.

Basic disks are:

- Disks initialized for basic storage
- The default storage for Windows operating system

Dynamic disks can:

- Be modified without restarting Windows
- Provide several options for configuring volumes

Disk volume requirements include:

- A system volume for hardware-specific files that are required to start the server
- A boot volume for the Windows operating system files

By default, when you initialize a disk in the Windows operating system, the disk is configured as a basic disk. You can easily convert basic disks to dynamic disks without any loss of data; however, when converting a dynamic disk to basic disk, all data on the disk is lost.

There is no performance gain by converting basic disks to dynamic disks, and some applications cannot address data that is stored on dynamic disks. For these reasons, most administrators do not convert basic disks to dynamic disks unless they need to use some of the additional volume configuration options that are available with dynamic disks.

Dynamic Disk

Dynamic storage was introduced in the Microsoft Windows 2000 Server operating system. With dynamic storage, you can perform disk and volume management without the need to restart computers running Windows operating systems. A disk that is initialized for dynamic storage is called a *dynamic disk*. A dynamic disk contains dynamic volumes.

When you configure dynamic disks, you create volumes rather than partitions. A *volume* is a storage unit that is made from free space on one or more disks. You can format the volume with a file system, and can assign it a drive letter or configure it with a mount point.

The following is a list of the dynamic volumes that are available:

- *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 extended to 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 from multiple disks that is linked together. 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 will lose the entire spanned volume.
- *Striped volumes.* A striped volume has data that 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. This means that the loss of one disk causes the immediate loss of all the data. Striping is also known as *RAID-0*.
- *Mirrored volumes.* A mirrored volume is a fault-tolerant volume that has all data duplicated onto two physical disks. All of 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*.
- *RAID-5 volumes.* A RAID-5 volume is a fault-tolerant volume that has data striped across a minimum of three or more disks. Parity 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 you use, you must configure both 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 operating system (for example, Bootmgr and BOOTSECT.bak). The system volume can—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—but does not have to—be the same as the system volume.



Note: When you install the Windows 8 operating system or the Windows Server 2012 operating system in a clean installation, a separate system volume is created to enable encrypting the boot volume by using Windows BitLocker® drive encryption.



Additional Reading:

- For more information about how basic disks and volumes work, see <http://go.microsoft.com/fwlink/?LinkID=199648>.
- For more information about dynamic disks and volumes, see <http://go.microsoft.com/fwlink/?LinkID=199649>.

Selecting a File System

When you configure your disks in Windows Server 2012, you can choose between file allocation table (FAT), the NTFS file system, and Resilient File System (ReFS) file systems.

FAT

The FAT file system is the most simplistic of the file systems that Windows operating systems support. The FAT file system is characterized by a table that resides at the very top of the volume. To protect the volume, two copies of the FAT file system are maintained in case one becomes damaged. In addition, the file allocation tables and the root directory must be stored in a fixed location so that the system's boot files can be correctly located.

When selecting a file system, consider the differences between FAT, NTFS, and ReFS

FAT provides:

- Basic file system
- Partition size limitations
- FAT32 to enable larger disks
- exFAT developed for flash drives

NTFS provides:

- Metadata
- Auditing and journaling
- Security (ACLs and encryption)

ReFS provides:

- Backward compatibility support for NTFS
- Enhanced data verification and error correction
- Support for larger files, directories, volumes, and so on

A disk formatted with the FAT file system is allocated in clusters, whose sizes are determined by the size of the volume. When a file is created, an entry is created in the directory, and the first cluster number containing data is established. This entry in the table indicates either that this is the last cluster of the file, or points to the next cluster. There is no organization to the FAT directory structure, and files are given the first open location on the drive.

Because of the size limitation with the file allocation table, the original release of FAT could only access partitions that were less than 2 GB in size. To enable larger disks, Microsoft developed FAT32. FAT32 supports partitions of up to 2 TB.

FAT does not provide any security for files on the partition. You should never use FAT or FAT32 as the file system for disks attached to Windows Server 2012 servers. You might consider using FAT or FAT32 to format external media such as USB flash media.

The file system designed especially for flash drives is Extended FAT (exFAT). It can be used where FAT32 is not suitable, such as when you need a disk format that works with a television, which requires a disk that is larger than 2 TB. A number of media devices support exFAT, such as modern flat panel TVs, media centers, and portable media players.

NTFS

NTFS is the standard file system for all Windows operating systems beginning with Windows NT® Server 3.1. Unlike FAT, there are no special objects on the disk, and there is no dependence on the underlying hardware, such as 512-byte sectors. In addition, in NTFS there are no special locations on the disk, such as the tables.

NTFS is an improvement over FAT in several ways, such as better support for metadata, and the use of advanced data structures to improve performance, reliability, and disk space utilization. NTFS also has additional extensions such as security access control lists (ACLs), which you can use for auditing, file system journaling, and encryption.

NTFS is required for a number of Windows Server 2012 roles and features such as Active Directory® Domain Services (AD DS), Volume Shadow Services (VSS), Distributed File System (DFS) and File Replication Services (FRS). NTFS also provides a much higher level of security than FAT or FAT 32.

Resilient File System (ReFS)

ReFS was introduced with Windows Server 2012 to enhance the capabilities of NTFS. ReFS was developed to improve upon NTFS by offering larger maximum sizes for individual files, directories, disk volumes, and other items. Additionally, ReFS offers greater resiliency, meaning better data verification, error correction, and scalability.

You should use ReFS with very large volumes and very large file shares to overcome the NTFS limitation of error checking and correction. Because ReFS was not available prior to Windows Server 2012 (the only choice was NTFS), it makes sense to use ReFS with Windows Server 2012 instead of NTFS to achieve better error checking, better reliability, and less corruption. ReFS cannot be used for the boot volume.



Additional Reading: For more information on how FAT works, see <http://go.microsoft.com/fwlink/?LinkID=199652>.



Additional Reading: For more information on how NTFS works, see <http://go.microsoft.com/fwlink/?LinkID=199654>.

Question: What file system do you currently use on your file server? Will you continue to use it?

What Is ReFS?

ReFS is a new feature in Windows Server 2012. ReFS is based on the NTFS file system, and provides the following advantages:

- Metadata integrity with checksums
- Expanded protection against data corruption
- Maximizes reliability, especially during a loss of power (while NTFS has been known to experience corruption in similar circumstances)
- Large volume, file, and directory sizes
- Storage pooling and virtualization, which makes creating and managing file systems easier
- Redundancy for fault tolerance
- Disk scrubbing for protection against latent disk errors
- Resiliency to corruptions with recovery for maximum volume availability
- Shared storage pools across machines for additional failure tolerance and load balancing

ReFS inherits some features from NTFS, including the following:

- BitLocker drive encryption
- ACLs for security
- Update sequence number (USN) journal
- Change notifications

ReFS is a new file system that is built in to Windows Server 2012. Advantages include:

- Metadata integrity with checksums
- Integrity streams with user data integrity
- Allocation on write transactional model
- Large volume, file, and directory sizes (2⁶⁴ 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 recovery
- Shared storage pools across machines

- Symbolic links, junction points, mount points and reparse points
- Volume snapshots
- File IDs

Because ReFS uses a subset of features from NTFS, it is designed to maintain backward compatibility with NTFS. Therefore, applications that run on Windows Server 2012 can access files on ReFS just as they would on NTFS. However, an ReFS-formatted drive is not recognized when placed in computers that are running Windows Server operating systems previous to Windows Server 2012. ReFS drives can be used with Windows 8.1, but not with Windows 8.

With NTFS, you can change the size of a cluster; however, with ReFS, each cluster has a fixed size of 64 KB, which you cannot change. Encrypted File System (EFS) for files is not supported in ReFS.

As its name implies, the new file system offers greater 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	Approximately 16 EB (18.446.744.073.709.551.616 bytes)
Maximum size of a single volume	2 ⁷⁸ bytes with 16 KB 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	32,000 Unicode characters
Maximum path length	32,000
Maximum size of any storage pool	4 petabytes (PB)
Maximum number of storage pools in a system	No limit
Maximum number of spaces in a storage pool	No limit



Additional Reading: For more information about ReFS, see <http://go.microsoft.com/fwlink/?linkID=270872>.

What Are Mount Points and Links?

With NTFS and ReFS file systems, you can create mount points and links to refer to files, directories, and volumes.

Mount Points

Mount points are used in Windows operating systems to make a portion of a disk or the entire disk useable by the operating system. Most commonly, mount points are associated with drive-letter mappings so that the operating system can gain access to the disk through the drive letter.

A mount point is a reference to a location on a disk that enables Windows operating system access to disk resources

Use volume mount points:

- To mount volumes or disks as folders instead of using drive letters
- When you do not have drive letters available for creating new volumes
- To add disk space without changing the folder structure

A link file contains a reference to another file or directory

Link options:

- Symbolic file link (or, soft link)
- Symbolic directory link (or, directory junctions)

Since the Windows 2000 Server was first introduced, you have been able to enable volume mount points, which you can then use to mount a hard disk to an empty folder that is located on another drive. For example, if you add a new hard disk to a server, rather than mounting the drive using a drive letter, you can assign a folder name such as C:\datadrive to the drive. When you do this, any time you access the C:\datadrive folder, you are actually accessing the new hard disk.

Volume mount points can be useful in the following scenarios:

- If you are running out of drive space on a server and you want to add disk space without modifying the folder structure. You can add the hard disk, and configure a folder to point to the hard disk.
- If you are running out of available letters to assign to partitions or volumes. If you have several hard disks that are attached to the server, you may run out of available letters in the alphabet to which you can assign drive letters. By using a volume mount point, you can add additional partitions or volumes without using more drive letters.
- If you need to separate disk input/output (I/O) within a folder structure. For example, if you are using an application that requires a specific file structure, but which uses the hard disks extensively, you can separate the disk I/O by creating a volume mount point within the folder structure.



Note: You can assign volume mount points only to empty folders on an NTFS partition. This means that if you want to use an existing folder name, you must first rename the folder, create and mount the hard disk using the required folder name, and then copy the data to the mounted folder.

Links

A *link* is a special type of file that contains a reference to another file or directory in the form of an absolute or relative path. Windows supports the following two types of links:

- A symbolic file link (also known as a *soft link*)
- A symbolic directory link (also known as a *directory junction*)

A link that is stored on a server share could refer back to a directory on a client that is not actually accessible from the server where the link is stored. Because the link processing is done from the client, the link would work correctly to access the client, even though the server cannot access the client.

Links operate transparently. Applications that read or write to files that are named by a link behave as if they are operating directly on the target file. For example, you can use a symbolic link to link to a Hyper-V® parent virtual hard disk file (.vhd) from another location. Hyper-V uses the link to work with

the parent virtual hard disk (VHD) as it would the original file. The benefit of using symbolic links is that you do not need to modify the properties of your differencing VHD.



Note: In Hyper-V, you can use a differencing VHD to save space by making changes only to the child VHD, when the child VHD is part of a parent/child VHD relationship.

Links are sometimes easier to manage than mount points. Mount points force you to place the files on the root of the volumes, whereas with links, you can be more flexible with where you save files.

You can create links by using the mklink.exe command-line tool.

Demonstration: Creating Mount Points and Links

In this demonstration, you will see how to:

- Create a mount point.
- Create a directory junction for a folder.
- Create a hard link for a file.

Demonstration Steps

Create a mount point

1. Sign in to **LON-SVR1** with the username **Adatum\Administrator** and the password **Pa\$\$w0rd**.
2. Open **Computer Management**, and then expand **Disk Management**.
3. In Disk Management, initialize **Disk2** with **GPT (GUID Partition Table)**.
4. On Disk 2, create a Simple Volume with the following parameters:
 - Size: **4000 MB**
 - **Do not assign a drive letter or drive path**
 - File system: **NTFS**
 - Volume label: **MountPoint**
5. Wait until the volume is created, right-click **MountPoint**, and then click **Change Drive Letter and Paths**.
6. Change the drive letter as follows:
 - **Mount in the following empty NTFS folder**
 - Create new Folder **C:\MountPointFolde**r and use it as mount point.
7. On the taskbar, open a File Explorer window, and then click **Local Disk (C:)**. You should now see the MountPoint folder with a size of 4,095,996 KB assigned to it. Notice the icon that is assigned to the mount point.

Create a directory junction for a folder

1. Open a Command Prompt window.
2. Create a folder in C:\ called **CustomApp**, and run the following:
copy C:\windows\system32\notepad.exe C:\CustomApp.

3. At the command prompt, type **mklink /j AppLink CustomApp**, and then press Enter.
4. In a File Explorer window, browse to **C:\AppLink**. Notice that because it is a link, the directory path in the address bar is not updated to C:\CustomApp.

Create a hard link for a file

1. At a command prompt, type **mklink /h C:\AppLink\notepad2.exe C:\AppLink\notepad.exe**.
2. In File Explorer, notice that Notepad2.exe appears exactly the same as Notepad.exe. Both file names point to the same file.

Extending and Shrinking Volumes

In versions of Windows prior to Windows Server 2008 or Windows Vista®, you required additional software to shrink or extend a volume on your disk. Since Windows Server 2008 and Windows Vista, this functionality is included in the Windows operating system so you can use the Disk Management snap-in to resize NTFS volumes.

When you want to resize a volume, you must be aware of the following:

- You only have the ability to shrink or extend NTFS volumes. FAT, FAT32 or exFAT volumes cannot be resized.
- You can only extend ReFS volumes, not shrink them.
- You can extend a volume using free space on the same disk and on other disks. When you extend a volume with other disks, you create a dynamic disk with a spanned volume. In a spanned volume, if one disk fails, all data on the volume is lost. In addition, a spanned volume cannot contain boot or system partitions, thus you cannot extend your boot partitions by using another disk.
- When you want to shrink a partition, immovable files such as page files are not relocated. This means that you cannot reclaim space beyond the location where these files are on the volume. If you have the requirement to shrink a partition further, you need to delete or move the immovable files. For example, you can remove the page file, shrink the volume, and then add the page file back again.

You can resize NTFS volumes from the Windows operating system, beginning with Windows Vista and Windows Server 2008

When you want to resize a disk, consider the following:

- You can extend or shrink NTFS volumes
- ReFS volumes can only be extended
- FAT/FAT32/exFAT cannot be resized
- You can shrink a volume only up to immovable files
- Bad clusters on a disk prevent you from shrinking a volume



Note: As a best practice for shrinking volumes, you should defragment the files on the volume before you shrink it. This procedure returns the maximum amount of free disk space. During the defragmenting process, you can identify any immovable files.

- If bad clusters exist on the partition, you cannot shrink it.

To modify a volume, you can use Disk Management, the Diskpart.exe tool, or the **Resize-Partition** cmdlet in Windows PowerShell®.



Additional Reading: For more information about how to extend a basic volume, see <http://go.microsoft.com/fwlink/?LinkID=266749>.



Additional Reading: For more information about how to shrink a basic volume, see <http://go.microsoft.com/fwlink/?LinkID=266750>.

Managing Virtual Hard Disks

Starting with Windows 7 and Windows 2008 R2, you can manage virtual hard disks within the operating system in much the same way that you can manage physical disks. For example, you can create and attach a virtual hard disk and use it for storing data. The virtual hard disk appears as another drive letter in the disk or folder management tools.

Virtual hard disks are files that represent a traditional hard disk drive. Typically, virtual hard disks are used with Hyper-V as the operating system disk and the storage disks for virtual machines. In Windows 7 and Windows Server 2008 R2, you can access the same virtual hard disks from within the operating system. The virtual hard disks have the following characteristics:

- In Windows 7 and Windows Server 2008 R2, you can only work with .vhd files.
- In Windows 8 or Windows Server 2012 or later, you can also create and manage .vhdx files, which enable much larger disk sizes as well as provide other benefits.



Note: For details on the differences between .vhd and .vhdx files, see Module 13, which covers the use of virtual hard disks in Hyper-V.

Virtual hard disks are files that can be managed just like physical hard disks

You can:

- Create and manage virtual hard disks using Disk Management and Diskpart
- Configure .vhd or .vhdx files
- Configure computers to start from the virtual hard disk
- Transfer virtual hard disks from Hyper-V servers and start computers from the virtual hard disk
- Use virtual hard disks as a deployment technology

- You can create and attach virtual hard disks using disk management tools such as Disk Management and Diskpart. After creating and attaching the virtual hard disk, you can create volumes on the drive and format the partition.
- You can configure Windows 7 or Windows Server 2008 R2 or later versions to start from a virtual hard disks using the native VHD boot feature. This feature enables you to configure multiple operating systems on a single computer and choose which operating system to use when you start the computer.
- You can attach virtual hard disks that have been created using Hyper-V or that have been created on another computer. For example, if you create a virtual hard disk in Hyper-V, you can copy that virtual hard disk to another computer, and then use the native VHD boot feature to start the computer using the virtual disk that you created in Hyper-V.
- You can use virtual hard disks as a deployment technology. For example, you can use Hyper-V to create a standard image for desktop or server computers, and then distribute the image to other computers.

Demonstration: Managing Virtual Hard Disks

In this demonstration, you will see how to:

- Create a virtual hard disk.
- Manage a virtual hard disk.

Demonstration Steps

Create a virtual hard disk

1. In Server Manager, open **Disk Management**.
2. Create a new .vhdx file named **DiskF.vhdx** in the **Documents** folder. Assign a size of **10 MB**, and configure the file as **dynamically expanding**.
3. Verify that the .vhdx file was created in the documents folder.

Manage a virtual hard disk

1. In Disk Management, initialize the disk and create and format a new volume using all the space on the disk and a volume label **Data**.
2. Verify that the new disk appears in File Manager.

Lesson 3

Implementing Storage Spaces

Managing physical disks that are attached directly to a server has proven to be a tedious task for administrators. To overcome this problem, many organizations use SANs that essentially group physical disks together.

SANs require specialized configuration and sometimes specialized hardware, which makes them expensive. To overcome these issues, you can use Storage Spaces, which is a Windows Server 2012 feature that pools disks together and presents them to the operating system as a single disk. This lesson explains how to configure and implement the Storage Spaces feature.

Lesson Objectives


After completing this lesson, you will be able to:

- Describe the Storage Spaces feature.
- Describe various options for configuring virtual disks.
- Describe advanced management options for Storage Spaces.
- Describe how to configure Storage Spaces.
- Compare Storage Spaces with other storage methods.

What Is the Storage Spaces Feature?

Storage Spaces is a storage virtualization capability that is built into Windows Server 2012 and Windows 8 and later. It is a feature that is available for both NTFS and ReFS volumes, providing redundancy and pooled storage for numerous internal and external drives of differing sizes and interfaces. You can use Storage Spaces to add physical disks of any type and size to a storage pool, and then create highly available virtual disks from the storage pool. The primary advantage of Storage Spaces is that you do not manage single disks, but can manage multiple disks as one unit.

- Use storage spaces to add physical disks of any type and size to a storage pool, and then create highly-available virtual disks from the storage pool
- To create a virtual disk, you need the following:
 - One or more physical disks
 - Storage pool that includes the disks
 - Virtual disk that are created with disks from the storage pool
 - Disk drives that are based on virtual drives
- Virtual disks are not virtual hard disks; they should be considered a drive in Disk Manager
- Windows Server 2012 R2 enables Storage Space tiering and write-back caching



To create a highly-available virtual disk, you need the following:

- Physical disk. Physical disks are disks such as SATA or SAS disks. If you want to add physical disks to a storage pool, the disks need to satisfy the following requirements:
 - One physical disk is required to create a storage pool; a minimum of two physical disks is required to create a resilient mirror virtual disk.
 - A minimum of three physical disks are required to create a virtual disk with resiliency through parity.
 - Three-way mirroring requires at least five physical disks.
 - Disks must be blank and unformatted; no volume must exist on them.
 - Disks can be attached using a variety of bus interfaces including SAS, SATA, SCSI, and USB. If you want to use failover clustering with storage pools, you cannot use SATA, USB or SCSI disks.

- **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.
- **Virtual disk (or storage space).** This is similar to a physical disk from the perspective of users and applications. However, virtual disks are more flexible because they include thin provisioning or just-in-time (JIT) allocations, and they include resiliency to physical disk failures with built-in functionality such as mirroring.
- **Disk drive.** This is a volume that you can access from your Windows operating system, for example, by using a drive letter.

New Features of Windows Server 2012 R2 Storage Spaces

Storage Spaces were first introduced in Windows 2012. Windows Server 2012 R2 provides the following enhancements to Storage Spaces:

- **Tiered Storage Spaces.** Tiered Storage Spaces allow you to use a combination of disks in a Storage Space: very fast, but small-capacity hard disks (such as SSDs) alongside slower, but large-capacity hard disks. When you use this combination of disks, Storage Spaces automatically moves frequently-accessed data to the faster hard disks and moves less frequently-accessed data to the slower disks. By default, Storage Spaces moves data once day at 01:00 A.M. You can also configure where files will be stored. The advantage to this is if you have files that are frequently accessed, you can pin them to the faster disk. The goal of tiering is to balance capacity against performance. Windows Server 2012 R2 only supports two levels of disk tiers.
- **Write-back caching.** The purpose of write-back caching is to optimize the process of writing data to the disks in a Storage Space. Write-back caching typically works with tiered Storage Spaces. If the server running the Storage Space detects a peak in disk-writing activity, it automatically starts writing data to the faster disks. Write-back caching is enabled by default. Write-back caching is limited to 1 GB by default.

Virtual Disk Configuration Options

You can create virtual disks from storage pools. If your storage pool contains more than one disk, you can also create redundant virtual disks. To configure virtual disks or Storage Spaces in Server Manager or Windows PowerShell, you need to consider the following features and their redundancy functionalities.

Storage Layout

Configure this feature to define the number of disks from the storage pool that are allocated. Valid options include:

Feature	Options
Storage Layout	<ul style="list-style-type: none"> • Simple • Two-way or three-way mirror • Parity
Disk sector size	<ul style="list-style-type: none"> • 512 or 512e
Drive allocation	<ul style="list-style-type: none"> • Automatic • Manual • Hot Spare
Provisioning schemes	<ul style="list-style-type: none"> • Thin vs. fixed provisioning

- **Simple.** 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 concurrently. Do not host important data on a simple volume, because it provides no failover capabilities when the disk that is storing the data fails.
- **Two-way and three-way mirrors.** Mirror spaces maintain two or three copies of the data that they host (two data copies for two-way mirrors and three data copies for three-way mirrors). Data

duplication happens with every write to ensure that all data copies are always current. Mirror spaces also stripe the data across multiple physical drives. Mirror spaces provide the benefit of greater data throughput and lower access latency. They also do not introduce a risk of corrupting at-rest data, and do not require the extra journaling stage when writing data.

- **Parity.** A parity space is similar to RAID 5. 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 I/O optimization. Storage spaces require a minimum of three physical drives for parity spaces. Parity spaces have increased resiliency through journaling.

Disk Sector Size

A storage pool's sector size is set when it is created. If the list of drives being used contains only 512 and/or 512e drives, then the pool is defaulted to 512e. A 512 disk uses 512 byte sectors. A 512e drive is a hard disk with 4,096 byte sectors that emulates 512 byte sectors. If the list contains at least one 4 KB drive, then the pool sector size is defaulted to 4 KB. Optionally, an administrator can explicitly define the sector size that is inherited by all contained spaces in the pool. After an administrator defines this, the Windows operating system only permits you to add 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.

Drive Allocation

This defines how the drive is allocated to the pool. Options are:

- **Automatic.** 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 JIT allocation.
- **Manual.** Administrators can choose to specify Manual as the usage type for drives that are added to a pool. A manual drive is not used automatically as part of a storage space unless it is specifically selected at the creation of that storage space. This usage property makes it possible for administrators to specify particular types of drives for use by only certain Storage Spaces.
- **Hot Spare.** 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 upon to replace the failed drive.

Provisioning Schemes

You can provision a virtual disk by using two different schemes:

- **Thin provisioning space.** Thin provisioning is a mechanism that allows storage to be easily allocated on a just-enough and JIT basis. Storage capacity in the pool is organized into provisioning slabs that are not allocated until the point in time when datasets grow to require the storage. As opposed to the traditional fixed storage allocation method—where large pools of storage capacity are allocated but may remain unused—thin provisioning optimizes utilization of available storage. Organizations are also able to save on operating costs such as electricity and floor space that are associated with keeping unused drives operating. The downside of using thin provisioning is lower disk performance.
- **Fixed provisioning space.** With Storage Spaces, fixed provisioned spaces also employ the flexible provisioning slabs. The difference between thin provisioning and a fixed provisioning space is that the storage capacity in the fixed provisioning space is allocated at the same time that the space is created.

Cluster Disk Requirement

Failover clustering prevents interruption to workloads or data in the event of a machine failure. For a pool to support failover, clustering all assigned drives must support a multi-initiator protocol, such as SAS.



Note: You can use Storage Spaces to create both thin and fixed provisioning virtual disks within the same storage pool. Having both provisioned types in the same storage pool is convenient, particularly 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.

Question: What is the name for a virtual disk that is larger than the amount of disk space available on the physical disks portion of the storage pool?

Advanced Management Options for Storage Spaces

Server Manager provides you with basic management of virtual disks and storage pools. In Server Manager, you can create storage pools, add and remove physical disks from pools, and create, manage, and delete virtual disks. For example, in Server Manager you can view the physical disks that are attached to a virtual disk. If any of these disks are unhealthy, you will see an unhealthy disk icon next to the disk name.

To correct a failed disk in a virtual disk or storage pool, you must remove the disk that is causing the problem. Tools such as defragmenting, scan disk, or **chkdsk** cannot repair a storage pool. To replace a failed disk, you add a new disk to the pool. The new disk resynchronizes automatically when disk maintenance occurs during daily maintenance. Alternatively, you can trigger disk maintenance manually.

Windows PowerShell provides advanced management options for virtual disks and storage pools. Some examples of management cmdlets are listed in the following table.

- Basic Management for Storage Spaces is available in Server Manager
- For disk failure:
 - Do not use **chkdsk** or scan disk
 - Remove the drive and add a new one
- Advanced management requires Windows PowerShell



Windows PowerShell cmdlet	Description
Get-StoragePool	Lists storage pools
Get-VirtualDisk	Lists virtual disks
Repair-VirtualDisk	Repairs a virtual disk
Get-PhysicalDisk Where{\$_.HealthStatus -ne "Healthy"}	Lists unhealthy physical disks
Reset-PhysicalDisk	Removes a physical disk from a storage pool
Get-VirtualDisk Get-PhysicalDisk	Lists physical disks that are used for a virtual disk



Additional Reading: To learn more about storage cmdlets in Windows PowerShell, see <http://go.microsoft.com/fwlink/?LinkID=266751>.

Demonstration: Configuring Storage Spaces

In this demonstration, you will see how to:

- Create a storage pool.
- Create a virtual disk and a volume.

Demonstration Steps

Create a storage pool

1. Sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On LON-SVR1, in Server Manager, access **File and Storage Services** and **Storage Pools**.
3. In the STORAGE POOLS pane, create a **New Storage Pool** named **StoragePool1**, and add all of the available disks.

Create a virtual disk and a volume

1. In the VIRTUAL DISKS pane, create a **New Virtual Disk** with the following 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 task completes, and then ensure that the **Create a volume when this wizard closes** check box is selected.
3. In the New Volume Wizard, create a volume with these settings:
 - Virtual disk: **Simple vDisk**
 - File system: **ReFS**
 - Volume label: **Simple Volume**
4. Wait until the task completes, and then click **Close**.

Discussion: Comparing Storage Spaces with Other Storage Solutions

Storage Spaces in Windows Server 2012 provides an alternative to using more traditional storage solutions such as SANs and NAS.

Consider the questions on the slide and in the table below. Write your responses to the questions in the following table to prepare for the class discussion.

- 1. Does your organization currently use SANs or NAS?
- 2. What are the advantages of using Storage Spaces compared to using SANs or NAS?
- 3. What are the disadvantages of using Storage Spaces compared to using SANs or NAS?
- 4. In what scenarios would you recommend each option?



10 minutes

Question	Your Response
Does your organization currently use SANs or NAS?	
What are the advantages of using Storage Spaces compared to using SANs or NAS?	
What are the disadvantages of using Storage Spaces compared to using SANs or NAS?	
In what scenarios would you recommend each option?	

Lab: Implementing Local Storage

Scenario

Your manager has asked to add disk space to a file server. After creating volumes, your manager has also asked you to resize those volumes based on updated information he has been given. Finally, you need to make data storage redundant by creating a three-way mirrored virtual disk.

Objectives

After completing this lab, you will be able to:

- Install and configure a new disk.
- Resize volumes.
- Configure a redundant storage space.

Lab Setup

Estimated Time: 30 minutes

Virtual machines	20410C-LON-DC1 20410C-LON-SVR1
User name	Adatum\Administrator
Password	Pa\$\$w0rd

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

On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.

1. In Hyper-V® Manager, click **20410C-LON-DC1**, and in the Actions pane, click **Start**.
2. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
3. Sign in using the following credentials:
 - User name: **Administrator**
 - Password: **Pa\$\$w0rd**
 - Domain: **Adatum**
4. Repeat steps 1 through 3 for 20410C-LON-SVR1.

Exercise 1: Installing and Configuring a New Disk

Scenario

The file server in your branch office is low on disk space. You need to add a new disk to the server and create volumes based on specifications provided by your manager.

The main tasks for this exercise are as follows:

1. Initialize a new disk.
2. Create and format two simple volumes on the disk.
3. Verify the drive letter in a File Explorer window.

► Task 1: Initialize a new disk

1. Sign in to **LON-SVR1** with the username **Adatum\Administrator** and the password **Pa\$\$w0rd**.
2. In Server Manager, open **Computer Management**, and then access **Disk Management**.
3. Initialize **Disk2** and configure it to use **GPT (GUID Partition Table)**.

► Task 2: Create and format two simple volumes on the disk

1. In the Computer Management console, on Disk 2, create a **Simple Volume** with the following attributes:
 - Volume size: **4000 MB**
 - Drive Letter: **F**
 - File system: **NTFS**
 - Volume label: **Volume1**
2. In the Computer Management console, on Disk 2, create a **Simple Volume** with the following attributes:
 - Volume size: **5000 MB**
 - Drive Letter: **G**
 - File system: **ReFS**
 - Volume label: **Volume2**

► Task 3: Verify the drive letter in a File Explorer window

1. Use File Explorer to make sure you can access the following volumes:
 - **Volume1 (F:)**
 - **Volume2 (G:)**
2. On Volume2 (G:), create a folder named **Folder1**.

Results: After you complete this lab, you should have initialized a new disk, created two simple volumes, and then formatted them. You should also have verified that the drive letters you assigned are available in File Explorer.

Exercise 2: Resizing Volumes

Scenario

After installing the new disk in your file server, you are contacted by your manager who indicates that the information he gave you was incorrect. He now needs you to resize the volumes without losing any data.

The main tasks for this exercise are as follows:

1. Shrink Volume1.
2. Extend Volume2.

► **Task 1: Shrink Volume1**

- Use Disk Management to shrink **Volume1 (F:)** to **3000 MB**.

► **Task 2: Extend Volume2**

1. Use Disk Management to extend **Volume2 (G:)** by **1000 MB**.
2. Use File Explorer to verify that the folder **Folder1** is still on drive **G**.

Results: After this lab, you should have made one volume smaller, and extended another.

Exercise 3: Configuring a Redundant Storage Space

Scenario

Your server does not have a hardware-based RAID card, but you have been asked to configure redundant storage. To support this feature, you need to create a storage pool.

After creating the storage pool, you also need to create a redundant virtual disk. Because the data is critical, the request for redundant storage specifies that you must use a three-way mirrored volume. Shortly after the volume is in use, a disk fails and you have to add another disk to the storage pool to replace it.

The main tasks for this exercise are as follows:

1. Create a storage pool from five disks that are attached to the server.
2. Create a three-way mirrored virtual disk.
3. Copy a file to the volume, and verify that it is visible in File Explorer.
4. Remove a physical drive.
5. Verify that the write.exe file is still accessible.
6. Add a new disk to the storage pool and remove a broken disk.

► **Task 1: Create a storage pool from five disks that are attached to the server**

1. On LON-SVR1, open **Server Manager**.
2. In the left pane, click **File and Storage Services**, and then in the Servers pane, click **Storage Pools**.
3. Create a storage pool with the following settings:
 - Name: **StoragePool1**
 - Physical disks:
 - **PhysicalDisk3**
 - **PhysicalDisk4**
 - **PhysicalDisk5**
 - **PhysicalDisk6**
 - **PhysicalDisk7**

► Task 2: Create a three-way mirrored virtual disk

1. On LON-SVR1, in Server Manager, in the VIRTUAL DISKS pane, create a virtual disk with the following settings:
 - Storage pool: **StoragePool1**
 - Name: **Mirrored Disk**
 - Storage Layout: **Mirror**
 - Resiliency settings: **Three-way mirror**
 - Provisioning type: **Thin**
 - Virtual disk size: **10 GB**
2. In the New Volume Wizard, create a volume with the following settings:
 - Virtual disk: **Mirrored Disk**
 - Drive letter: **H**
 - File system: **ReFS**
 - Volume label: **Mirrored Volume**

► Task 3: Copy a file to the volume, and verify that it is visible in File Explorer

1. Open a Command Prompt window.
2. Type the following command:

```
Copy C:\windows\system32\write.exe H:\
```

3. Open File Explorer from the taskbar, and access **Mirrored Volume (H:)**. You should now see write.exe in the file list.

► Task 4: Remove a physical drive

- On Host machine, in Hyper-V Manager, in the Virtual Machines pane, change **20410C-LON-SVR1** settings to the following:
 - Remove Hard Drive that begins with **20410C-LON-SVR1-Disk5**.

► Task 5: Verify that the write.exe file is still accessible

1. Switch to LON-SVR1.
2. Open File Explorer and browse to **H:\write.exe** to ensure access to the file 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 displays next to Mirrored Disk.

4. Open the **Mirrored Disk Properties** dialog box, and access the Health pane.

Notice that the Health Status indicates a Warning. The Operational Status should indicate **Incomplete** or **Degraded**.

► Task 6: Add a new disk to the storage pool and remove a broken disk

1. Switch to **LON-SVR1**.
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**, click **Add Physical Disk**, and then click **PhysicalDisk8 (LON-SVR1)**.
4. Open Windows PowerShell and run the following commands to remove the disconnected disk.
 - a. **Get-PhysicalDisk** Note the FriendlyName for the disk that shows an OperationalStatus of Lost Communication.
 - b. **\$Disk = Get-PhysicalDisk -FriendlyName *diskname*** Replace *diskname* with the name of the disk you noted in previously.
 - c. **Remove-PhysicalDisk -PhysicalDisks \$disk _StoragePoolFriendlyName StoragePool1**
5. In Server Manager, refresh the storage pools view to see the warnings disappear.

Results: After completing this lab, you should have created a storage pool and added five disks to it. Then you should have created a three-way mirrored, thinly provisioned virtual disk from the storage pool. You should have also copied a file to the new volume and verified that it is accessible. Next, after removing a physical drive, you should have verified that the virtual disk was still available and could be accessed. Finally, you should have added another physical disk to the storage pool.

► Prepare for the next module

After you finish the lab, revert the virtual machines back 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 **20410C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20410C-LON-SVR1**.

Lab Review Questions

Question: At a minimum, how many disks must you add to a storage pool to create a three-way mirrored virtual disk?

Question: You have a USB-attached disk, four SAS disks, and one SATA disk that are attached to a Windows Server 2012 server. You want to provide a single volume to your users that they can use for file storage. What would you use?

Module Review and Takeaways

Module Review Questions

Question: Your current volume runs out of disk space. You have another disk available in the same server. What actions in the Windows operating system can you perform to help you add disk space?

Question: What are the two different types of disks in Disk Management?

Question: What are the most important implementations of RAID?

Question: You attach five 2 TB disks to your Windows Server 2012 computer. You want to manage them almost automatically, and if one disk fails, you want to make sure the data is not lost. What feature can you implement to accomplish this?

Best Practices

The following are recommended best practices:

- If you want to shrink a volume, defragment the volume first so you can reclaim more space from the volume.
- Use the GPT partition table format for disks larger than 2 TB.
- For very large volumes, use ReFS.
- Do not use FAT or FAT32 on Windows Server operating system disks.
- Use the Storage Spaces feature to have the Windows operating system manage your disks.

Tools

Tool	Use	Where to find it
Disk Management	<ul style="list-style-type: none"> • Initialize disks • Create and modify volumes 	In Server Manager on the Tools menu (part of Computer Management)
Diskpart.exe	<ul style="list-style-type: none"> • Initialize disks • Create and modify volumes from a command prompt 	Command prompt
Mklink.exe	<ul style="list-style-type: none"> • Create a symbolic link to a file or folder 	Command prompt
Chkdsk.exe	<ul style="list-style-type: none"> • Check a disk for a NTFS-formatted volume • Cannot be used for ReFS or virtual disks 	Command prompt
Defrag.exe	<ul style="list-style-type: none"> • Disk defragmentation tool for NTFS-formatted volumes. • Cannot be used for ReFS or virtual disks 	Command prompt

Module 10

Implementing File and Print Services

Contents:

Module Overview	10-1
Lesson 1: Securing Files and Folders	10-2
Lesson 2: Protecting Shared Files and Folders by Using Shadow Copies	10-15
Lesson 3: Configuring Work Folders	10-18
Lesson 4: Configuring Network Printing	10-25
Lab: Implementing File and Print Services	10-31
Module Review and Takeaways	10-38

Module Overview

Accessing files and printers on the network is one of the most common activities in the Windows Server® environment. Reliable, secure access to files and folders and print resources is often the first requirement of a Windows Server 2012-based network. To provide access to file and print resources on your network, you must understand how to configure these resources within Windows Server 2012 server, and how to configure appropriate access to the resources for users in your environment.

This module discusses how to provide these important file and print resources with Windows Server 2012. It describes how to secure files and folders, how to protect previous versions of files and folders by using shadow copies, and how to give workers remote access to corporate files by implementing the new Work Folders role service. It also describes new network printing features that help manage the network printing environment.

Objectives

After completing this module, you will be able to:

- Secure shared files and folders.
- Protect shared files and folders by using shadow copies.
- Configure the Work Folders role service.
- Configure network printing.

Lesson 1

Securing Files and Folders

The files and folders that your servers store typically contain your organization's business and functional data. Providing appropriate access to these files and folders, usually over the network, is an important part of managing file and print services in Windows Server 2012.

This lesson gives you information necessary to secure files and folders on your Windows Server 2012 servers, so that your organization's data is available yet protected.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe NTFS file system permissions.
- Describe a shared folder.
- Describe permissions inheritance.
- Explain how effective access and permissions work when you access shared folders.
- Describe access-based enumeration.
- Describe Offline files.
- Explain how to create and configure a shared folder.

What Are NTFS Permissions?

NTFS permissions are assigned to files or folders on a storage volume that is formatted with NTFS. The permissions that you assign to NTFS files and folders govern user access to these files and folders.

The following points describe the key aspects of NTFS permissions:

- NTFS permissions can be configured for an individual file or folder, or sets of files or folders.
- NTFS permissions can be assigned individually to objects that include users, groups, and computers.
- NTFS permissions are controlled by granting or denying specific types of NTFS file and folder access, such as Read or Write.
- NTFS permissions can be inherited from parent folders. By default, the NTFS permissions that are assigned to a folder are also assigned to newly created folders or files within that parent folder.

- NTFS permissions control access for files and folders on NTFS-formatted storage volumes

- NTFS Permissions:

- Are configured for files or folders
- Can be granted or denied
- Are inherited from parent folders

- Permissions conflict precedence:

1. Explicitly assigned Deny
2. Explicitly assigned Allow
3. Inherited Deny
4. Inherited Allow

NTFS Permission Types

There are two assignable NTFS permissions types: standard, and advanced.

Standard Permissions

Standard permissions provide the most commonly used permission settings for files and folders. You assign standard permissions in the Permissions for *folder name* dialog box.

The following table details the standard permissions options for NTFS files and folders.

File permissions	Description
Full Control	Grants the user complete control of the file or folder, including control of permissions.
Modify	Grants the user permission to read, write, or delete a file or folder, including creating a file or folder. It also grants permission to execute files.
Read and Execute	Grants the user permission to read a file and start apps.
Read	Grants the user permission to view file or folder content.
Write	Grants the user permission to write to a file.
List folder contents (folders only)	Grants the user permission to view a list of the folder's contents.



Note: Granting users Full Control permissions on a file or a folder gives them the ability to perform any file system operation on the object, and the ability to change permissions on the object. They can also remove permissions on the resource for any or all users, including you.

Advanced Permissions

Advanced permissions can provide a much greater level of control over NTFS files and folders. Advanced permissions are accessible by clicking the **Advanced** button from the **Security** tab of a file or folder's **Properties** dialog box.

The following table details the Advanced permissions for NTFS files and folders.

File permissions	Description
Traverse Folder/Execute File	<p>The Traverse Folder permission applies only to folders. This permission grants or denies users the right to browse through folders to reach other files or folders, even if the user has no permissions for the traversed folders. The Traverse Folder permission takes effect only when the group or user is not granted the Bypass Traverse Checking user right. By default, the Everyone group is given the Bypass Traverse Checking user right.</p> <p>The Execute File permission grants or denies access to run program files.</p> <p>If you set the Traverse Folder permission on a folder, the Execute File permission is not automatically set on all files in that folder.</p>
List Folder/Read Data	<p>The List Folder permission grants the user permission to view file names and subfolder names. This permission applies only to folders and affects only the contents of that folder—it does not affect whether the folder itself is listed. In addition, this setting has no effect on viewing the file structure from a command-line interface.</p> <p>The Read Data permission grants or denies the user permission to view data in files. The Read Data permission applies only to files.</p>
Read Attributes	The Read Attributes permission grants the user permission to view the basic attributes of a file or a folder such as Read-only and Hidden attributes. Attributes are defined by NTFS.

(continued from previous page)

File permissions	Description
Read Extended Attributes	The Read Extended Attributes permission grants the user permission to view the extended attributes of a file or folder. Extended attributes are defined by apps, and can vary by app.
Create Files/Write Data	The Create Files permission applies only to folders, and grants the user permission to create files in the folder. The Write Data permission grants the user permission to make changes to the file and overwrite existing content by NTFS. The Write Data permission applies only to files.
Create Folders /Append Data	The Create Folders permission grants the user permission to create folders within the folder. The Create Folders permission applies only to folders. The Append Data permission grants the user permission to make changes to the end of the file, but not to delete or overwrite existing data. The Append Data permission applies only to files.
Write Attributes	The Write Attributes permission grants the user permission to change the basic attributes of a file or folder, such as Read-only or Hidden. Attributes are defined by NTFS. The Write Attributes permission does not imply that you can create or delete files or folders; it includes only the permission to make changes to the attributes of a file or folder. To grant Create or Delete permissions, see the Create Files/Write Data, Create Folders/Append Data, Delete Subfolders and Files, and Delete entries in this table.
Write Extended Attributes	The Write Extended Attributes permission grants the user permission to change the extended attributes of a file or folder. Extended attributes are defined by programs and apps, and can vary by each one. The Write Extended Attributes permission does not imply that the user can create or delete files or folders; it includes only the permission to make changes to the attributes of a file or folder. To grant Create or Delete permissions, see the Create Files/Write Data, Create Folders/Append Data, Delete Subfolders and Files, and Delete entries in this table.
Delete Subfolders and Files	The Delete Subfolders and Files permission grants the user permission to delete subfolders and files, even if the Delete permission is not granted on the subfolder or file. The Delete Subfolders and Files permission applies only to folders.
Delete	The Delete permission grants the user permission to delete the file or folder. If you have not been assigned Delete permission on a file or folder, you can still delete the file or folder if you are granted Delete Subfolders and Files permissions on the parent folder.
Read Permissions	Read Permissions grants the user permission to read permissions about the file or folder, such as Full Control, Read, and Write.
Change Permissions	Change Permissions grants the user permission to change permissions on the file or folder, such as Full Control, Read, and Write.

(continued from previous page)

File permissions	Description
Take Ownership	The Take Ownership permission grants the user permission to take ownership of the file or folder. The owner of a file or folder can change permissions on it, regardless of any existing permissions that protect the file or folder.
Synchronize	The Synchronize permission assigns different threads to wait on the handle for the file or folder, and then synchronize with another thread that may signal it. This permission applies only to multiple-threaded, multiple-process programs and apps.



Note: Standard permissions are combinations of several individual Advanced permissions that are grouped into commonly used file and folder scenarios.

NTFS Permissions Examples

The following are basic examples of assigning NTFS permissions:

- For a folder called Marketing Pictures, an administrator has chosen to assign Adam Carter Allow permissions for the Read permission type. Under default NTFS permissions behavior, Adam Carter will have Read access to the files and folders that are contained in the Marketing Pictures folder.
- When applying NTFS permissions, the results are cumulative. For example, in the previous example, say that Adam Carter is also a part of the Marketing group. The Marketing group has been given Write permissions on the Marketing Pictures folder. When we combine the permissions assigned to Adam Carter's user account with the permissions assigned to the Marketing group, Adam will have both Read and Write permissions for the Marketing Pictures folder.

Important Rules for NTFS Permissions

There are two important groups of NTFS permissions:

- Explicit vs. Inherited. When you apply NTFS permissions, permissions that are explicitly applied to a file or a folder take precedence over those that are inherited from a parent folder.
- Deny vs. Allow. After NTFS permissions have been divided into explicit and inherited permissions, any existing Deny permissions override conflicting Allow permissions within the group.

Therefore, taking these rules into account, NTFS permissions are applied in the following order:

1. Explicit Deny
2. Explicit Allow
3. Inherited Deny
4. Inherited Allow

It is important to remember that NTFS permissions are cumulative, and these rules apply only when two NTFS permission settings conflict with each other.

How to Configure NTFS Permissions

You can view and configure NTFS permissions by following these steps:

1. Right-click the file or folder for which you want to assign permissions, and then click **Properties**.
2. In the **Properties** dialog box, click the **Security** tab.

3. In the **Security** tab, select the user or group that you want to view or edit the specific permissions of.
4. To modify existing permissions or add new users or groups, click the **Edit** button.

This opens the **Permissions** dialog box.

What Are Shared Folders?

Shared folders are a key component to granting access to files on your server from the network. When you share a folder, the folder and all of its contents are made available to multiple users simultaneously over the network. Shared folders maintain a separate set of permissions from the NTFS permissions, which apply to the folder's contents. These permissions provide an extra level of security for files and folders that are made available on the network.

Most organizations deploy dedicated file servers to host shared folders. You can store files in shared folders according to categories or functions. For example, you can put shared files for the Sales department in one shared folder, and shared files for the Marketing department in another.

- Shared folders are folders that grant network access to their contents



- Folders can be shared, but individual files cannot
- Accessing a shared folder using the UNC path:
 - \\LON-SVR1\Sales (standard share)
 - \\LON-SVR1\Sales\$ (hidden share)



Note: The sharing process applies only to the folder level. You cannot share an individual file or a group of files.

Accessing a Shared Folder

Users typically access a shared folder over the network by using its Universal Naming Convention (UNC) address. The UNC address contains the name of the server on which the folder is hosted, and the actual shared folder name, separated by a backward slash (\) and preceded by two backward slashes (\\). For example, the UNC path for the Sales shared folder on the LON-SVR1 server is \\LON-SVR1\Sales.

Sharing a Folder on the Network

Windows Server 2012 provides different ways to share a folder:

- Click the appropriate drive, and then in the Files and Storage Services section in Server Manager, click the **New Share** task.
- Use the File Sharing Wizard, either from the folder's shortcut menu, or by clicking the **Share** button on the **Sharing** tab of the folder's **Properties** dialog box.
- Use Advanced Sharing by clicking the **Advanced Sharing** button on the **Sharing** tab of the folder's **Properties** dialog box.
- Use the **Net share** command-line tool from a command-line window.
- Use the **New-SMBShare** cmdlet in Windows PowerShell®.



Note: When you are setting up a shared folder, you will be asked to give it a name. This name does not have to be the same name as the actual folder; it can be a descriptive name that better describes the folder contents to network users.

Administrative Shares

If you have shared folders that need to be available from the network, but should be hidden from users browsing the network, you can create administrative (or hidden) shared folders. You can access an administrative shared folder by typing in its UNC path, but the folder will not be visible if you browse the server by using File Explorer. Administrative shared folders also typically have a more restrictive set of permissions to reflect the administrative nature of the folder's contents.

To hide a shared folder, append the dollar symbol (\$) to the folder's name. For example, a shared folder on LON-SVR1 named Sales can be made into a hidden shared folder by naming it Sales\$. The shared folder is accessible over the network by using the UNC path \\LON-SVR1\Sales\$.



Note: Shared folder permissions apply only to users who access the folder over the network. They do not affect users who access the folder locally on the computer where the folder is stored.

Shared Folder Permissions

Just like NTFS permissions, you can assign shared folder permissions to users, groups, or computers. However, unlike NTFS permissions, shared folder permissions are not configurable for individual files or folders within the shared folder. Shared folder permissions are set once for the shared folder itself, and apply universally to the entire contents of the shared folder for users who access the folder over the network.

When you create a shared folder, the default assigned shared permission for the Everyone group is set to Read.

The following table lists the permissions that you can grant to a shared folder.

Shared folder permission	Description
Read	Users can view folder and file names, view file data and attributes, run program files and scripts, and navigate the folder structure within the shared folder.
Change	Users can create folders, add files to folders, change data in files, append data to files, change file attributes, delete folders and files, and perform all tasks permitted by the Read permission.
Full Control	Users can change file permissions, take ownership of files, and perform all tasks permitted by the Change permission.



Note: When you assign Full Control permissions on a shared folder to a user, that user can modify permissions on the shared folder, which includes removing all users (including administrators), from the shared folder's permissions list. In most cases, you should grant Change Permission instead of Full Control permission.

Permissions Inheritance

By default, NTFS and shared folders use inheritance to propagate permissions throughout a folder structure. When you create a file or a folder, it is automatically assigned the permissions that are set on any folders that exist above it (parent folders) in the hierarchy of the folder structure.

How Inheritance Is Applied

Consider the following example. Adam Carter is a member of the Marketing group and the New York Editors group. The following table is a summary of the permissions for this example:

- Inheritance is used to manage access to resources without assigning explicit permissions to each object
- By default, permissions are inherited in a parent/child relationship
- Blocking inheritance:
 - You can block permission inheritance
 - You can apply blocking at the file or folder level
 - You can set blocking on a folder to propagate the new permissions to child objects

Folder or File	Assigned Permissions	Adam's Permissions
Marketing (folder)	Read – Marketing	Read
Marketing Pictures (folder)	None set	Read (inherited)
New York (folder)	Write – New York Editors	Read(i) + Write
Fall_Composite.jpg (file)	None set	Read(i) + Write(i)

In this example, Adam is a member of two groups that are assigned permissions for files or folders within the folder structure. They are as follows:

- The top-level folder, Marketing, has an assigned permission for the Marketing Group giving them Read access.
- In the next level, the Marketing Pictures folder has no explicit permissions set, but because of permissions inheritance Adam has Read access to this folder and its contents from the permissions that are set on the Marketing folder.
- In the third level, the New York folder has Write permissions assigned to one of Adam's groups—New York Editors. In addition to this explicitly assigned Write permission, the New York folder also inherits the Read permission from the Marketing folder. These permissions pass down to file and folder objects, cumulating with any explicit Read and Write permissions set on those files.
- The fourth and last level is the Fall_Composite.jpg file. Even though no explicit permissions have been set for this file, Adam has both Read and Write access to the file due to the inherited permissions from both the Marketing folder and the New York folder.

Permission Conflicts

Sometimes, explicitly set permissions on a file or folder conflict with permissions inherited from a parent folder. In these cases, the explicitly assigned permissions always override the inherited permissions. In the given example, if Adam Carter was denied Write access to the parent Marketing folder, but then explicitly granted Write access to the New York folder, the granted Write access permissions take precedence over the inherited deny Write access permission.

Blocking Inheritance

You can also disable the inheritance behavior for a file or a folder (and its contents) on an NTFS drive. You do this when you want to explicitly define permissions for a set of objects without including any of the

inherited permissions from any parent folders. Windows Server 2012 provides an option for blocking inheritance on a file or a folder. To block inheritance on a file or folder, complete the following steps:

1. Right-click the file or folder where you want to block inheritance, and then click **Properties**.
2. In the **Properties** dialog box, click the **Security** tab, and then click **Advanced**.
3. In the **Advanced Security Settings** dialog box, click **Change Permissions**.
4. In the next **Advanced Security Settings** dialog box, click **Disable inheritance**.
5. At this point, you are prompted to either convert the inherited permissions into explicit permissions or remove all inherited permissions from the object to start with a blank permissions slate.

Resetting Default Inheritance Behavior

After you block inheritance, changes made to permissions on the parent folder structure no longer have an effect on the permissions for the child object (and its contents) that has blocked inheritance, unless you reset that behavior from one of the parent folders by selecting the Replace All Child Objects With Inheritable Permissions From This Object option. When you select this option, the existing set of permissions on the current folder are propagated down to all child objects in the tree structure, and override all explicitly assigned permissions for those files and folders. This check box is located directly under the Include Inheritable Permissions From This Object's Parent check box.

Effective Permissions

Access to a file or folder in Windows Server 2012 is granted based on a combination of permissions. When a user attempts to access a file or folder, the permission that applies is dependent on various factors, including:

- Explicitly defined and inherited permissions that apply to the user
- Explicitly defined and inherited permissions that apply to the groups to which the user belongs
- How the user is accessing the file or folders: locally, or over the network

- When combining shared folder and NTFS permissions, the most restrictive permission is applied
 - Example: If a user or group is given the shared folder permission of Read and the NTFS permission of Write, the user or group will only be able to read the file because it is the more restrictive permission
- Both the share and the NTFS file and folder permissions must have the correct permissions, otherwise the user or group will be denied access to the resource

Effective NTFS permissions are the cumulative permissions that are assigned to a user for a file or folder based on the factors listed above. The following principles determine effective NTFS permissions:

- Cumulative permissions are the combination of the highest NTFS permissions granted to the user and to all the groups of which the user is a member. For example, if a user is a member of a group that has Read permission and is a member of a group that has Modify permission, the user is assigned cumulative Modify permissions.
- Deny permissions override equivalent Allow permissions. However, an explicit Allow permission can override an inherited Deny permission. For example, if a user is denied Write access to a folder via an inherited Deny permission, but is explicitly granted Write access to a subfolder or a particular file, the explicit Allow overrides the inherited Deny for the particular subfolder or file.
- You can apply permissions to a user or to a group. Assigning permissions to groups is preferred because they are more efficient than managing permissions that are set for many individuals.

- NTFS file permissions take priority over folder permissions. For example, if a user has Read permission to a folder, but has been granted Modify permission to certain files in that folder, the effective permission for those files will be set to Modify.
- Every object in an NTFS drive or in Active Directory® Domain Services (AD DS) is owned. The owner controls how permissions are set on the object and to whom permissions are granted. For example, a user who creates a file in a folder where they have Modify permissions can change the permissions on the file to Full Control.

Effective Access Tool

Windows Server 2012 provides an Effective Access tool that shows the effective NTFS permissions on a file or folder for a user, based on permissions assigned to the user account and groups to which the user account belongs. You can access Effective Access tool by the following steps:

1. Right-click the file or folder for which you want to analyze permissions, and then click **Properties**.
2. In the **Properties** dialog box, click the **Advanced** button.
3. In the **Advanced Security Settings** dialog box, click the **Effective Access** tab.
4. Choose a user or group to evaluate by using **Select a user**.

Combining NTFS Permissions and Shared Folder Permissions

NTFS permissions and shared folder permissions work together to control access to file and folder resources that are accessed from a network. When you configure access to network resources on an NTFS drive, use the most restrictive NTFS permissions to control access to folders and files, and combine them with the most restrictive shared folder permissions to control access to the network.

How Combining NTFS and Shared Folder Permissions Works

When you apply both NTFS and shared folder permissions, remember that the more restrictive of the two permissions dictates the access that a user will have to a file or folder. The following two examples explain this further:

- If you set the NTFS permissions on a folder to Full Control, but you set the shared folder permissions to Read, then that user has only Read permission when accessing the folder over the network. Access is restricted at the shared folder level, and any greater access at the NTFS permissions level does not apply.
- Likewise, if you set the shared folder permission to Full Control, and you set the NTFS permissions to Write, then the user will have no restrictions at the shared folder level, but the NTFS permissions on the folder grants only Write permissions to that folder.

The user must have appropriate permissions on both the NTFS file or folder and the shared folder. If no permissions exist for the user (either as an individual or as the member of a group) on either resource, access is denied.

Considerations for Combined NTFS and Shared Folder Permissions

The following are several considerations that make administering permissions more manageable:

- Grant permissions to groups instead of users. Groups can always have individuals added or deleted, while permissions on a case-by-case basis are difficult to track and cumbersome to manage.
- Use Deny permissions only when necessary. Because Deny permissions are inherited, assigning deny permissions to a folder can result in users not being able to access files further down in the folder structure tree. You should assign Deny permissions only in the following situations:
 - To exclude a subset of a group that has Allow permissions

- To exclude one specific permission when you have granted Full Control permissions to a user or a group
- Never deny the Everyone group access to an object. If you deny the Everyone group access to an object, you deny Administrators access—including yourself. Instead, remove the Everyone group from the permissions list, as long as you grant permissions for the object to other users, groups, or computers.
- Grant permissions to an object that is as high in the folder structure as possible, so that the security settings are propagated throughout the tree. For example, instead of bringing groups representing all departments of the company together into a Read folder, assign Domain Users (which is a default group for all user accounts on the domain) to the share. In this manner, you eliminate the need to update department groups before new users receive the shared folder.
- Use NTFS permissions instead of shared permissions for fine-grained access. Configuring both NTFS and shared folder permissions can be difficult. Consider assigning the most restrictive permissions for a group that contains many users at the shared folder level, and then use NTFS permissions to assign permissions that are more specific.

What Is Access-Based Enumeration?

With access-based enumeration, users see only the files and folders which they have permission to access. Access-based enumeration provides a better user experience because it displays a less complex view of the contents of a shared folder, making it easier for users to find the files that they need. Windows Server 2012 allows access-based enumeration of folders that a server shares over the network.

- Access-based enumeration allows an administrator to control the visibility of shared folders according to the permissions set on the shared folder
- Access Based Enumeration is:
 - Built into Windows Server 2012
 - Available for shared folders
 - Configurable on a per shared folder basis

Enabling Access-Based Enumeration

To enable access-based enumeration for a shared folder:

1. Open Server Manager.
2. In the navigation pane, click **File and Storage Services**.
3. In the navigation pane, click **Shares**.
4. In the Shares pane, right-click the shared folder for which you want to enable access-based enumeration, and then click **Properties**.
5. In the **Properties** dialog box, click **Settings**, and then select **Enable access-based enumeration**.

When **Enable access-based enumeration** is selected, access-based enumeration is enabled on the shared folder. This setting is unique to each shared folder on the server.



Note: The File and Storage Services console is the only place in the Windows Server 2012 interface where you can configure access-based enumeration for a shared folder. Access-based enumeration is not available in any of the properties dialog boxes that are accessible by right-clicking the shared folder in File Explorer.

What Are Offline Files?

An *offline file* is a copy of a network file that is stored on a client computer. By using offline files, users can access network-based files when their client computer is disconnected from the network.

If offline files and folders have been edited or modified by the client, then the changes are synchronized with the network copy of the files the next time the client reconnects to the network. The synchronization schedule and behavior of offline files is controlled by the Windows client operating system.

Offline files are available with the following operating systems:

- Windows 8.1
- Windows 8
- Windows Server 2012 R2
- Windows Server 2012
- Windows 7
- Windows Server 2008 R2
- Windows Server 2008
- Windows Vista®
- Windows Server 2003

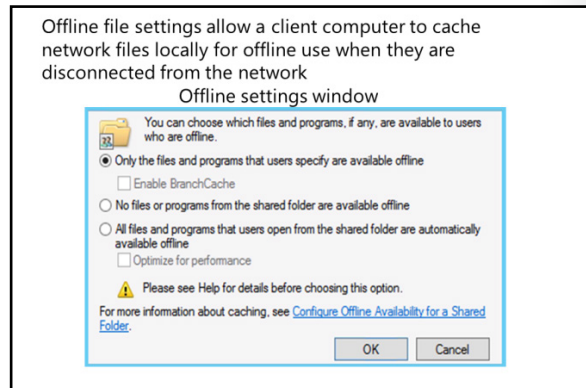


Note: Offline files is not available in home versions of Windows operating systems.

Offline Settings

With Windows Server 2012, you view the Offline Settings dialog box for a shared folder by clicking the Caching button in the Advanced Sharing dialog box. The following options are available within the Offline Settings dialog box:

- Only the files and programs that users specify are available offline. This is the default option when you set up a shared folder. When you use this option, no files or programs are available offline by default, and users control which files and programs they want to access when they are not connected to the network. Alternatively, you can choose the **Enable BranchCache** option. This option enables computers that are accessing the files to cache files downloaded from the folder using Windows BranchCache®. You must install and configure BranchCache on the Windows Server 2012 server to select this option.
- No files or programs from the shared folder are available offline. This option blocks client computers from making copies of the files and programs on the shared folder.
- All files and programs that users open from the shared folder are automatically available offline. Whenever a user accesses the shared folder or drive and opens a file or program in it, that file or program is made automatically available offline to that user. Files and programs that are made automatically available offline remain in the offline files cache and synchronize with the version on the server until the cache is full or the user deletes the files. Files and programs that are not opened are not available offline.



- Optimized for performance. If you select this option, executable files (.exe, .dll) that are run from the shared folder by a client computer are cached on that client computer automatically. The next time the client computer runs the executable files, it will access its local cache instead of the shared folder on the server.



Note: The Offline Files feature must be enabled on the client computer for files and programs to be cached automatically. In addition, the Optimized For Performance option does not have any effect on client computers that use Windows Vista or older Windows operating systems, because these operating systems perform the program-level caching automatically, as specified by this option.

The Always Offline Mode

You can configure Windows Server 2012 and Windows 8 computers to use the Always Available Offline Mode when accessing shared folders. When you configure this option, client computers always use the locally cached version of the files from a network share, even if they are connected to the file server by a high-speed network connection.

This configuration typically results in faster access to files for client computers, especially when connectivity or speed of a network connection is intermittent. Synchronization with the files on the server occurs according to the offline files configuration of the client computer.

How to Enable the Always Offline Mode

To enable Always Offline mode, use Group Policy to enable the **Configure slow-link mode** setting, and set the latency value to **1**. The **Configure slow-link mode** setting is located in Group Policy under the Computer Configuration\Administrative Policies\Network\Offline Files node.

Demonstration: Creating and Configuring a Shared Folder

Creating and configuring a shared folder is typically done within File Explorer, from the file or folder's **Properties** dialog box on the **Sharing** tab. When creating a shared folder, always ensure that you set permissions that are appropriate for all of the files and folders within the shared folder location.

In this demonstration, you will see how to:

- Create a shared folder.
- Assign permissions for the shared folder.
- Configure access-based enumeration.
- Configure offline files.

Demonstration Steps

Create a shared folder

1. Sign in to LON-SVR1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On drive E, create a folder named **Data**.
3. Share the Data folder.

Assign permissions for the shared folder

- Grant the Authenticated Users Change permissions for the Data folder.

Configure access-based enumeration

1. Open Server Manager.
2. Navigate to the Share pane in the File and Storage Services management console.
3. Open the **Data Properties** dialog box for **\\LON-SVR1\Data**, and enable access-based enumeration.

Configure offline files

1. Open the **Data Properties** dialog box for **E:\Data**.
2. Navigate to the **Sharing** tab, and open the Advanced Sharing settings.
3. Open the Caching settings, and then disable offline files.

Lesson 2

Protecting Shared Files and Folders by Using Shadow Copies

You use shadow copies to restore previous versions of files and folders. It is much faster to restore a previous version of a file from a shadow copy than from a traditional backup copy, which might be stored offsite. Files and folders can be recovered by administrators, or directly by end users.

This lesson introduces you to shadow copies, and shows you how to configure a schedule of shadow copies in Windows Server 2012.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe shadow copies.
- Describe considerations for scheduling shadow copies.
- Identify methods for restoring data from shadow copies.
- Restore data from a shadow copy.

What Are Shadow Copies?

A *shadow copy* is a static image (or a *snapshot*) of a set of data, such as a file or folder. Shadow copies provide the capability to recover files and folders based on snapshots of storage drives. After a snapshot is taken, you can view and potentially restore previous versions of files and folders from that snapshot.

A shadow copy does not make a complete copy of all files for each snapshot. Instead, after a snapshot is taken, Windows Server 2012 tracks changes to the drive. A specific amount of disk space is allocated for tracking the changed disk blocks. When you access a previous version of a file, some of the content might be in the current version of the file, and some might be in the snapshot.

By default, the changed disk blocks are stored on the same drive as the original file, but you can modify where they are stored. You can also define how much disk space is allocated for shadow copies. Multiple snapshots are retained until the allocated disk space is full, after which, older snapshots are removed to make room for new snapshots. The amount of disk space that is used by a snapshot is based on the size of disk changes between snapshots.

Because a snapshot is not a complete copy of files, you cannot use shadow copies as a replacement for traditional backups. If the disk containing a drive is lost or damaged, then the snapshots of that drive are also lost.

Shadow copies are suitable for recovering data files, but not for more complex data (such as databases), that need to be logically consistent before a backup is performed. A database that is restored from previous versions is likely to be corrupt and require database repairs.

- Allow access to previous versions of files
- Are based on tracking disk changes
 - Disk space is allocated on the same volume
 - When the space is full, older shadow copies are removed
- Are not a replacement for backups
- Are not suitable for recovering databases

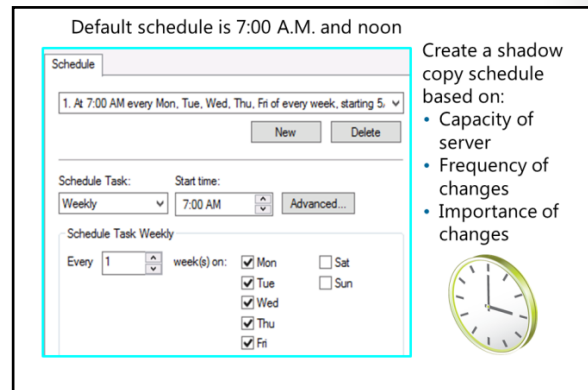


Considerations for Scheduling Shadow Copies

The default schedule for creating shadow copies is Monday through Friday at 07:00 A.M., and again at noon. You can modify the default schedule as desired for your organization.

When scheduling shadow copies:

- Consider that increasing the frequency of shadow copies increases the load on the server. As a best practice, you should not schedule drive shadow copies more than once each hour.
- Increase the frequency of shadow copies for frequently changing data. This increases the likelihood that recent file changes are captured.
- Increase the frequency of shadow copies for important data. This increases the likelihood that recent file changes are captured.



Restoring Data from a Shadow Copy


Previous versions of files can be restored by either users or administrators. Most users are unaware that they can do this, and they will need instructions on how to restore a previous version of a file.

Administrators can access and restore previous versions of files directly on the server that stores the files. Users can access and restore previous versions of files over the network from a file share. In both cases, previous versions are accessed from the Properties dialog box of the file or folder.

When viewing previous versions of a folder, you can browse the available files and select only the file that you need. If multiple versions of files are available, you can review each version before deciding which one to restore. Finally, you can copy a previous version of a file to an alternate location instead of restoring it to its previous location. This prevents overwriting the current file version.

Windows Vista and Windows 7 operating system clients can access previous file versions without installing any additional software. The ability to access previous file versions is no longer supported in Windows operating systems before Windows Vista.

- Previous versions are accessible from the Properties dialog box of a file or folder
 - Administrators can restore previous versions directly on the server
 - Users can restore previous versions over the network
- All users can:
 - Restore a file or folder
 - Browse previous versions to select the correct one
 - Copy a file or folder to an alternate location



Demonstration: Restoring Data from a Shadow Copy

You can create shadow copies using the default schedule, or you can modify the schedule to provide more frequent snapshots. In either case, you will only see the versions of the file as it has changed. Taking a shadow copy of a file that has not changed has no actual effect on the shadow copy. No additional versions are available, and no space is used in the snapshot for that particular file.

In this demonstration, you will see how to:

- Configure shadow copies.
- Create a new file.
- Create a shadow copy.
- Modify the file.
- Restore the previous version.

Demonstration Steps

Configure shadow copies

1. On LON-SVR1, open File Explorer.
2. Enable Shadow Copies for **Local Disk (C:)**.

Create a new file

1. Open File Explorer.
2. Create a folder on drive C named **Data**.
3. Create a text file named **TestFile.txt** in the Data folder.
4. Change the contents of **TestFile.txt** by adding and saving the text **Version 1**.

Create a shadow copy

1. In File Explorer, right-click **Local Disk (C:)**, and then click **Configure Shadow Copies**.
2. In the **Shadow Copies** dialog box, click **Create Now**.
3. When the shadow copy is complete, click **OK**.

Modify the file

1. Open **TestFile.txt** as a Notepad document.
2. In Notepad, type **Version 2**.
3. Save the changes.

Restore the previous version

1. In File Explorer, right-click **TestFile.txt**, and then click **Restore previous versions**.
2. Choose the most recent version.
3. In the **Are you sure you want to restore** message, click **Restore**.
4. Open **TestFile.txt** and verify that the previous version is restored.

Lesson 3

Configuring Work Folders

More and more, information workers want the ability to use their own devices such as smart phones and tablets to access corporate data files while out of the office. The Work Folders role service address this by allowing users to store and access work files from anywhere while complying with corporate policies. This is accomplished by synchronizing corporate data to user devices from a centralized, on-premise server using a new synchronization protocol. The corporate organization still maintains control of the data by implementing policies such as encryption.

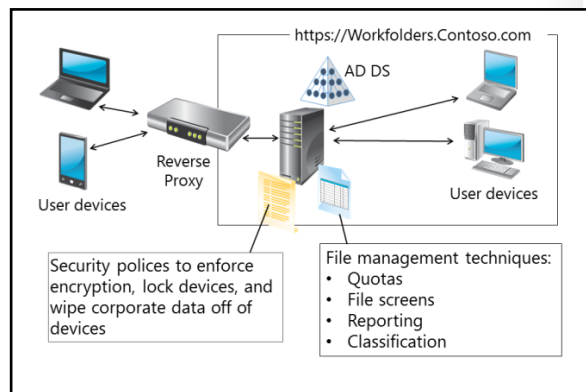
Lesson Objectives

After completing the lesson, you will be able to:

- Describe the Work Folders role service.
- Discuss the benefits and limitations of Work Folders.
- Describe Work Folders components.
- Configure Work Folders.

What is the Work Folders Role Service?

Work Folders is a new role service of the File and Storage Services role and is only available in Windows Server 2012 R2. Work Folders allows users to synchronize corporate data to all of their devices. When a user creates or modifies a file in a Work Folders folder on any device or PC, it is automatically replicated (using Secure Sockets Layer (SSL) connections on port 443) to a folder known as the *sync share* on the corporate file server. The changes in the sync share are then securely replicated to that user's other devices if those devices are configured to use Work Folders. A sync share maps to a physical location on the file server where files are stored. New folders or existing shared folders can be mapped to sync shares.



Client computers can be configured to connect to the sync share either manually or automatically. Once configured on the client computer, the work folder appears to the user like any other folder in File Explorer. Users can create files and folders in the work folder just as they do in any other network shared folder. These files and folders will be synchronized to all other devices configured to use Work Folders.

Other factors to keep in mind when working with Work Folders are:

- Corporate security policies can be applied to the data to enforce encryption, lock devices, and wipe corporate data off of devices.
- File management technologies such as quotas, file screens, reporting, and classification can be applied to files and folders held in Work Folders.
- Client devices are limited to one synchronize partnership per user per device.

How Files Stay In Synchronization

Once the synchronize partnership is established between the client and the server, a data directory is created on the NTFS volume on the client device. A hidden version database is also created and stored in the user profile—this database tracks the metadata of the files and folders stored in the work folders and detects when changes occur. A hidden download staging directory is used to accept updated files from the Work Folders server.

The first time a user synchronizes a device, a data directory and upload staging directory is created on the server for that user. One version-database is created on the sync share for each user. Synchronization occurs through change detection on the client or by polling. Polling is done every 10 minutes by default.

When a local change is detected on a device, the client connects to the server and uploads the change to the upload staging directory. Then the change is applied to the user's data directory on the server. Synchronization is always initiated by the client device.

Conflict Resolution

If a file is edited and saved on different devices at the same time, both copies will be uploaded to the server. One of the file names will have the name of the device it was saved on appended to it. For example, a user opens, edits, and saves a file named *Doc1* on his office PC; he then edits the offline version on his tablet. When the tablet version synchronizes, the file will be saved as *Doc1 name of tablet*. There will now be two versions of the file in the sync share.

Backup and Recovery

Files can be selectively restored either on the server or the client. The restored file is seen as just another change and becomes the authoritative version that will be synchronized to the other devices.

When you are backing up client systems, do not backup the version database; it will rebuild itself from the server.

For server disaster scenarios, the Volume Shadow Copy Service (VSS) writer supports a full server restore. Because synchronizations are initiated by the client the database becomes current automatically after receiving updates from clients.

Comparing Work Folders to Cloud-based Storage

For organizations that want to maintain data storage on-premise and already have established practices around data management and storage, Work Folders provides a solution that users will find familiar. Cloud-based technologies such as SkyDrive Pro® are good solutions for organizations that use SharePoint and need the collaboration features of Office 365.

Benefits and Limitations of Work Folders

Work Folders provides a number of benefits that cannot be achieved using existing technologies, but there are limitations to what Work Folders can do.

Benefits

Work Folders provides the following benefits:

- It is available for domain joined and non-domain joined systems. Users need to provide credentials to connect from non-domain joined devices.

- The benefits of Work Folders include:
 - Works on domain joined and non-domain joined systems
 - Provides a single point of access to work files
 - Provides offline access to work files
 - Synchronizes files for users
 - Data can be encrypted
 - Compatible with existing data management technologies
- Work Folders has the following limitations:
 - Currently supported on Windows Server 2012 R2 and Windows 8.1
 - Does not support collaboration scenarios
 - Files cannot be selectively synchronized
 - Does not support synchronizing multiple file shares

- It provides a single point of access to work files on a user's work and personal computers and devices.
- It provides access to work files while users are offline.
- It synchronizes files for the users when the computer or device next has Internet or network access.
- It can be deployed alongside existing technologies such as Folder Redirection and Offline Files.
- Data can be encrypted while in-transit as well as when it is sitting on the device itself.
- Administrators can configure security policies. These policies may include to instruct user computers and devices to encrypt work folders and to use a lock-screen password.
- It can use existing file server management technologies such as file classification and folder quotas to manage user data.
- Failover Clustering can be used to ensure high-availability.

Limitations

Work Folders has the following limitations:

- Work Folders are currently supported on Windows Server 2012 R2 and Windows 8.1 only.
- It does not provide collaboration functionality such as sharing synchronized files or folders with other users.
- There is no ability to selectively synchronize files in work folders; all files are synchronized.
- Users synchronized to their own folder on the file server—there is no support for synchronizing to other file shares.

Work Folder Components

In order to implement Work Folders there are specific software requirements and both server and client side components that need to be configured.

Software Requirements

Work Folders has the following software requirements for file servers:

- A server running Windows Server 2012 R2 for hosting sync shares and user data
- An NTFS formatted volume to store user files
- A server certificate from a certification authority (CA) that is trusted by your users – a public CA is best

- Software Requirements
 - Windows Server 2012 R2 file server
 - Windows 8.1 client
 - SSL certificates
 - NTFS file system for both client and server
- Server Components
 - Work Folders role service
 - File Server role service
 - Web Server Role (IIS) role
 - IIS Management Console role service
 - IIS Hostable Web Core role service
- Client Components
 - Manual deployment using built-in Control Panel item
 - Automatic deployment via Group Policy, Configuration Manager, or Intune

To enable users to synchronize across the Internet, there are additional requirements:

- The file server must be accessible from the Internet
- A publicly registered domain name and associated Domain Name System (DNS) records

Work Folders has the following software requirements for client computers:

- Windows 8.1
- Windows RT 8.1
- An NTFS formatted volume to store user files



Note: A Windows Server 2012 R2 cannot be a client of Work Folders.

Server Components

Work Folders is a role service of the File and Storage Services role. The Work Folders role service can be installed on any edition of Windows Server 2012 R2 and can be installed alongside any other roles or programs. For example, a domain controller or Exchange server can also host Work Folders.

Installing the Work Folders role service also installs the following roles and role services:

- File Server role service
- Web Server Role (Internet Information Services (IIS)) role
- IIS Management Console role service
- IIS Hostable Web Core role service

Once the role service is installed, the sync share must be created. You can create multiple sync shares on a file server. Each one maps to different file system locations with different users and groups having access and different policies defined per share.

Client Components

The ability to connect to and manage Work Folders folders and files is built into Windows 8.1. Deployment can be manual or automatic.

Manual Deployment

A built-in item in Control Panel named Work Folders is used to supply the user's corporate email address. This email address is used to construct the URL for the Work Folders server and that URL is used to connect to the Work Folders folder. If there is no corporate email address the URL can be entered manually.

Opt-in Deployment

Work Folders settings can be delivered via Group Policy, Microsoft System Center 2012 Configuration Manager or by Windows Intune™. After the settings are delivered, the user can then decide if they want to use Work Folders on that device.

Mandatory Deployment

Settings are delivered via Group Policy, System Center 2012 Configuration Manager or by Windows Intune™. No user action is required. Work Folders is configured on the device.

Configuring Work Folders

There are a number of steps on both the server and a client that must be completed in order to successfully configure Work Folders.

Server Configuration

Server configuration begins with the addition of the Work Folders role service and then the configuration of the sync share as outlined in the following steps:

1. Use Server Manager or Windows PowerShell to add the Work Folders role service and dependent role services.

The following Windows PowerShell command adds the Work Folders role service:

```
Add-WindowsFeature FS-SyncShareService
```

2. Use the New Sync Share Wizard or Windows PowerShell to create a sync share. You must provide the following information:
 - The name of the server that will host the sync share.
 - The path to the sync share. This is a path to a local folder or an existing shared folder on the local server. If you are using an existing shared folder then the work folders can also be accessed by the UNC path.
 - The folder naming format. This is in the form of an email address or a user alias. The user alias is compatible with technologies such as home folders. You can also specify that only a subfolder of the sync share will be synchronized.
 - The name of the sync share. This is the friendly name the sync share is known by.
 - The names of the users or groups that will have access to the sync share. By default, inherited permissions on the user folders is disabled and the user is granted exclusive access to the folder, but you can change that.
 - You can specify whether to encrypt the work folders and whether to automatically lock the screen and require a password.

Windows PowerShell cmdlets **New-SyncShare** and **Set-SyncShare** are used to create and modify the sync share. The following example creates a sync share named SalesShare at the local path of C:\SalesShare, grants access to the Sales group, and sets the conflict resolution method to keep the latest file saved.

```
New-SyncShare SalesShare -path C:\SalesShare -User Contoso\Sales -ConflictResolution KeepLatest
```

The server must have an SSL certificate installed in the computer's Trusted Root Certification Authority. The common name (CN) in the certificate must match the name used as the Work Folders URL. For example, if the client is making a request to <https://syncsvr.contoso.com>, then the CN must also be <https://syncsvr.contoso.com>.



Note: A single file server can host multiple sync shares, which will require a certificate with multiple hostnames (a Subject Alternative Name (SAN) certificate can be used).

- Server configuration
 - Install the Work Folders role service
 - Create a sync share
 - Install a server certificate named to support the Work Folders URL
- Client configuration
 - Manual configuration requires the user to manually enter their email address
 - Automatic configuration can be done by using Group Policy

Client Configuration

Clients can be configured either manually or automatically. In either case the Work Folders connection uses SSL, therefore clients must trust the server certificate. Although it is possible to use an internal CA, those certificates are not normally trusted by non-domain joined devices. As a best practice, purchase the server certificate from a public CA.



Reference Links: For more information about certificates for Work Folders see the article "Work Folders Certificate Management" at <http://go.microsoft.com/fwlink/?LinkID=331094>.

Manual Configuration

This requires the user to launch the Work Folders item in Control Panel and enter their corporate email address. This address is used to build the URL (by default HTTPS://FQDN) of the file server, which will connect the user to Work Folders. If the URL cannot be discovered by using the user's email address, the URL can be input manually.

Automatic Configuration using Group Policy

Automatic configuration can be accomplished via Group Policy. The following Group Policy settings are used for configuring client computers:

Setting	Description
Force automatic setup for all users	This computer configuration setting specifies whether Work Folders will be set up automatically for all users on this computer. This prevents users from manually specifying the local folder in which files are stored. Work Folders uses the settings specified in the user Group Policy configuration for Work Folders.
Specify Work Folders settings	This user configuration setting specifies the Work Folders server as well as whether or not the users are allowed to change settings on domain-joined computers. When enabled, users receive settings for the Work Folders URL and they can be prevented from manually specifying the local folder in which work folders are stored. The default location is %userprofile%\Work Folders.



Note: Automatic configuration using System Center 2012 Configuration Manager or Windows Intune is beyond the scope of this course.

Demonstration: How to Configure Work Folders

In this demonstration, you will see how to:

- Install the Work Folders role service.
- Create a sync share for work folders on a file server.
- Configure Work Folder access on a Windows 8.1 client.
- Create a file in the work folder.
- Configure Work Folders to synchronize data on a second Windows 8.1 client.

Demonstration Steps

Install the Work Folders role service

- On LON-SVR1, install the Work Folders role service.

Create a sync share on a file server

- In Server Manager, in File and Storage Services use the New Sync Share Wizard to create a new sync share with the following parameters:
 - Server Name: **LON-SVR1**
 - Select by file share: **Data**
 - Structure for user folders: **User alias**
 - Sync share name: **WorkFolders**
 - Grant synchronize access to groups: **Domain Users**
 - Device policies: **Automatically lock screen, and require a password**

Configure Work Folder access on a Windows 8.1 client

1. Sign in to **LON-CL1** as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Navigate to **C:\Labfiles\Mod10** and run the **WorkFolders.bat**.
This batch file adds a registry entry that allows unsecured connections to work folders.
3. Open **Control Panel** and, in System and Security, open the **Work Folders** item.
4. Setup Work Folders as follows:
Click **Enter a Work Folders URL instead**.
 - Work Folders URL: **http://lon-svr1.adatum.com**
Normally this requires a secure connection
 - Work Folders location: **Accept default**
 - Policies: **Accept the policies**
5. Set up the Work Folders folder.
6. Open **File Explorer**.
Notice there is now a Work Folders folder under the This PC folder.

Create a file in the work folder

- Open the Work Folders folder and create a new text document.

Synchronize data on a second client computer

1. Sign in to **LON-CL2** as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Navigate to **C:\Labfiles\Mod10** and double-click **SetIP.bat**.
This configures the IP address of the client to be on the correct subnet.
3. Repeat steps 2 through 6 from the **Configure Work Folder access on a Windows 8.1 Client** task.
4. Open the **Work Folders** folder and notice the file you created is available from this computer.
5. Close all open windows.
6. Use Hyper-V Manager on the host computer to Revert **20410C-LON-CL2**.

Lesson 4

Configuring Network Printing

By using the Print and Document Services role in Windows Server 2012, you can share printers on a network and centralize print server and network printer management. By using the Print Management console, you can monitor print queues, and receive important notifications regarding print server activity.

Windows Server 2012 introduces new features and important changes to the Print and Document Services role that you can use to manage your network printing environment better. This lesson explains the important aspects of network printing, and introduces new network printing features that are available in Windows Server 2012.

Lesson Objectives

After completing the lesson, you will be able to:

- Identify the benefits of network printing.
- Describe Enhanced Point and Print.
- Identify security options for network printing.
- Create multiple configurations for a print device.
- Describe printer pooling.
- Describe Branch Office Direct Printing.
- Identify methods for deploying printers to clients.

Benefits of Network Printing

You can configure network printing by using Windows Server 2012 as a print server for users. In this configuration, client computers submit print jobs to the print server for delivery to a printer that is connected to the network.

Benefits of Network Printing

- Centralized management. The biggest benefit of using Windows Server 2012 as a print server is centralized management of printing. Instead of managing client connections to many individual devices, you manage their connection to the server. You install printer drivers centrally on the server, and then distribute them to workstations.
- Simplified troubleshooting. By installing printer drivers centrally on a server, you also simplify troubleshooting. It is relatively easy to determine whether printing problems are caused by the printer, server, or client computer.
- Lower costs. A network printer is more expensive than those typically used for local printing, but it also has significantly lower consumables costs and better quality printing. Therefore, the cost of printing is still minimized, because the initial cost of the printer is spread over all the computers that connect to that printer. For example, a single network printer could service 100 users or more.
- Users can search for printers easily. You can also publish network printers in AD DS, which allows users to search for printers in their domain.

- Centralized management via the Print Management Console
- Simplified troubleshooting
- Lower total cost of ownership
- Listing in AD DS

Enterprise Print Management

Windows Server 2012 provides a Print Management console that allows administrators to manage printing for the entire enterprise from a single interface. The Print Management console provides real-time information about the status of printers and print servers on the network and can take actions such as sending notifications or running scripts when printers need attention. With this console you can connect to and manage printers on print servers running Microsoft Windows 2000 and higher.

The Print Services Tools are not installed by default; you install the role using Server Manager or Windows PowerShell. Once installed, Print Services can detect print devices that exist on the same subnet as the print server, install the appropriate printer drivers, set up print queues and share the printers. Printers can then be deployed to users or computers through either existing or new group policies directly from the Print Management console.



Reference Links: For more information see the TechNet article Print Management Step-by-Step Guide at <http://go.microsoft.com/fwlink/?LinkID=331093>.

What Is Enhanced Point and Print?

Enhanced Point and Print is a new feature in Windows Server 2012 that makes it easier to install drivers for network printers. Enhanced Point and Print uses the new version 4 (v4) driver type that is introduced in Windows Server 2012 and Windows 8.

Understanding V3 Drivers and V4 Drivers

The Windows printer driver standard that is used in previous versions of Windows Server has existed in relatively the same form since the introduction of version 3 (v3) drivers in the Microsoft Windows 2000 operating systems. With v3 drivers, printer manufacturers create customized print drivers for each specific device that they produce, to ensure that Windows apps can use all of their printer's features. With the v3 model, printer infrastructure management requires administrators to maintain drivers for each print device in the environment, and separate 32-bit and 64-bit drivers for a single print device, to support both platforms.

- Enhanced Point and Print uses the v4 driver model to provide a simplified management structure for network printer drivers
- Benefits of Enhanced Point and Print :
 - Print servers do not need to store client print drivers
 - Driver files are isolated, preventing file naming conflicts
 - A single driver can support multiple devices
 - Driver packages are smaller and install faster
 - The print driver and the printer user interface can be deployed independently

Introducing the V4 Printer Driver

Windows Server 2012 and Windows 8 include support for v4 print drivers, which enables improved print device driver management and installation. Under the v4 model, print device manufacturers can create Print Class Drivers that support similar printing features and printing language that may be common to a large set of devices. Common printing languages may include Printer Control Language (PCL), .ps, or XML Paper Specification (XPS).

V4 drivers are typically delivered by using Windows Update or Windows Software Update Services. Unlike v3 drivers, v4 drivers are not delivered from a printer store that is hosted on the print server.

The v4 driver model provides the following benefits:

- Sharing a printer does not require provisioning drivers that match the client architecture.
- Driver files are isolated on a per-driver basis, preventing driver file naming conflicts.
- A single driver can support multiple devices.
- Driver packages are smaller and more streamlined than v3 drivers, resulting in faster driver installation times.
- The printer driver and the printer user interface can be deployed independently.

Using Enhanced Point and Print for Driver Installation

Under the v4 model, printer sharing and driver installation operates automatically under Enhanced Point and Print. When a network printer is installed on a client computer, the server and client work together to identify the print device. The driver then installs directly from the driver store on the client machine, or from Windows Update or Windows Software Update Services.

With Enhanced Point and Print, the print device drivers no longer need to be maintained on the print server. Driver installation for network print devices becomes faster because printer drivers no longer need to be transferred over the network from server to client.

If the driver store on the client machine does not contain a driver for the network printer that is being installed, and if an appropriate driver cannot be obtained from Windows Update or Windows Server Update Services, Windows uses a fallback mechanism to enable cross-platform printing using the print driver from the print server.

Security Options for Network Printing

When a printer is shared over a network, in many cases no security is required. The printer is considered open-access, meaning everyone is allowed to print on it. This is the default configuration for a printer that is shared on a Windows server.

The permissions that are available for shared printing include:

- **Print:** This permission allows users to print documents on the printer. By default, the Everyone group is assigned this permission.
- **Manage this printer:** This permission allows users to modify printer settings, including updating drivers. By default, this permission is given to Administrators, Server Operators, and Print Operators.
- **Manage documents:** This permission allows users to modify and delete print jobs in the queue. This permission is assigned to CREATOR OWNER, which means that the user who creates a print job manages that job. Administrators, Server Operators, and Print Operators also have this permission for all print jobs.

- The default security allows everyone to:
 - Print
 - Manage their own print jobs
- The available permissions are:
 - Print
 - Manage this printer
 - Manage documents

Demonstration: Creating Multiple Configurations for a Print Device

Creating multiple configurations for a print device enables you to assign print queues to specific users or groups so that they can print high priority jobs to a printer that is being used by other users. When a print job is sent to the high priority print queue, the print server will process the job before any jobs coming from the normal priority queue.

In this demonstration, you will see how to:

- Create a shared printer.
- Create a second shared printer using the same port.
- Increase printing priority for a high priority print queue.

Demonstration Steps

Create a shared printer

1. On LON-SVR1, open the Devices and Printers window.
2. Add a printer using the **LPT1** local port, and the **Brother Color Leg Type1 Class** driver.
3. Name the printer **AllUsers**.
4. Share the printer using the default settings.

Create a second shared printer using the same port

1. On LON-SVR1, open the Devices and Printers window.
2. Add a printer using the **LPT1** local port, and the **Brother Color Leg Type1 Class** driver.
3. Name the printer **Executives**.
4. Share the printer using the default settings.

Increase printing priority for a high priority print queue

1. Open the Executives Printer properties window.
2. Increase the Priority to **10**.

What Is Printer Pooling?

Printer pooling is a way to combine multiple physical printers into a single logical unit. To client computers, the printer pool appears to be a single printer. When jobs are submitted to the printer pool, any available printer in the printer pool can process them.

Printer pooling increases the scalability and availability of network printing by using a printer pool. If one printer in the pool is unavailable (for example, from a large print job, a paper jam, or being offline), all jobs are distributed to the remaining printers. If a printer pool does not have sufficient capacity, you can add another printer to the printer pool without performing any client configuration.

- Printer pooling combines multiple physical printers into a single logical unit
- A printer pool:
 - Increases availability and scalability
 - Requires that all printers use the same driver
 - Requires that all printers are in the same location

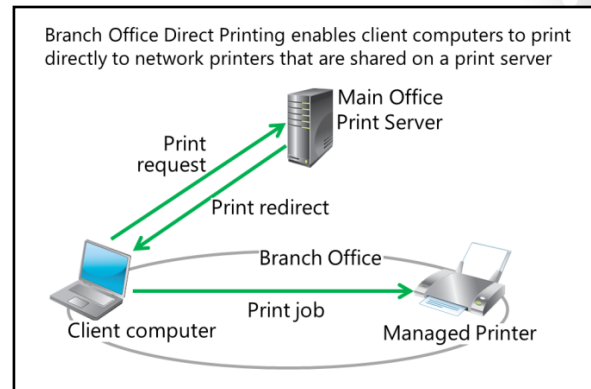
A printer pool is configured on a server by specifying multiple ports for a printer. Each port is the location of one physical printer. In most cases, the ports are an IP address on the network, instead of a local LPT or USB connection.

The requirements for a printer pool are as follows:

- Printers must use the same driver: Clients use a single printer driver for generating print jobs. All printers must accept print jobs in the same format. In many cases, this means that a single printer model is used.
- Printers should be in the same location: The printers in a printer pool should be located physically close together. When users retrieve their print jobs, they must check all printers in the printer pool to find their document. There is no way for users to know which printer has printed their document.

What Is Branch Office Direct Printing?

Branch Office Direct Printing reduces network costs for organizations that have centralized their Windows Server roles. When Branch Office Direct Printing is enabled, Windows clients obtain printer information from the print server, but send the print jobs directly to the printer. The print data no longer travels to the central server and then back to the branch office printer. This configuration reduces traffic between the client computer, the print server, and the branch office printer, and results in increased network efficiency.



Branch Office Direct Printing is transparent to the user. In addition, the user can print even if the print server is unavailable for some reason (for example if the wide area network (WAN) link to the data center is down). This is because the printer information is cached on the client computer in the branch office.

Configuring Branch Office Direct Printing

Branch Office Direct Printing is configured by an administrator using the Print Management console or a Windows PowerShell command-line interface.

To configure Branch Office Direct Printing from the Print Management console, use the following steps:

1. In Server Manager, open the Print Management console.
2. In the navigation pane, expand **Print Servers**, and then expand the print server that is hosting the network printer for which Branch Office Direct Printing will be enabled.
3. Click the **Printers** node, right-click the desired printer, and then click **Enable Branch Office Direct Printing**.

To configure Branch Office Direct Printing using Windows PowerShell, type the following cmdlet at a Windows PowerShell prompt:

```
Set-Printer -name "<Printer Name Here>" -ComputerName <Print Server Name Here>
-RenderingMode BranchOffice
```

Deploying Printers to Clients

Deploying printers to clients is a critical part of managing printing services on the network. A well-designed system for deploying printers is scalable and can be used to manage hundreds or thousands of computers.

The options for deploying printers are:

- Group Policy preferences. You can use Group Policy preferences to deploy shared printers to Windows XP, Windows Vista, Windows 7, and Windows 8 clients. The printer can be associated with either the user account or computer account, and can be targeted by group. For Windows XP computers, you must install the Group Policy Preference Client Extension.
- Group Policy Object (GPO) created by Print Management. The Print Management administrative tool can add printers to a GPO for distribution to client computers based on either a user account or a computer account. Windows XP computers must be configured to run PushPrinterConnections.exe.
- Manual installation. Each user can add printers manually by either browsing the network or using the Add Printer Wizard. It is important to note that network printers that are installed manually are available only to the user that installed them. If multiple users share a computer, they must each install the printer manually.

- You can deploy printers to clients by using:
 - Group Policy preferences
 - GPO created by Print Management
 - Manual installation



Lab: Implementing File and Print Services

Scenario

Your manager has recently asked you to configure file and print services for the branch office. This requires you to configure a new shared folder that will be used by multiple departments, configure shadow copies on the file servers, and configure a printer pool.

Many other users want to be able to work on their data files while they are connected across the Internet to other devices, for example, Windows RT-based tablets. You must ensure that these users are able to access their work-related data files from other locations when offline.

Objectives

After performing this lab you will be able to:

- Create and configure a file share.
- Configure shadow copies.
- Enable and configure Work Folders.
- Create and configure a printer pool.

Lab Setup

Estimated time: 60 minutes

Virtual machines	20410C-LON-CL1 20410C-LON-DC1 20410C-LON-SVR1
User name	Adatum\Administrator
Password	Pa\$\$w0rd

For this lab, you will use the available virtual machine environment. Before beginning 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 **20410C-LON-DC1** and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**.
Wait until the virtual machine starts.
4. Sign in using the following credentials:
 - User name: **Administrator**
 - Password: **Pa\$\$w0rd**
 - Domain: **Adatum**
5. Repeat steps 2 through 4 for **20410C-LON-SVR1**.
6. Repeat steps 2 and 3 for **20410C-LON-CL1**. Do not sign in to LON-CL1 until directed to do so.

Exercise 1: Creating and Configuring a File Share

Scenario

Your manager has asked you to create a new shared folder for use by all departments. There will be a single file share with separate folders for each department. To ensure that users only see files to which they have access, you need to enable access-based enumeration on the share.

There have been problems in other branch offices with conflicts when offline files are used for shared data structures. To avoid conflicts, you need to disable Offline Files for this share.

The main tasks for this exercise are as follows:

1. Create the folder structure for the new share.
2. Configure NTFS permissions on the folder structure.
3. Create the shared folder.
4. Test access to the shared folder.
5. Enable access-based enumeration.
6. Test access to the share.
7. Disable Offline Files for the share.

► Task 1: Create the folder structure for the new share

- On LON-SVR1, open File Explorer and create the following folders:
 - **E:\Data**
 - **E:\Data\Development**
 - **E:\Data\Marketing**
 - **E:\Data\Research**
 - **E:\Data\Sales**

► Task 2: Configure NTFS permissions on the folder structure

1. In File Explorer, block the NTFS permissions inheritance for **E:\Data**, and when prompted, convert inherited permissions into explicit permissions.
2. In File Explorer, remove permissions for **LON-SVR1\Users** on subdirectories in **E:\Data**.
3. In File Explorer, add the following NTFS permissions for the folder structure:

Folder	Permissions
E:\Data	No change
E:\Data\Development	Modify: Adatum\Development
E:\Data\Marketing	Modify: Adatum\Marketing
E:\Data\Research	Modify: Adatum\Research
E:\Data\Sales	Modify: Adatum\Sales

► Task 3: Create the shared folder

1. In File Explorer, share the **E:\Data** folder.
2. Assign the following permissions to the shared folder:
 - Change: **Adatum\Authenticated Users**

► Task 4: Test access to the shared folder

1. Sign in to LON-CL1 as **Adatum\Bernard** with the password **Pa\$\$w0rd**.



Note: Bernard is a member of the Development group.

2. Open File Explorer.
3. Navigate to **\\LON-SVR1\Data**.
4. Attempt to open the **Development**, **Marketing**, **Research**, and **Sales** folders.



Note: Bernard should have access to the Development folder. However, although Bernard can still see the other folders, he does not have access to their contents.

5. Sign out of LON-CL1.

► Task 5: Enable access-based enumeration

1. Switch to LON-SVR1.
2. Open Server Manager.
3. Click **File and Storage Management**.
4. Click **Shares**.
5. Open the **Properties** dialog box for the **Data** share, and on the **Settings** page, enable **Access-based enumeration**.

► Task 6: Test access to the share

1. Sign in to LON-CL1 as **Adatum\Bernard** with the password **Pa\$\$w0rd**.
2. Open File Explorer, and navigate to **\\LON-SVR1\Data**.



Note: Bernard can now view only the Development folder, the folder for which he has been assigned permissions.

3. Open the **Development** folder to confirm access.
4. Sign out of LON-CL1.

► Task 7: Disable Offline Files for the share

1. Switch to LON-SVR1.
2. Open File Explorer.

3. Navigate to drive **E**.
4. Open the **Properties** dialog box for the Data folder, and disable Offline file caching.

Results: After completing this exercise, you will have created a new shared folder for use by multiple departments.

Exercise 2: Configuring Shadow Copies

Scenario

A. Datum Corporation stores daily backups offsite for disaster recovery. Every morning the backup from the previous night is taken offsite. To recover a file from backup requires the backup tapes to be shipped back onsite. The overall time to recover a file from backup can be a day or more.

Your manager has asked you to ensure that shadow copies are enabled on the file server so you can restore recently modified or deleted files without using a backup tape. Because the data in this branch office changes frequently, you have been asked to configure a shadow copy to be created once per hour.

The main tasks for this exercise are as follows:

1. Configure shadow copies for the file share.
2. Create multiple shadow copies of a file.
3. Recover a deleted file from a shadow copy.

► Task 1: Configure shadow copies for the file share

1. On LON-SVR1.
2. Open File Explorer.
3. Navigate to drive **E**, right-click **Allfiles (E:)**, and then click **Configure Shadow Copies**.
4. Enable Shadow Copies for drive E.
5. Configure the settings to schedule hourly shadow copies for drive E.

► Task 2: Create multiple shadow copies of a file

1. On LON-SVR1, switch to **File Explorer**, and navigate to **E:\Data\Development**.
2. Create a new text file named **Report.txt**.
3. Switch back to the **Allfiles (E:) Properties** dialog box; it should still be opened on the **Shadow Copies** tab. Click **Create Now**.

► Task 3: Recover a deleted file from a shadow copy

1. On LON-SVR1, switch back to File Explorer.
2. Delete the **Report.txt** file.
3. Open the **Properties** dialog box for **E:\Data\Development**, and then click the **Previous Versions** tab.
4. Open the most recent version of the **Development** folder, and then copy the **Report.txt** file.

5. Paste the file back into the **Development** folder.
6. Close File Explorer and all open windows.

Results: After completing this exercise, you will have enabled shadow copies on the file server.

Exercise 3: Enabling and Configuring Work Folders

Scenario

You must now enable and configure Work Folders to support the requirements of your users. Domain users have their own Windows 8.1 and Windows RT 8.1 tablet devices and want access to their work data from anywhere. When they return to work, they want to be able to synchronize these data files. You will use Group Policy to force the Work Folders settings to users and test the settings.

The main tasks for this exercise are as follows:

1. Install the Work Folders role service.
2. Create a Sync Share on the File Server.
3. Automate settings for users via Group Policy.
4. Test synchronization.

► Task 1: Install the Work Folders role service

- On LON-SVR1, use Windows PowerShell to run the following command to install the Work Folders role service:

Add-WindowsFeature FS-SyncShareService

Note that the name of the feature is case-sensitive.

► Task 2: Create a Sync Share on the File Server

1. On LON-SVR1, use Windows PowerShell to run the following command to create the sync share named **Corp**:

New-SyncShare Corp -path C:\CorpData -User "Adatum\Domain Users"

2. Open Server Manager and view the Work Folders to ensure the sync share was created.

► Task 3: Automate settings for users via Group Policy

1. On LON-DC1, create a GPO named **Work Folders** and link it to the **Adatum.com** domain.
2. Edit the **Work Folders** GPO as follows:
 - Navigate to **User Configuration\Policies\Administrative Templates\Windows Components\Work Folders**.
 - Enable the Specify Work Folders settings policy and specify the Work Folders URL as **http://lon-svr1.Adatum.com**.
 - Select **Force automatic setup** to force automatic setup.
3. Close all open windows.

► Task 4: Test synchronization

1. Sign in to LON-CL1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Use File Explorer to navigate to **C:\Labfiles\Mod10** and double-click **WorkFolders.bat**

This adds a registry entry to allow unsecured connections to the work folders.

3. Sign out of LON-CL1.
4. Sign in to LON-CL1 as **Adatum\Administrator**.
5. In File Explorer, open Work Folders and create a new text document named **TestFile2**.
6. Switch to LON-SVR1 and use File Explorer to open **C:\CorpData\Administrator**.

Ensure the new text file you created exists.

Results: After completing this exercise, you will have installed the Work Folders role service, created a sync share, and created a Group Policy Object to deliver the settings to the users automatically. You will have also tested the settings.

Exercise 4: Creating and Configuring a Printer Pool

Scenario

Your manager has asked you to create a new shared printer for your branch office. However, instead of creating the shared printer on the local server in the branch office, he has asked you to create the shared printer in the head office and use Branch Office Direct Printing. This allows the printer to be managed in the head office, but prevents print jobs from traversing WAN links.

To ensure high availability of this printer, you need to format it as a pooled printer. Two physical print devices of the same model have been installed in the branch office for this purpose.

The main tasks for this exercise are as follows:

1. Install the Print and Document Services server role.
2. Install a printer.
3. Configure printer pooling.
4. Install a printer on a client computer.

► Task 1: Install the Print and Document Services server role

1. On LON-SVR1, open **Server Manager**.
2. Install the Print and Document Services role, and accept the default settings.

► Task 2: Install a printer

1. On LON-SVR1, use the Print Management console to install a printer with following parameters:
 - IP Address: **172.16.0.200**
 - Driver: **Microsoft XPS Class Driver**
 - Name: **Branch Office Printer**
2. Share the printer.
3. List the printer in AD DS.
4. Enable Branch Office Direct Printing.

► Task 3: Configure printer pooling

1. On LON-SVR1, in the Print Management console, create a new port with the following configuration:
 - Type: **Standard TCP/IP port**
 - IP Address: **172.16.0.201**
 - Connection: **Generic Network Card**
2. Open the **Branch Office Printer Properties** dialog box, and on the **Ports** tab, enable printer pooling.
3. Select port **172.16.0.201** as the second port.

► Task 4: Install a printer on a client computer

- On LON-CL1, add a printer, selecting the **Branch Office Printer on LON-SVR1** printer.

Results: After completing this exercise, you will have installed the Print and Document Services server role and installed a printer with printer pooling.

► Prepare for the next module

After 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 **20410C-LON-SVR1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20410C-LON-CL1** and **20410C-LON-DC1**.

Lab Review Questions

Question: How does implementing access-based enumeration benefit the users of the Data shared folder in this lab?

Question: Is there another way you could recover the file in the shadow copy exercise? What benefit do shadow copies provide in comparison?

Question: In Exercise 3, how could you configure Branch Office Direct Printing if you were in a remote location and did not have access to the Windows Server 2012 GUI for the print server?

Module Review and Takeaways

Review Questions

Question: How does inheritance affect explicitly assigned permissions on a file?

Question: Why should you not use shadow copies as a means for data backup?

Question: In which scenarios could Branch Office Direct Printing be beneficial?

Tools

Tool	Use	Where to find it
Effective Access Tool	Assessing combined permissions for a file, folder, or shared folder.	Under Advanced , on the Security tab of the Properties dialog box of a file, folder or shared folder.
Net share command-line tool	Configuring Windows Server 2012 networking components.	Command Prompt window.
Print Management console	Managing the print environment in Windows Server 2012.	The Tools menu in Server Manager.

Module 11

Implementing Group Policy

Contents:

Module Overview	11-1
Lesson 1: Overview of Group Policy	11-2
Lesson 2: Group Policy Processing	11-10
Lesson 3: Implementing a Central Store for Administrative Templates	11-16
Lab: Implementing Group Policy	11-20
Module Review and Takeaways	11-24

Module Overview

Maintaining a consistent computing environment across an organization is challenging. Administrators need a mechanism to configure and enforce user and computer settings and restrictions. Group Policy can provide that consistency by enabling administrators to centrally manage and apply configuration settings.

This module provides an overview of Group Policy and provides details about how to implement Group Policy.

Objectives

After completing this module, you will be able to:

- Create and manage Group Policy Objects.
- Describe Group Policy processing.
- Implement a Central Store for administrative templates.

Lesson 1

Overview of Group Policy

You can use Group Policy to control the settings of the computing environment. It is important to understand how Group Policy functions, so you can apply Group Policy correctly. This lesson provides an overview of Group Policy structure, and defines local and domain-based Group Policy Objects (GPOs). It also describes the types of settings available for users and groups.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the components of Group Policy.
- Describe multiple local GPOs.
- Describe storage options for domain GPOs.
- Describe GPO policies and preferences.
- Describe starter GPOs.
- Describe the process of delegating GPO management.
- Describe the process of creating and managing GPOs.

Components of Group Policy

Group Policy settings are configuration settings that allow administrators to enforce settings by modifying the computer-specific and user-specific registry settings on domain-based computers. You can group together Group Policy settings to make GPOs, which you can then apply to users or computers.

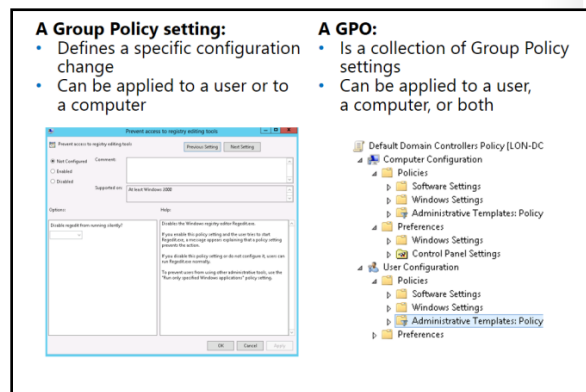
GPOs

A *Group Policy Object (GPO)* is an object that contains one or more policy settings that apply configuration setting for users, computers, or both. GPO templates are stored in SYSVOL, and GPO container objects are stored in Active Directory® Domain Services (AD DS). You can manage GPOs by using the Group Policy Management Console (GPMC). Within the GPMC, you can open and edit a GPO by using the Group Policy Management Editor window. GPOs are linked to Active Directory containers to apply settings to the objects in those containers.

Group Policy Settings

A *Group Policy setting* is the most granular component of Group Policy. It defines a specific configuration setting to apply to an object (a computer or a user, or both) within AD DS. Group Policy has thousands of configurable settings. These settings can affect nearly every area of the computing environment.

Not all settings can be applied to all older versions of Windows Server® and Windows® operating systems. Each new version introduces new settings and capabilities that only apply to that specific version. If a computer has a Group Policy setting applied that it cannot process, the setting is simply ignored.



Most policy settings have the following three states:

- Not Configured. The GPO does not modify the existing configuration of the particular setting for the user or computer.
- Enabled. The policy setting is applied.
- Disabled. The policy setting is specifically reversed.

By default, most settings are set to Not Configured.



Note: Some settings are multi-valued or have text string values. These are typically used to provide specific configuration details to apps or operating system components. For example, a setting might provide the URL of the home page used in Windows Internet Explorer® or the path to blocked apps.

The effect of a configuration change depends on the policy setting. For example, if you enable the Prohibit Access to Control Panel policy setting, users cannot open Control Panel. If you disable the policy setting, you ensure that users can open Control Panel. Notice the double negative in this policy setting: You disable a policy that prevents an action, thereby allowing the action.

Group Policy Settings Structure

There are two distinct areas of Group Policy settings:

- User settings. These are settings that modify the HKey Current User hive of the registry.
- Computer settings. These are settings that modify the HKEY Local Machine hive of the registry.

User and computer settings each have three areas of configuration, as described in the following table.

Section	Description
Software settings	Contain software settings that can be deployed to either the user or the computer. Software that is deployed to a user is specific to that user. Software that is deployed to the computer is available to all users of that computer.
Windows operating system settings	Contain script settings and security settings for both user and computer, and Internet Explorer maintenance settings for the user configuration.
Administrative templates	Contain hundreds of settings that modify the registry to control various aspects of the user and computer environment. New administrative templates might be created by Microsoft or other vendors. You can add these new templates to the GPMC. For example, there are Microsoft Office templates that you can add to the GPME. You can download these templates from the Microsoft website.

Group Policy Management Editor Window

The Group Policy Management Editor window displays the individual Group Policy settings that are available in a GPO. These are displayed in an organized hierarchy that begins with the division between computer settings and user settings, and then expands to show the Computer Configuration node and the User Configuration node. The Group Policy Management Editor window is where you configure all Group Policy settings and preferences.

Group Policy Preferences

In addition to the Group Policy sections shown in the previous table, a Preferences node is present under both the Computer Configuration and User Configuration nodes in the Group Policy Management Editor window. Preferences provide even more capabilities with which to configure the environment, and are discussed later in this module.

Local Group Policy

All systems that are running Microsoft Windows client or server operating systems also have available local GPOs. Local policy settings only apply to the local machine, but you can export and import them to other computers.

New in Windows Server 2012 R2

Windows Server 2012 R2 offers several new or updated Group Policy settings and features for computers that run Windows Server 2012 R2 or Windows 8.1. These settings and features include the following:

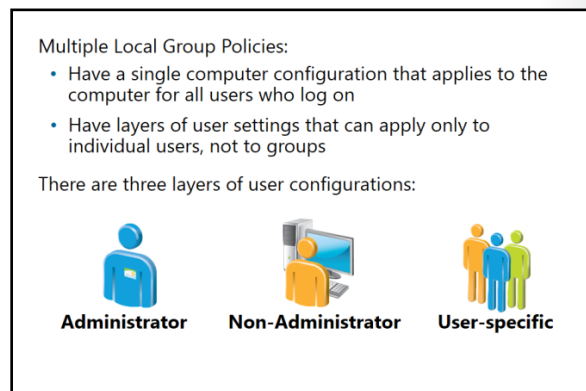
- Faster processing by using a Group Policy Caching setting. This new setting allows computers to rely on a local cache of a GPO when running in synchronous mode (which is the default mode for Group Policy processing).
- Increased support for IPv6. New IPv6 settings include the ability to push IPv6 printers and IPv6 VPN connections to computers. In addition, item-level targeting is available for IPv6.
- Extended logging for Group Policy operations. The Group Policy Operational event log contains more details of operational events, including the length of processing time and the amount of time for downloading policies, than previous versions. You can find this log here: Event Viewer\Applications and Services\Microsoft\Windows\GroupPolicy\Operational.
- There are many new settings for Windows 8.1 and Windows Server 2012 R2, including settings for managing the Start screen layout, configuring charms, and customizing background colors.


What Are Multiple Local GPOs?

In Windows operating systems prior to Windows Vista®, there was only one available user configuration in the local Group Policy. That configuration was applied to all users who logged on from that local computer. This is still true, but Windows Vista and newer Windows client operating systems, and Windows Server 2008 and newer Windows Server operating systems, have an added feature: multiple local GPOs. Since Windows 8 and Windows Server 2012, you also can have different user settings for different local users, but this is only available for users' configurations that are in Group Policy. In fact, there is only one set of computer configurations available that affects all users of the computer.

Since Windows 8 and Windows Server 2012, Computers that run Windows provide this ability with the following three layers of local GPOs:

- Local Group Policy (contains the computer configuration settings)
- Administrators and Non-Administrators Local Group Policy
- User-specific Local Group Policy




 **Note:** The exception to this feature is domain controllers. Due to the nature of their role, domain controllers cannot have local GPOs.

How the Layers Are Processed

The layers of local GPOs are processed in the following order:

1. Local Group Policy
2. Administrators and Non-Administrators Group Policy
3. User-specific Local Group Policy

With the exception of the Administrator or Non-Administrator categories, it is not possible to apply local GPOs to groups, but only to individual local user accounts. Domain users are subject to the local Group Policy, or to the Administrator or Non-Administrator settings, as appropriate.

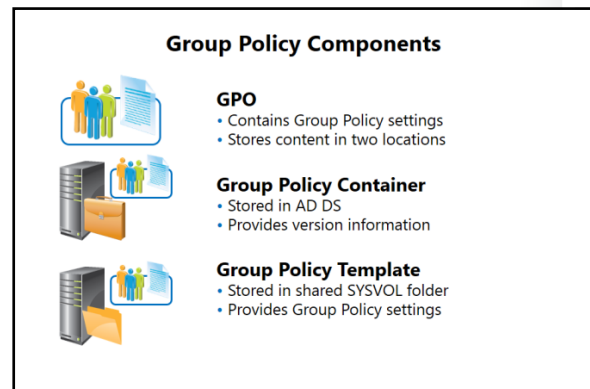
 **Note:** Domain administrators can disable processing local GPOs on clients that are running Windows client operating systems and Windows Server operating systems by enabling the Turn Off Local Group Policy Objects Processing policy setting.

Storage of Domain GPOs

Group Policy settings are presented as GPOs in the GPMC, but a GPO is actually two components: a Group Policy template, and a Group Policy container.

Group Policy Template

Group Policy templates are the actual collection of settings that you can change. The Group Policy template is a collection of files stored in the SYSVOL of each domain controller. SYSVOL is located in the %SystemRoot%\SYSVOL\Domain\Policies\GPOGUID path, where GPOGUID is the globally unique identifier (GUID) of the Group Policy container. When you create a GPO, a new Group Policy template is created in the SYSVOL folder, and a new Group Policy container is created in AD DS.



Group Policy Container

The Group Policy container is an Active Directory object that is stored in the Active Directory database. Each Group Policy container includes a GUID attribute that uniquely identifies the object within AD DS. The Group Policy container defines basic attributes of the GPO such as links and version numbers, but it does not contain any of the settings.

By default, when a Group Policy refresh occurs, the Group Policy client-side extensions (CSEs) only apply settings in a GPO if the GPO has been updated.

The Group Policy client can identify an updated GPO by its version number. Each GPO has a version number that is incremented each time a change is made. The version number is stored as an attribute of the Group Policy container, and in a text file, GPT.ini, in the Group Policy Template folder. The Group Policy Client knows the version number of each GPO that it has previously applied. If, during Group Policy refresh, the Group Policy client discovers that the version number of the Group Policy container has been changed, the CSEs are informed that the GPO is updated.

When editing a Group Policy Object, the version you are editing is the version on the domain controller that has the primary domain controller (PDC) emulator Flexible Single Master Operations (FSMO) role. It does not matter what computer you are using to perform the editing, the GPMC is focused on the PDC emulator by default. You can, however, change the focus of the GPMC to edit a version on a different domain controller.

What Are Group Policy Preferences?

Group Policy preferences are a Group Policy feature, which includes more than 20 Group Policy extensions that expand the range of configurable settings within a GPO. Configuring preferences helps reduce the need for logon scripts.

Characteristics of Preferences

Preferences have the following characteristics:

- Preferences exist for both computers and users.
- Unlike Group Policy settings, preferences are not enforced, and users can change the configurations that are established by preferences.
- Preferences can be managed through the Remote Server Administration Tool (RSAT).
- Preferences can be applied only once at startup or sign in, and can be refreshed at intervals.
- Unlike Group Policy settings, preferences are not removed when the GPO is no longer applied, but you can change this behavior.
- You can target preferences easily to certain users or computers by using a variety of ways, such as by the user's security group membership or by the operating system version.
- Preferences are not available for local GPOs.
- Unlike a Group Policy setting, the user interface of a Group Policy preference setting is not disabled.

Use Group Policy preferences to:

- Configure, deploy, and manage operating system and application settings that are not manageable by using Group Policy
 - Apply once at startup or sign in, optionally refresh at intervals
 - Target to users or computers
- Expand the range of configurable settings within a GPO

Group Policy preferences:

- Are not enforced
- Are not removed when the GPO no longer applies
- Do not disable the interface of the setting; users can change the setting
- Cannot be used in local group policies

Common Uses for Group Policy Preferences

Although you can configure many settings through Group Policy preferences, the following are some of the more common uses:

- Map network drives for users
- Configure desktop shortcuts for users or computers
- Set environment variables
- Map printers
- Set power options
- Configure Start menus
- Configure data sources
- Configure Internet options
- Schedule tasks

What Are Starter GPOs?


Starter GPOs are templates that assist in the creation of GPOs. When creating new GPOs, you can choose to use a starter GPO as the source. This makes it easier and faster to create multiple GPOs with the same baseline configuration.

Available Settings

Starter GPOs contain settings from only the Administrative Templates node of either the User Configuration section or the Computer Configuration section. The Software Settings and Windows Settings nodes of Group Policy are not available, because these nodes involve interaction of services and are more complex and domain-dependent.

A starter GPO:

- Has preconfigured administrative template settings upon which new GPOs can be based
- Can be exported to .cab files
- Can be imported into other areas of the enterprise



Exporting Starter GPOs

You can export starter GPOs to a Cabinet file (.cab) and then load that .cab file into another environment that is completely independent of the source domain or forest. By exporting a starter GPO, you can send the .cab file to other administrators, who can then use it in other areas. For example, you might create a GPO that defines Internet Explorer security settings. If you want all sites and domains to employ the same settings, then you could export the starter GPO to a .cab file, and then distribute it.

When to Use Starter GPOs

The most common situation in which you would use a starter GPO is when you want a group of settings for a type of computer role. For example, you might want all corporate laptops to have the same desktop restrictions, or all file servers to have the same baseline Group Policy settings, but enable variations for different departments.

Included Starter GPOs

The GPMC includes a link to create a Starter GPO folder, which contains a number of predefined starter GPOs. These policies provide preconfigured, security-oriented settings for Enterprise Clients (EC), in addition to Specialized Security–Limited Functionality (SSLF) clients for both user and computer settings on Windows Vista and Windows XP with Service Pack 2 (SP2) operating systems. You can use these policies as starting points when you design security policies.

Delegating Management of GPOs

Administrators can delegate some of the Group Policy administrative tasks to other users. These users do not have to be domain administrators—they can be users that are granted certain rights to GPOs.

For example, a user who manages a particular Organizational Unit (OU) could be tasked with performing reporting and analysis duties, while the help desk group is allowed to edit GPOs for that OU. A third group made up of developers might be put in charge of creating Windows Management Instrumentation (WMI) filters.

Delegation of GPO-related tasks allows the administrative workload to be distributed across the enterprise

The following Group Policy tasks can be independently delegated:

- Creating GPOs, including Starter GPOs
- Editing GPOs
- Managing Group Policy links for a site, domain, or OU
- Performing Group Policy Modeling analysis in a domain or OU
- Reading Group Policy Results data in a domain or OU
- Creating WMI filters on a domain

The following Group Policy administrative tasks can be delegated independently:

- Creating GPOs, including creating Starter GPOs
- Editing GPOs
- Managing Group Policy links for a site, domain, or OU
- Performing Group Policy modeling analysis
- Reading Group Policy results data
- Creating WMI filters

Members of the Group Policy Creator Owners group can create new GPOs and edit or delete GPOs that they have created.

Group Policy Default Permissions

By default, the following users and groups have full access to manage Group Policy:

- Domain Admins
- Enterprise Admins
- Creator Owner
- Local System

The Authenticated User group has Read and Apply Group Policy permissions only.

Permissions for Creating GPOs

By default, only Domain Admins, Enterprise Admins, and Group Policy Creator Owners can create new GPOs. You can use two methods to grant a group or user this right:

- Add the user to the Group Policy Creator Owners group
- Explicitly grant the group or user permission to create GPOs by using the GPMC

Permissions for Editing GPOs

To edit a GPO, the user must have both Read and Write access to the GPO. You can grant this permission by using the GPMC.

Managing GPO Links

The ability to link GPOs to a container is a permission that is specific to that container. In the GPMC, you can manage this permission by using the Delegation tab on the container. You can also delegate it through the Delegation of Control Wizard in Active Directory Users and Computers.

Group Policy Modeling and Group Policy Results

You can delegate the ability to use the reporting tools either through the GPMC or through the Delegation of Control Wizard in Active Directory Users and Computers.

Creating WMI Filters

You can delegate the ability to create and manage WMI filters either through the GPMC or through the Delegation of Control Wizard in Active Directory Users and Computers.

Demonstration: Creating and Managing GPOs

In this demonstration, you will see how to:

- Create a GPO by using the GPMC.
- Edit a GPO in the Group Policy Management Editor window.
- Use Windows PowerShell to create a GPO.

Demonstration Steps

Create a GPO by using the GPMC

- Sign in to **LON-DC1** as **Administrator** with a password of **Pa\$\$w0rd**, and create a policy named Prohibit Windows Messenger.

Edit a GPO in the Group Policy Management Editor window

1. Edit the policy to prohibit the use of Windows Messenger.
2. Link the Prohibit Windows Messenger GPO to the domain.

Use Windows PowerShell to create a GPO named Desktop Lockdown

- In Windows PowerShell, import the **grouppolicy** module, and then use the following **New-GPO** cmdlet:

```
New-GPO -Name "Desktop Lockdown"
```

Lesson 2

Group Policy Processing

Understanding how Group Policy is applied is the key to being able to develop a Group Policy strategy. This lesson shows you how Group Policy is associated with Active Directory objects, how it is processed, and how to control the application of Group Policy. After creating the GPOs and configuring the settings you want to apply, they must be linked to containers. GPOs are applied in a specific order. This order might determine what settings are applied to objects. There are two default policies that are automatically created. You can use these policies to deliver password and security settings for the domain and for domain controllers. The application of policies can also be controlled through security filtering.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe a GPO link.
- Explain how to apply GPOs to containers and objects.
- Describe the Group Policy processing order.
- Describe the default GPOs.
- Describe GPO security filtering.

GPO Links

Once you have created a GPO and defined all the settings that you want it to deliver, the next step is to link the policy to an Active Directory container. A GPO link is the logical connection of the policy to a container. You can link a single GPO to multiple containers by using the GPMC. You can link GPOs to the following types of containers:

- Sites
- Domains
- Organizational Unit (OUs)

When linking GPOs, remember that:

- To deliver settings to an object, a GPO must be linked to a container
- Disabling a link removes the settings from the container
- Deleting a link does not delete the GPO

GPOs can be linked to:

- Sites
- Domains
- OUs

GPOs cannot be linked to:

- Users
- Groups
- Computers
- System containers

Once a GPO is linked to a container, by default the policy is applied to all the objects in the container, and subsequently all the child containers under that parent object. This is because the default permissions of the GPO are such that Authenticated Users have Read and Apply Group Policy permission. You can modify this behavior by managing permissions on the GPO.

You can disable links to containers, which removes the configuration settings. You can also delete links. Deleting links does not delete the actual GPO, only the logical connection to the container.

GPOs cannot be linked directly to users, groups, or computers. In addition, GPOs cannot be linked to the system containers in AD DS, including Builtin, Computers, Users, or Managed Service Accounts. The AD DS system containers receive Group Policy settings from GPOs that are linked to the domain level only.

Applying GPOs

Computer configuration settings are applied at startup, and then are refreshed at regular intervals. Any startup scripts are run at computer startup. The default interval is every 90 minutes, but this is configurable. The exceptions to this default interval are domain controllers, which have their settings refreshed every five minutes.

User settings are applied at logon and are refreshed at regular, configurable intervals. The default for this is also 90 minutes. Any logon scripts are run at sign in.

When you apply GPOs, remember that:

- Computer settings apply at startup
- User settings apply at sign in
- Policies refresh at regular, configurable intervals
- Security settings refresh at least every 16 hours
- Policies refresh manually by using:
 - The Gpupdate command
 - The Windows PowerShell cmdlet Invoke-Gpupdate
- Since Windows Server 2012 and Windows 8, a new Remote Policy Refresh feature allows you to remotely refresh policies



Note: A number of user settings require two sign ins before the user sees the effect of the GPO. This is because multiple users signing in to the same computer use cached credentials to speed up sign ins. This means that, although the policy settings are being delivered to the computer, the user is already signed in and thus the settings do not take effect until the next sign in. The Folder Redirection setting is an example of this.

You can change the refresh interval by configuring a Group Policy setting. For computer settings, the refresh interval setting is found in the Computer Configuration\Policies\Administrative Templates\System\Group Policy node. For user settings, the refresh interval is found at the corresponding settings under User Configuration. An exception to the refresh interval is the security settings. The security settings section of the Group Policy is refreshed at least every 16 hours, regardless of the interval that you set for the refresh interval.

You can also refresh Group Policy manually. The command-line utility **Gpupdate** refreshes and delivers any new Group Policy configurations. The **Gpupdate /force** command refreshes all the Group Policy settings. There is also a new Windows PowerShell **Invoke-Gpupdate** cmdlet, which performs the same function.

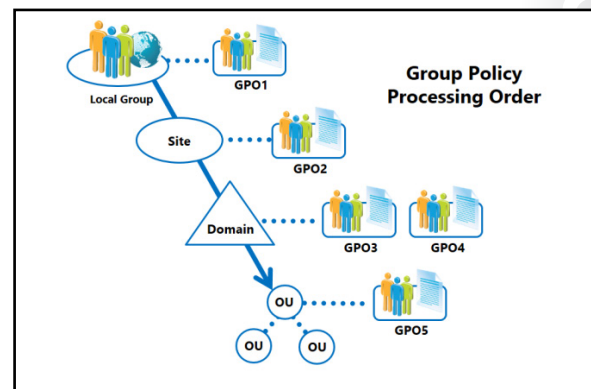
A new feature in Windows Server 2012 and in Windows 8 is Remote Policy Refresh. This feature allows administrators to use the GPMC to target an OU and force Group Policy refresh on all of its computers and their currently signed-in users. To force a Group Policy refresh, right-click any OU, and then click **Group Policy Update**. The update occurs within 10 minutes.

Group Policy Processing Order

GPOs are not applied simultaneously; rather, they are applied in a logical order. GPOs that are applied later in the process of applying GPOs overwrite any conflicting policy settings that were applied earlier.

GPOs are applied in the following order:

1. Local GPOs. Local GPOs are processed first. Computers running Windows operating systems already have a configured local Group Policy.



2. Site GPOs. Policies that are linked to sites are processed next.
3. Domain GPOs. Policies that are linked to the domain are processed next. There are often multiple policies at the domain level. These policies are processed in order of preference.
4. OU GPOs. Policies linked to OUs are processed next. These policies contain settings that are unique to the objects in that OU. For example, the Sales users might have special required settings. You can link a policy to the Sales OU to deliver those settings.
5. Child OU policies. Any policies that are linked to child OUs are processed last.

Objects in the containers receive the cumulative effect of all policies in their processing order. In the case of a conflict between settings, the last policy applied takes effect. For example, a domain-level policy might restrict access to registry editing tools, but you could configure an OU-level policy and link it to the IT OU to reverse that policy. Because the OU-level policy is applied later in the process, access to registry tools would be available.



Note: Other methods such as Enforcement and Inheritance Blocking can change the effect of policies on containers.

If multiple policies are applied at the same level, the administrator can assign a preference value to control the order of processing. The default preference order is the order in which the policies were linked.

The administrator can also disable the user or computer configuration of a particular GPO. If one section of a policy is known to be empty, it should be disabled to speed up policy processing. For example, if there is a policy that only delivers user desktop configuration, the administrator could disable the computer side of the policy.

What Are the Default GPOs?

During the installation of the AD DS role, two default GPOs are created: Default Domain Policy, and Default Domain Controllers Policy.

Default Domain Policy

The Default Domain Policy is linked to the domain and affects all security principles in the domain. It contains the default password policy settings, the account lockout settings, and the Kerberos protocol. As a best practice, this policy should not have other settings configured. If you need to configure other settings to apply to the entire domain, then you should create new policies to deliver the settings, and then link those policies to the domain.

There are two default GPOs:

- Default Domain Policy
 - Used to define the account policies for the domain:
 - Password
 - Account lockout
 - Kerberos protocol
- Default Domain Controllers Policy
 - Used to define auditing policies
 - Defines user rights on domain controllers



Note: Currently, fine-grained password policies are the typical enterprise method, although they are beyond the scope of this module.

Default Domain Controllers Policy

The Default Domain Controllers Policy is linked to the Domain Controllers OU, and should only affect domain controllers. This policy is designed to provide auditing settings and user rights, and should not be used for other purposes.

GPO Security Filtering

By nature, a GPO applies to all the security principles in the container, and all child containers below the parent. However, you might want to change that behavior and have certain GPOs apply only to particular security principles. For example, you might want to exempt certain users in an OU from a restrictive desktop policy. You can accomplish this through security filtering.

Each GPO has an Access Control List (ACL) that defines permissions to that GPO. The default permission is for Authenticated Users to have the Read and Apply Group Policy permissions applied.

Apply Group Policy permissions

- GPO has an ACL (Delegation tab, click Advanced)
- Default: Authenticated Users have Allow Apply Group Policy

Scope only to users in selected global or universal groups

- Remove Authenticated Users
- Add appropriate global or universal groups (GPOs do not scope to domain local groups)

Scope to users except for those in selected groups

- On the Delegation tab, click Advanced
- Add appropriate global groups
- Deny the Apply Group Policy permission

By adjusting the permissions in the ACL, you can control which security principles receive permission to have the GPO settings applied. There are two approaches you might take to do this:

- Deny access to the Group Policy
- Limit permissions to Group Policy



Note: The Authenticated Users group includes all user and computer accounts that have been authenticated to AD DS.

Deny Access to Group Policy

If most security principles in the container should receive the policy settings but some should not, then you can exempt particular security principles by denying them access to the Group Policy. For example, you might have a Group Policy that all the users in the Sales OU should receive except the Sales Managers group. You can exempt that group (or user) by adding that group (or user) to the ACL of the GPO, and then setting the permission to Deny.

Limit Permissions to Group Policy

Alternatively, if you have created a GPO that should only be applied to a few security principles in a container, you can remove the Authenticated Users group from the ACL, add the security principles that should receive the GPO settings, and then grant the security principles the Read and Apply Group Policy permissions. For example, you might have a GPO with computer configuration settings that should only apply to laptop computers. You could remove the Authenticated Users group from the ACL, add the computer accounts of the laptops, and then grant the security principles the Read and Apply Group Policy permission.

The ACL of a GPO is accessed in the GPMC by selecting the GPO in the Group Policy Object folder, and then clicking the **Delegation>Advanced** tab.



Note: As a best practice, you should never deny access to the Authenticated User group. If you do, then security principles would never receive the GPO settings.

Discussion: Identifying Group Policy Application

For this discussion, review the AD DS structure in the graphic, read the scenario, and then answer the questions on the slide.

Scenario

The following illustration represents a portion of the A. Datum Corporation's AD DS structure, which contains the Sales OU with its child OUs and the Servers OU.

Review the scenario and the AD DS structure graphic in the handbook to answer the following questions:

- What power options will the servers in the Servers OU receive?
- What power options will the laptops in the Sales Laptops OU receive?
- What power options will all other computers in the domain receive?
- Will users in the Sales Users OU who have created local policies to grant access to Control Panel be able to access Control Panel?
- If you needed to grant access to Control Panel to some users, how would you do it?
- Can GPO2 be applied to other department OUs?



20 minutes

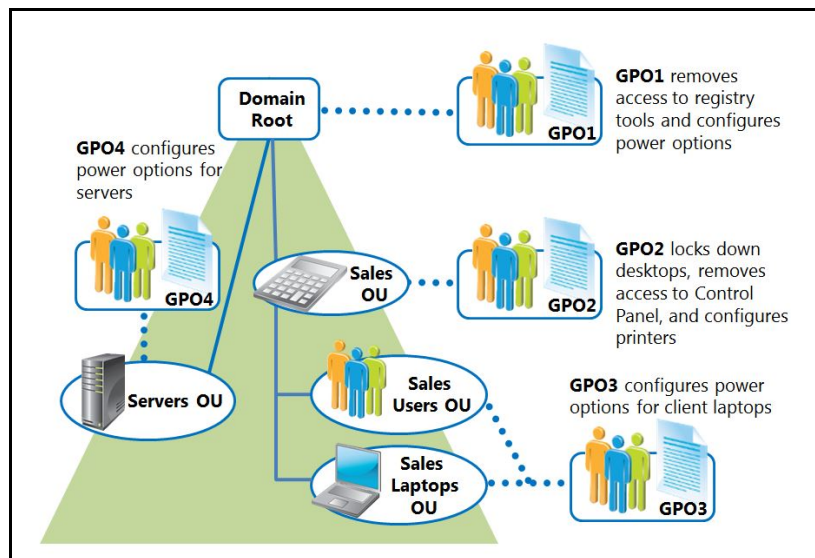


FIGURE 11.1: AD DS STRUCTURE

- GPO1 is linked to the Adatum domain container. The GPO configures power options that turn off the monitors and disks after 30 minutes of inactivity, and restricts access to registry editing tools.
- GPO2 has settings to lock down the desktops of the Sales Users OU, and configure printers for Sales Users.
- GPO3 configures power options for laptops in the Sales Laptops OU.
- GPO4 configures a different set of power options to ensure that the servers never go into power save mode.

Some users in the Sales OU have administrative rights on their computers, and have created local policies to specifically grant access to Control Panel.

Discussion Questions

Based on this scenario, answer the following questions:

Question: What power options will the servers in the Servers OU receive?

Question: What power options will the laptops in the Sales Laptops OU receive?

Question: What power options will all other computers in the domain receive?

Question: Will users in the Sales Users OU who have created local policies to grant access to Control Panel be able to access Control Panel?

Question: If you needed to grant access to Control Panel to some users, how would you do it?

Question: Can GPO2 be applied to other department OUs?

Demonstration: Using Group Policy Diagnostic Tools

In this demonstration, you will see how to:

- Use **Gpupdate** to refresh Group Policy.
- Use the **Gpresult** cmdlet to output the results to an HTML file.
- Use the Group Policy Modeling Wizard to test the policy.

Demonstration Steps

Use Gpupdate to refresh Group Policy

- On LON-DC1, use **Gpupdate** to refresh the GPOs.

Use the Gpresult cmdlet to output the results to an HTML file

1. Use **Gpresult /H** to create an HTML file that displays the current GPO settings.
2. Open the HTML report and review the results.

Use the Group Policy Modeling Wizard to test the policy

- Use the Group Policy Modeling Wizard to simulate a policy application for users in the Managers OU who sign in to any computer.

Lesson 3

Implementing a Central Store for Administrative Templates

Larger organizations might have many GPOs with multiple administrators managing them. When an administrator edits a GPO, the template files are pulled from the local workstation. The Central Store provides a single folder in SYSVOL that contains all of the templates required to create and edit GPOs.

This lesson discusses the files that make up the templates, and covers how to create a Central Store location to provide consistency in the templates that administrators use.

Lesson Objectives

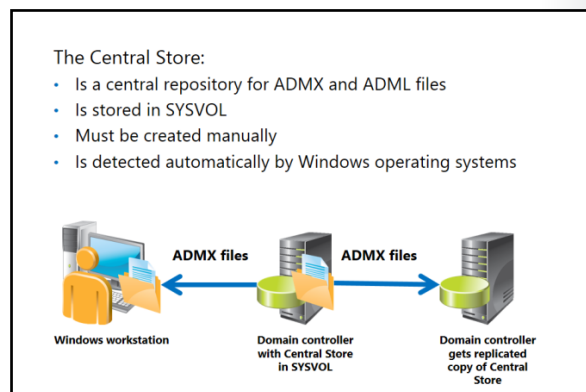
After completing this lesson, you will be able to:

- Describe the Central Store.
- Describe administrative templates.
- Describe how administrative templates work.
- Describe managed and unmanaged policy settings.

What Is the Central Store?

If your organization has multiple administration workstations, there could be potential issues when editing GPOs. If you do not have a Central Store that contains the template files, then the workstation from which you are editing will use the .admx (ADMX) and .adml (ADML) files that are stored in the local PolicyDefinitions folder. If different administration workstations have different operating systems or are at different service pack levels, there might be differences in the ADMX and ADML files. For example, the ADMX and ADML files that are stored on a workstation running Windows 7 with no service pack installed might not be the same as the files that are stored on a domain controller running Windows Server 2012. This could lead to administrators not seeing the same settings in a GPO.

The Central Store addresses this issue. The Central Store provides a single point from which administration workstations can download the same ADMX and ADML files when editing a GPO. The Central Store is detected automatically by Windows operating systems (Windows Vista or newer or Windows Server 2008 or newer). Because of this automatic behavior, the local workstation that the administrator uses to perform administration always checks to see if a Central Store exists before loading the local ADMX and ADML files in the Group Policy Management Editor window. When the local workstation detects a Central Store, it then downloads the template files from there. In this way, there is a consistent administration experience among multiple workstations.



Creating and Provisioning the Central Store

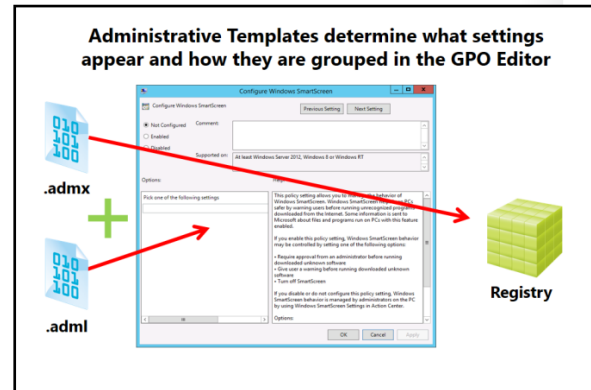
You must create and provision the Central Store manually. First you must create a folder on a domain controller, name the folder PolicyDefinitions, and store the folder at C:\Windows\SYSVOL\sysvol\{Domain Name}\Policies\. This folder is now your Central Store. You must then copy all the contents of the C:\Windows\PolicyDefinitions folder to the Central Store. The ADML files in this folder are also in a language-specific folder (such as en-US).

What Are Administrative Templates?

An administrative template is made up of two XML files types:

- ADMX files that specify the registry setting to change. ADMX files are language-neutral.
- ADML files that generate the user interface to configure the Administrative Templates policy settings in the Group Policy Management Editor window. ADML files are language-specific.

ADMX and ADML files are stored in the %SystemRoot%\PolicyDefinitions folder or in the Central Store. You can also create your own custom administrative templates in XML format. Administrative templates that control Microsoft Office products (such as Office Word, Office Excel and Office PowerPoint) are also available from the Microsoft website.



Administrative templates have the following characteristics:

- They are organized into subfolders that house configuration options for specific areas of the environment, such as network, system, and Windows components.
- The settings in the Computer section edit the HKEY_LOCAL_MACHINE registry hive, and settings in the User section edit the HKEY_CURRENT_USER registry hive.
- Some settings exist for both User and Computer. For example, there is a setting to prevent Windows Messenger from running in both the User and the Computer templates. In case of conflicting settings, the Computer setting prevails.
- Some settings are available only to certain versions of Windows operating systems. Double-clicking the settings displays the supported versions for that setting. Any setting that cannot be processed by an older Windows operating system is simply ignored by that system.

ADM Files

Prior to Windows Vista, administrative templates had an .adm file extension (ADM). ADM files were language-specific, and were difficult to customize. ADM files are stored in SYSVOL as part of the Group Policy template. If an ADM file is used in multiple GPOs, then the file is stored multiple times. This increases the size of SYSVOL, and therefore increases the size of Active Directory replication traffic.

How Administrative Templates Work

Administrative Templates have settings for almost every aspect of the computing environment. Each setting in the template corresponds to a registry setting that controls an aspect of the computing environment. For example, when you enable the setting that prevents access to Control Panel, the value in the registry key that controls that aspect also changes.

The Administrative Templates node is organized as shown in the following table.

Section	Nodes
Computer settings	<ul style="list-style-type: none"> Control Panel Network Printers System Windows Components All Settings
User settings	<ul style="list-style-type: none"> Control Panel Desktop Network Shared Folders Start Menu and Taskbar System Windows Components All Settings

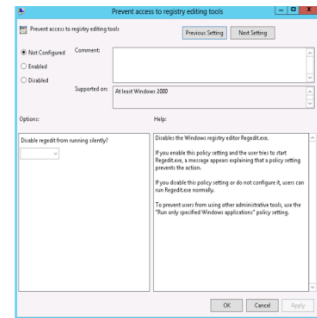
Most of the nodes contain multiple subfolders to further organize settings into logical groupings. Even with this organization, finding the setting that you need might be a daunting task.

To help you locate settings in the All Settings folder you can filter the entire list of settings in either the computer or the user section. The following filter options are available:

- Managed or unmanaged
- Configured or not configured
- Commented
- By keyword
- By platform

You can also combine multiple criteria. For example, you could filter to find all the configured settings that apply to Internet Explorer 10 by using the keyword **ActiveX**.

- Changing policy settings in the Administrative Templates node also changes the registry
- Changing the **Prevent access to registry editing tools** setting changes the value of HKLM\Software\Classes\Regedit



Managed and Unmanaged Policy Settings

There are two types of policy settings: managed, and unmanaged. All policy settings in a GPO's Administrative Templates are managed policies. The Group Policy service controls the managed policy settings and removes a policy setting when it is no longer within scope of the user or computer. The Group Policy service does not control unmanaged policy settings. These policy settings are persistent. The Group Policy service does not remove unmanaged policy settings.

Managed policy settings:

- UI is locked; user cannot make a change to the setting
- Changes are made in one of four reserved registry keys
- Change and UI locks are released when the user/computer falls out of scope

Unmanaged policy settings:

- UI is not locked
- Changes made are persistent: tattoos the registry
- Only managed settings are shown by default
- Set Filter options to view unmanaged settings

Managed Policy Settings

A managed policy setting has the following characteristics:

- The UI is locked so that a user cannot change the setting. Managed policy settings result in the appropriate UI being disabled. For example, if you configure the desktop wallpaper through a Group Policy setting, then those settings are grayed out in the user's local UI.
- Changes are made in the restricted areas of the registry to which only administrators have access. These reserved registry keys are:
 - HKLM\Software\Policies (computer settings)
 - HKCU\Software\Policies (user settings)
 - HKLM\Software\Microsoft\Windows\Current Version\Policies (computer settings)
 - HKCU\Software\Microsoft\Windows\Current Version\Policies (user settings)
- Changes made by a Group Policy setting and the UI lockout are released if the user or computer falls out of scope of the GPO. For example, if you delete a GPO, managed policy settings that had been applied to a user are released. Typically, the setting then resets to its previous state. Also, the UI interface for the setting is enabled.

Unmanaged Policy Settings

In contrast, an unmanaged policy setting makes a change that is persistent in the registry. If the GPO no longer applies, the setting remains. This is often called *tattooing* the registry—in other words, making a permanent change. To reverse the effect of the policy setting, you must deploy a change that reverts the configuration to the desired state. Additionally, an unmanaged policy setting does not lock the UI for that setting.

By default, the Group Policy Management Editor hides unmanaged policy settings to discourage administrators from implementing a configuration that is difficult to revert. Many of the settings that are available in Group Policy preferences are unmanaged settings.

Lab: Implementing Group Policy

Scenario

A. Datum Corporation is a global engineering and manufacturing company with a head office based in London, England. An IT office and a data center are located in London to support the London location and other locations. A. Datum has recently deployed a Windows Server 2012 infrastructure with Windows 8 clients.

In your role as a member of the server support team, you help to deploy and configure new servers and services into the existing infrastructure based on the instructions given to you by your IT manager.

Your manager has asked you to create a Central Store for ADMX files to ensure that everyone can edit GPOs that have been created with customized ADMX files. You also need to create a starter GPO that includes Internet Explorer settings, and then configure a GPO that applies GPO settings for the Marketing department and the IT department.

Objectives

After completing this lab, you will be able to:

- Configure a Central Store.
- Create GPOs.

Lab Setup

Estimated Time: 40 minutes

Virtual machines	20410C-LON-DC1 20410C-LON-CL1
User name	Adatum\Administrator
Password	Pa\$\$w0rd

Lab Setup Instructions

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 **20410C-LON-DC1**. In the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in using the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
5. Repeat steps 2 and 3 for 20410C-LON-CL1. Do not sign in until directed to do so.

Exercise 1: Configuring a Central Store

Scenario

A. Datum recently implemented a customized ADMX template to configure an app. A colleague obtained the ADMX files from the vendor before creating the GPO with the configurations settings. The settings were applied to the app as expected.

After implementation, you noticed that you are unable to modify the app settings in the GPO from any location other than the workstation that was originally used by your colleague. To resolve this issue, your manager has asked you to create a Central Store for administrative templates. After you create the Central Store, your colleague will copy the vendor ADMX template from the workstation into the Central Store.

The main tasks for this exercise are as follows:

1. View the location of administrative templates in a Group Policy Object (GPO).
2. Create a Central Store.
3. Copy administrative templates to the Central Store.
4. Verify the administrative template location in GPMC.

► Task 1: View the location of administrative templates in a Group Policy Object (GPO)

1. Sign in to LON-DC1 as **Administrator** with a password of **Pa\$\$w0rd**.
2. Start the **Group Policy Management Console** (GPMC).
3. In the Group Policy Object folder, open the **Default Domain Policy** and view the location of the administrative templates.

► Task 2: Create a Central Store

1. Open **File Explorer** and browse to **C:\Windows\SYSVOL\sysvol\Adatum.com\Policies**.
2. Create a folder to use for the Central Store and name it **PolicyDefinitions**.

► Task 3: Copy administrative templates to the Central Store

- Copy the contents of the default PolicyDefinitions folder located at **C:\Windows\PolicyDefinitions** to the new PolicyDefinitions folder located at **C:\Windows\SYSVOL\sysvol\Adatum.com\Policies**.

► Task 4: Verify the administrative template location in GPMC

1. Verify that the Group Policy Management Editor is using the ADMX files from the Central Store **PolicyDefinitions** folder by viewing the location information text of the Administrative templates folder.
2. Close the **Group Policy Management Editor** window.

Results: After completing this exercise, you should have configured a Central Store.

Exercise 2: Creating GPOs

Scenario

After a recent meeting of the IT Policy committee, management has decided that A. Datum will use Group Policy to restrict user access to the General page of Internet Explorer.

Your manager has asked you to create a starter GPO that can be used for all departments with default restriction settings for Internet Explorer. You then need to create the GPOs that will deliver the settings for members of all departments except for the IT department.

The main tasks for this exercise are as follows:

1. Create a Windows Internet Explorer Restriction default starter GPO.
2. Configure the Internet Explorer Restriction starter GPO.
3. Create an Internet Explorer Restrictions GPO from the Internet Explorer Restrictions starter GPO.
4. Test the GPO for Domain Users.
5. Use security filtering to exempt the IT Department from the Internet Explorer Restrictions policy.
6. Test the GPO app for IT department users.
7. Test Application of the GPO for other domain users.

► Task 1: Create a Windows Internet Explorer Restriction default starter GPO

1. Open the **GPMC** and create a starter GPO named **Internet Explorer Restrictions**.
2. Type a comment that states **This GPO disables the General page in Internet Options**.

► Task 2: Configure the Internet Explorer Restriction starter GPO

1. Configure the starter GPO to disable the **General** page of Internet Options and name it **Internet Explorer Restrictions**.



Hint: Select **All Settings** in Administrative Templates and filter for an exact match by using the keywords **General page**.

2. Close the Group Policy Management Editor window.

► Task 3: Create an Internet Explorer Restrictions GPO from the Internet Explorer Restrictions starter GPO

- Create a new GPO named **IE Restrictions** that is based on the Internet Explorer Restrictions starter GPO, and then link it to the **Adatum.com** domain.

► Task 4: Test the GPO for Domain Users

1. Sign in to **LON-CL1** as **Adatum\Brad**, with a password of **Pa\$\$w0rd**.
2. Open **Control Panel**.
3. Attempt to change your home page.
4. Open **Internet Options** to verify that the **General** tab has been restricted.
5. Sign out from **LON-CL1**.

► **Task 5: Use security filtering to exempt the IT Department from the Internet Explorer Restrictions policy**

1. On LON-DC1, open the **GPMC**.
2. Configure security filtering on the **Internet Explorer Restrictions** policy to deny access to the IT department.

► **Task 6: Test the GPO app for IT department users**

1. Sign in to **LON-CL1** as **Brad**, with a password of **Pa\$\$w0rd**.
2. Open **Control Panel**.
3. Attempt to change your home page. Verify that the **Internet Properties** dialog box opens to the **General** tab, and all settings are available.
4. Sign out from **LON-CL1**.

► **Task 7: Test Application of the GPO for other domain users**

1. Sign in to **LON-CL1** as **Boris**, with a password of **Pa\$\$w0rd**.
2. Open **Control Panel**.
3. Attempt to change your home page.
4. Open **Internet Options** to verify that the **General** tab has been restricted.
5. Sign out from **LON-CL1**.

Results: After completing this lab, you should have created a GPO.

► **Prepare for the next module**

After you finish the lab, revert the virtual machines back 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 **20410C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for 20410C-LON-CL1.

Lab Review Questions

Question: What is the difference between ADMX and ADML files?

Question: The Sales Managers group should be exempted from the desktop lockdown policy that is being applied to the entire Sales OU. All sales user accounts and sales groups reside in the Sales OU. How would you exempt the Sales Managers group?

Question: What Windows command can you use to force the immediate refresh of all GPOs on a client computer?

Module Review and Takeaways

Module Review Questions

Question: What are some of the advantages and disadvantages of using site-level GPOs?

Question: You have a number of logon scripts that map network drives for users. Not all users need these drive mappings, so you must ensure that only the desired users receive the mappings. You want to move away from using scripts. What is the best way to map network drives for selected users without using scripts?

Best Practices

The following are recommended best practices:

- Do not use the Default Domain and Default Domain Controllers policies for uses other than their default uses. Instead, create new policies.
- Limit the use of security filtering and other mechanisms that make diagnostics more complex.
- If they have no settings configured, disable the User or Computer sections of policies.
- If you have multiple administration workstations, create a Central Store.
- Add comments to your GPOs to explain what the policies are doing.
- Design your OU structure to support Group Policy application.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
A user is experiencing abnormal behavior on their workstation.	
All users in a particular OU are having issues, and the OU has multiple GPOs applied.	

Tools

Tool	Use	Where to find it
Group Policy Management Console (GPMC)	Controls all aspects of Group Policy	In Server Manager, on the Tools menu
Group Policy Management Editor snap-in	Configure settings in GPOs	Accessed by editing any GPO
Resultant Set of Policy (RSOP)	Determine what settings are applying to a user or computer	In the GPMC
Group Policy Modeling Wizard	Test what would occur if settings were applied to users or computers, prior to actually applying the settings	In the GPMC
Local Group Policy Editor	Configure Group Policy settings that apply only to the local computer	Accessed by creating a new Microsoft Management Console (MMC) on the local computer, and adding the Group Policy Management Editor snap-in

Module 12

Securing Windows Servers Using Group Policy Objects

Contents:

Module Overview	12-1
Lesson 1: Windows Operating Systems Security Overview	12-2
Lesson 2: Configuring Security Settings	12-6
Lab A: Increasing Security for Server Resources	12-17
Lesson 3: Restricting Software	12-24
Lesson 4: Configuring Windows Firewall with Advanced Security	12-28
Lab B: Configuring AppLocker and Windows Firewall	12-33
Module Review and Takeaways	12-39

Module Overview

Protecting IT infrastructure has always been a priority for organizations. Many security risks threaten companies and their critical data. When companies fail to have adequate security policies, they can lose data, experience server unavailability, and lose credibility.

To protect against security threats, companies must have well-designed security policies that include many organizational and IT-related components. Security policies must be evaluated on a regular basis because as security threats evolve, so must IT evolve.

Before you begin designing security policies to help protect your organization's data, services, and IT infrastructure, you must learn how to identify security threats, plan your strategy to mitigate security threats, and secure your Windows Server® 2012 infrastructure.

Objectives

After completing this module, you will be able to:

- Describe Windows Server operating system security.
- Configure security settings by using Group Policy.
- Increase security for server resources.
- Restrict unauthorized software from running on servers and clients.
- Configure Windows Firewall with Advanced Security.

Lesson 1

Windows Operating Systems Security Overview

As organizations expand their availability of network data, apps, programs, and systems, ensuring the security of network infrastructure becomes more challenging. Security technologies in Windows Server 2012 enable organizations to provide better protection for their network resources and organizational assets in increasingly complex environments and business scenarios. This lesson reviews the tools and concepts that are available for implementing security within a Windows 8® and Windows Server 2012 infrastructure.

Windows Server 2012 includes numerous features that provide different methods for implementing security. These features combine to form the core of the security functionality of Windows Server 2012. Understanding these features and their associated concepts and being familiar with their basic implementation are critical to maintaining a secure environment.

Lesson Objectives

After completing this lesson, you will be able to:

- Identify security risks for Windows Server 2012 and the costs associated with them.
- Apply the defense-in-depth model to increase security.
- Describe best practices for increasing Windows Server 2012 security.

Discussion: Identifying Security Risks and Costs

The first step in defending your systems is identifying the potential security risks and their associated costs. Once you do that, you can begin to make accurate decisions about how to allocate resources to mitigate those risks.

Review the question on the slide and participate in the discussion to identify some of the security risks in Windows-based networks, and their associated costs.

- What are some of security risks in Windows-based networks?

10 minutes



Applying Defense-In-Depth to Increase Security

You can mitigate risks to your organization's computer network by providing security at various infrastructure layers. The term *defense-in-depth* is often used to describe the use of multiple security technologies at different points throughout your organization.

Defense-in-depth technologies include layers of security that extend from user policies, to the app, and to the data itself.

Policies, Procedures, and Awareness

Security policy measures need to operate within the context of organizational policies regarding security best practices. For example, enforcing a strong user password policy is not helpful if users write down their passwords and place them next to their computer screens; users must be educated about how to protect their passwords. Another example of a security best practice is ensuring that users do not leave their desktop computer without first locking the desktop or signing off of the computer. When establishing a security foundation for your organization's network, it is a good idea to start by establishing appropriate policies and procedures and then make users aware of them. You can then progress to the other aspects of the defense-in-depth model.

Physical Security

If any unauthorized person can gain physical access to a computer on your network, then most other security measures can be more easily bypassed. You must ensure that computers containing the most sensitive data (such as servers) are physically secure, and that physical access is granted only to authorized personnel.

Perimeter

These days, no organization is an isolated enterprise. Organizations operate within the Internet, and many organizational network resources are available from the Internet. This might include a website that describes your organization's services, or internal services that you make available externally—such as web conferencing and email—so that users can work from home or from branch offices.

Perimeter networks mark the boundary between public and private networks. Providing reverse proxy servers in the perimeter network enables you to provide more secure corporate services across the public network.

Many organizations implement network access quarantine control, where computers that connect to the corporate network are checked for different security criteria, such as whether the computer has the latest security updates, antivirus updates, and other company-recommended security settings. If these criteria are met, the computer is allowed to connect to corporate network. If not, the computer is placed in an isolated network, called a *quarantine*, with no access to corporate resources. Once the computer has its security settings remediated, it is removed from the quarantine network and is allowed to connect to corporate resources.

Defense-in-depth uses a layered approach to security

- Reduces an attacker's chance of success
- Increases an attacker's risk of detection

Policies, procedures, and awareness	Security documents, user education
Physical security	Guards, locks, tracking devices
Perimeter	Firewalls, network access quarantine control
Networks	Network segments, IPsec, Forefront TMG 2010
Host	Hardening, authentication, update management
Application	Application hardening, antivirus
Data	ACLs, EFS, BitLocker, backup/restore procedures



Note: A reverse proxy, such as Microsoft Forefront® Threat Management Gateway 2010 (Forefront TMG), enables you to publish services such as email or web services, from the corporate intranet without placing the email or web servers in the perimeter or exposing them to external users. Forefront TMG acts as both reverse proxy and as a firewall solution.

Networks

Once you connect your computers to a network (either internal or public), they are susceptible to a number of threats including eavesdropping, spoofing, denial of service, and replay attacks. By implementing Internet Protocol Security (IPsec), you can encrypt network traffic and protect data while in transfer between computers.

When communication takes place over public networks, such as by employees who are working from home or from remote offices, as a best practice they should connect to a firewall solution such as Forefront TMG 2010 to protect from different types of network threats.

Host Computer Security Hardening

The next layer of defense is the layer that is used for the host computer. Together, the following steps form a process that is referred to as *host computer security hardening*. On your host computer, you must:

- Keep computers secure with the latest security updates.
- Configure security policies, such as password complexity.
- Configure the host firewall.
- Install antivirus software.

Application Security Hardening

Apps are only as secure as your latest security update. Together, the following steps form a process that is called *application security hardening*:

- Consistently use the Windows Update feature or the app vendor update web sites to keep your apps up-to-date.
- Test apps to determine if they have any security vulnerabilities that might allow an external attacker to compromise them or other network components.

Data Security

The final layer of security is data security. To help ensure the protection of your network, you should:

- Ensure the proper use of file user permissions by using Access Control Lists (ACLs).
- Implement the encryption of confidential data with Encrypting File System (EFS).
- Perform regular data backups.



Additional Reading:

- For the latest Microsoft security bulletin and advisory information, see Security for IT Pros at <http://go.microsoft.com/fwlink/?LinkID=266741>.
- For more information about common types of network attacks, see <http://go.microsoft.com/fwlink/?LinkID=266742>.

Question: How many layers of the defense-in-depth model should you implement in your organization?

Best Practices for Increasing Security


Consider the following best practices for increasing security:

- Apply all available security updates as quickly as possible following their release. You should strive to implement security updates as soon as possible to ensure that your systems are protected from known vulnerabilities. Microsoft publicly releases the details of any known vulnerabilities after an update has been released, which can lead to an increased volume of malware attempting to exploit the vulnerability. However, you must still ensure that you adequately test updates before they are applied widely within your organization.
- Follow the principle of least privilege. Provide users and service accounts with the lowest permission levels required to complete their necessary tasks. This ensures that any malware using those credentials is limited in its impact. It also ensures that users are limited in their ability to accidentally delete data or modify critical operating system settings.
- Mandate that administrators use separate administrative accounts for administration and configuration changes. This ensures that administrators, while browsing the Internet or reading email, are not exposing a user account that has virtually unlimited access to the IT environment.
- Restrict administrator console sign in. Signing in locally at a console is a greater risk to a server than accessing data remotely. This is because some malware can only infect a computer by using a user session at the desktop. If you allow administrators to use Remote Desktop Connection for server administration, ensure that enhanced security features such as User Account Control (UAC) are enabled.
- Restrict physical access. If someone has physical access to your servers, that person has virtually unlimited access to the data on that server. An unauthorized person could use a wide variety of tools to quickly reset the password on local administrator accounts and allow local access, or use a USB drive to introduce malware. BitLocker can be effective at limiting or reducing the effectiveness of some physical attacks.

Some best practices for increasing security are:

- Apply all available security updates quickly
- Follow the principle of least privilege
- Use separate administrative accounts
- Restrict administrator console sign-in
- Restrict physical access



 **Additional Reading:** For more information about best practices for enterprise security, see <http://go.microsoft.com/fwlink/?LinkID=266743>.

Lesson 2

Configuring Security Settings

Once you have learned about security threats, risks, and best practices for increasing security, you can start configuring security for your Windows 8 and Windows Server 2012 environment. This lesson explains how to configure security settings.

To apply those security settings to multiple users and computers in your organization, you can use Group Policy. For example, you can configure password policy settings by using Group Policy, and then deploy them to multiple users.

Group Policy has a large security component that you can use to configure security for both users and computers. You can apply security consistently across the organization in Active Directory® Domain Services (AD DS) by defining security settings in a Group Policy Object (GPO) that is associated with a site, domain, or organizational unit (OU).



Additional Reading: For a detailed list of Group Policy settings, see the document "Group Policy Settings Reference for Windows and Windows Server" at: <http://go.microsoft.com/fwlink/?LinkID=266744>.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe how to configure security templates.
- Describe what user rights are, and how to configure them.
- Describe how to configure Security Options.
- Describe how to configure UAC.
- Describe how to configure Security Auditing.
- Describe how to configure Restricted Groups.
- Describe how to configure Account Policy Settings.
- Describe Security Compliance Manager (SCM).
- Install and Use SCM.

Configuring Security Templates

Security templates are files that you use to manage and configure security settings on Windows-based computers. Depending on the various categories of security settings, security templates are divided into logical sections. You can configure each of the following sections according to a company's needs and requests:

- Account Policies. Password Policy, Account Lockout Policy, and Kerberos Policy
- Local Policies. Audit Policy, User Rights Assignment, and Security Options

Security Templates categories:

- Account Policies
- Local Policies
- Event Log
- Restricted Groups
- System Services
- Registry
- File System

How Security Templates are distributed:

- Secedit.exe
- Security Templates Snap-in
- Security Configuration and Analysis Wizard
- Group Policy
- Security Compliance Manager (SCM)

- Event Log. App, system, and security event log settings
- Restricted Groups. Membership of groups that have special rights and permissions
- System Services. Startup and permissions for system services
- Registry. Permissions for registry keys
- File System. Permissions for folders and files

When you configure a security template, you can use it to configure a single computer or to configure multiple computers on the network. The following are a few ways that you can configure and distribute security templates:

- Secedit.exe. You can use the secedit.exe command-line tool to compare the current configuration of a computer running Windows Server 2012 to specific security templates.
- Security Templates snap-in. You can use this snap-in to create a security policy by using security templates.
- Security Configuration and Analysis Wizard. You can use this wizard to analyze and configure computer security.
- Group Policy. You can use Group Policy to analyze and configure computer settings, and to distribute specific security settings.
- SCM. You can use the SCM to view security settings, compare settings to security baselines (which are groups of settings designed on the basis of Microsoft security guides and best practices), customize settings, and import or export GPO backups. SCM is described in detail in a later topic.

Configuring User Rights

User rights assignment refers to the ability to perform actions in the operating system. Each computer has its own set of user rights, such as the right to change the system time. By default, most rights are granted either to the Local System or to the Administrator.

Privileges and Logon rights are two types of user rights:

- *Privileges* define access to computer and domain resources. For example, rights to back up files and directories.
- *Logon rights* define who is authorized to sign in to a computer, and how they can sign in. For example, logon rights may define the right to sign in to a system locally.

You can configure rights through Group Policy. Initially, the default domain policy does not have user rights defined.

You can configure settings for User Rights by accessing the following location from the Group Policy Management Console (GPMC):

- Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\User Rights Assignment

User Rights Types:

- Privileges
- Logon Rights

Examples of common user rights:

- Add workstations to domain
- Allow log on locally
- Allow log on through Remote Desktop Services
- Back up files and directories
- Change the system time
- Force shutdown from a remote computer
- Shut down the system



Some examples of commonly used user rights (and policies configured by them) are:

- Add workstations to domain. Determines which users or groups can add workstations to the domain.
- Allow log on locally. Determines which users can sign in to the computer.
- Allow log on through Remote Desktop Services. Determines which users or groups have permission to sign in by using Remote Desktop Services Client.
- Back up files and directories. Determines which users have permissions to back up files and folders on a computer.
- Change the system time. Determines which users or groups have the rights to change the time and date on the internal clock of the computer.
- Force shutdown from a remote system. Determines which users are allowed to shut down a computer from a remote location on the network.
- Shut down the system. Determines which of the users who are locally signed in to a computer are allowed to shut down the computer.

Configuring Security Options

You can also use Group Policy to access and configure security options. The computer security settings that you can configure in Security Options include the following:

- Administrator and Guest account names
- Access to CD/DVD drives
- Digital data signatures
- Driver installation behavior
- Logon prompts
- UAC

Security options settings:

- Administrator and Guest account names
- Access to CD/DVD drives
- Digital data signatures
- Driver installation behavior
- Logon prompts
- User Account Control

Examples:

- Prompt user to change password before expiration
- Do not display last user name
- Rename administrator account
- Restrict CD-ROM access to locally logged-on users only

You can configure settings for Security Options by accessing the following location from the GPMC:

- Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\Security Options

The following are examples of commonly used Security Options:

- Prompt user to change password before expiration. Determines how many days before a user's password expiring that the operating system provides a warning.
- Interactive logon: Do not display last user name. Determines whether the name of the last user to sign in to the computer is displayed in the Windows logon window.
- Accounts: Rename administrator account. Determines whether a different account name is associated with the security identifier (SID) for the administrator account.
- Devices: Restrict CD-ROM access to locally logged-on user only. Determines whether a CD-ROM is accessible to both local and remote users simultaneously.

Configuring User Account Control

Administrative accounts carry with them a higher degree of security risk. When an administrative account is signed in, its privileges allow access to the entire Windows operating system, including the registry, system files, and configuration settings. As long as an administrative account is signed in, the system is vulnerable to attack and has the potential to be compromised.

User Account Control (UAC) is a security feature that helps prevent unauthorized changes to a computer. It does this by asking the user for permission or for administrator credentials before performing actions that could affect the computer's operation or that could change settings that affect multiple users.

By default, both standard users and administrators run apps and access resources in the security context of a standard user. The UAC prompt provides a way for a user to elevate his or her status from a standard user account to an administrator account without signing out, switching users, or running an app by using different credentials. Because of this, UAC creates a more secure environment in which to run and install apps.

When an app requires administrator level permission, UAC notifies the user as follows:

- If the user is an administrator, the user confirms to elevate his or her permission level and continue. This process of requesting approval is known as *Admin Approval Mode*.



Note: Since Windows Server 2008, the built-in Administrator account does not run in Admin Approval Mode. The result is that no UAC prompts display when using the local Administrator account.

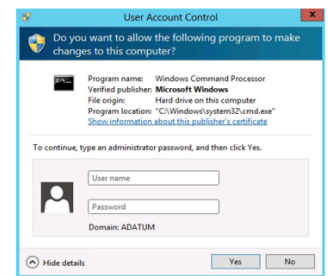
- If the user is not an administrator, then a username and password for an account that has administrative permissions needs to be entered. Providing administrative credentials temporarily gives the user administrative privileges, but only to complete the current task. After the task is complete, permissions change back to those of a standard user.

When using this process of notification and elevation to administrator account privileges, changes cannot be made to the computer without the user knowing, because a prompt asks the user for permission or for administrator credentials. This can help prevent malicious software (malware) and spyware from being installed on or making changes to a computer.

UAC allows the following system-level changes to occur without prompting, even when a user is signed in as a local user:

- Install updates from Windows Update
- Install drivers from Windows Update or those that are packaged with the operating system
- View Windows operating system settings
- Pair Bluetooth devices with the computer
- Reset the network adapter, and perform other network diagnostic and repair tasks

- UAC is a security feature that prompts the user for an administrative user's credentials if the task requires administrative permissions
- UAC enables users to perform common daily tasks as non-administrators



Modifying UAC Behavior

You can modify the UAC notification experience to adjust the frequency and behavior of UAC prompts. To modify UAC behavior on a single computer, access the Windows Server 2012 control panel in System and Security.

You can configure settings for UAC by accessing the following location from the GPMC:

- Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\Security Options

The following are examples of some GPO settings that you can configure for UAC:

- User Account Control: Run all administrators in Admin Approval Mode. Controls the behavior of all UAC policy settings for the computer. If this setting is disabled, UAC will not run on this computer.
- User Account Control: Administrator Approval Mode for the built-in Administrator account. When you enable this setting, the built-in Administrator account uses Admin Approval Mode.
- User Account Control: Detect app installations and prompt for elevation. This setting controls the behavior of app installation detection for the computer.
- User Account Control: Only elevate executables that are signed and validated. When you enable this setting, a Public Key Infrastructure (PKI) check is performed on the executable file to verify that it originates from a trusted source. If the file is verified, then the file is permitted to run.



Note: By default, UAC is not configured or enabled in Server Core installations of Windows Server 2012.

Configuring Security Auditing

Typically, one of the components of an organization's security strategy is recording user activities. The activities may include successful or unsuccessful attempts to access business-critical data that is stored in different folders, or successful or unsuccessful sign-in attempts on different servers. Recording these security-related events is called *Security Auditing*. Security auditing adds entries to the Security Event Log that administrators can then view in the Event Viewer.

Information in security event logs can help your organization audit their compliance with important business-related and security-related goals by tracking precisely defined activities such as:


- An administrator who modified settings or data on servers that contain highly confidential information.
- An employee within a defined group that has accessed an important folder containing data from different departments.
- A user who is trying to sign in to his or her account repeatedly without success from an internal company computer. You might find that the employee who owns that user account was on a vacation that week, which means some other employee was trying to sign in with a different user account.

You can configure settings for Security Auditing by accessing the following location from the GPMC:

- Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\Audit Policy

When using security auditing to log security-related events, you can:

- Configure security auditing according to your company's security regulations
- Filter the Security Event Log in Event Viewer to find specific security related events



The following are examples of some GPO settings that you can configure for auditing:

- Audit account logon events. Determines whether the operating system audits each time the computer validates an account's credentials.
- Audit accounting management. Determines whether to audit each event of account management, such as creating, changing, renaming, or deleting a user account, changing a password, or enabling or disabling a user account.
- Audit object access. Determines whether operating system audits have access to non-Active Directory objects, such as folders or files. Before configuring audit settings with Group Policy, you must configure system access control lists (SACLs) on folders or files to allow auditing for a specific type of action, such as write, read, or modify.
- Audit system events. Determines whether the operating system audits system-related events, such as attempting to change the system time, attempting a system startup or shutdown, or the security log size exceeding a configurable threshold warning.

When working with security auditing, be aware of the following concerns.

- Configuring Windows Server 2012 to audit activities generates a large amount of data that is difficult to analyze.
- A large amount of data might cause servers or domain controllers to run out of disk space because the Security Event Log can become very large. Recording a large amount of data may also cause poor performance on legacy servers.



Additional Reading: For more information about security auditing, see the article "What's New in Security Auditing" at <http://go.microsoft.com/fwlink/?LinkID=266747>.

Configuring Restricted Groups

In some cases, you may want to control the membership of certain groups in a domain—such as the local administrators group—to prevent other user accounts being added to those groups.

You can use the Restricted Groups policy to control group membership by using either of the following methods:

- You can specify which members are added to a group.

If you choose this option, then when you define a Restricted Groups policy and refresh Group Policy, all current members remain while the defined members in the policy are added to the existing membership.

- You can specify which members make up the total membership of a group.

If you choose this option, then when you define a Restricted Groups policy and refresh Group Policy, any current member of a group that is not on the Restricted Groups policy members list is removed, including default members such as the Domain Admins group.

Group Policy can control group membership:

- For any group on a domain-joined computer, by applying a Group Policy Object (GPO) to the Organizational Unit (OU) containing the computer account
- For any group in AD DS, by applying a GPO to the Domain Controller's OU

Although you can control domain groups by assigning Restricted Groups policies to domain controllers, you should use this setting primarily to configure membership of critical groups such as Enterprise Admins and Schema Admins. You can also use this setting to control the membership of built-in local groups on workstations and member servers. For example, you can place the Helpdesk group into the local Administrators group on all workstations.

You cannot specify local users in a domain GPO. Local users who currently are in the local group that the Restricted Groups policy controls will be removed, depending on the Restricted Groups policy option that you choose. The only exception to this is that the local Administrators account is always in the local Administrators group.

You can configure settings for Restricted Groups by accessing the following location from the GPMC:

- Computer Configuration\Policies\Windows Settings\Security Settings\Restricted Groups

Configuring Account Policy Settings

Account policies protect your organization's accounts and data by mitigating the threat of attacks that try to guess account username and password (these are sometimes referred to as *brute force attacks*). Securing your network environment requires that all users use strong passwords. You use password policy settings to control the complexity and lifetime of user passwords. You configure password policy settings through Group Policy.

Account policies mitigate the threat of brute force guessing of account passwords

Policies	Default settings
Password	<ul style="list-style-type: none"> • Controls complexity and lifetime of passwords • Max password age: 42 days • Min password age: 1 day • Min password length: 7 characters • Complex Password: enabled • Store password using reversible encryption: disabled
Account lockout	<ul style="list-style-type: none"> • Controls how many incorrect attempts can be made • Lockout duration: not defined • Lockout threshold: 0 invalid logon attempts • Reset account lockout after: not defined
Kerberos	<ul style="list-style-type: none"> • Subset of the attributes of domain security policy • Can only be applied at the domain level

Account Policies

Account policy components include password policies, account lockout policies, and Kerberos policies.

The policy settings under Account policies are implemented at the domain level. A Windows Server 2012 domain can have multiple password and account lockout policies, which are called *fine-grained password policies*. You can apply these multiple policies to a user or to a global security group in a domain, but not to an OU.



Note: If you need to apply a fine-grained password policy to users of an OU, you can use a *shadow group*, which is a global security group that is logically mapped to an OU.

You can configure settings for Account policies by accessing the following location from the GPMC:

- Computer Configuration\Policies\Windows Settings\Security Settings\Account Policies

Password Policy

Password policies that you can configure are listed in the following table.

Policy	Function	Best practice
Password must meet complexity requirements	<ul style="list-style-type: none"> Requires passwords to: Be at least as long as specified by the Minimum Password Length, with a minimum of 3 characters if the Minimum Password Length is set to 0. Contain a combination of at least three of the following types of characters: uppercase letters, lowercase letters, numbers, and symbols (punctuation marks). Must not contain the user's user name or screen name. 	<p>Enable this setting. These complexity requirements can help ensure a strong password. Strong passwords are more difficult to decrypt than those containing simple letters or numbers.</p> <p>Instruct users to use pass phrases to create long passwords that are easy to remember.</p>
Enforce password history	<p>Prevents users from creating a new password that is the same as their current password or a recently used password.</p> <p>To specify how many passwords are remembered, provide a value. For example, a value of 1 means that only the last password is remembered, and a value of 5 means that the previous five are remembered.</p>	<p>The greater number ensures better security. The default value is 24.</p> <p>Enforcing password history ensures that passwords that have been compromised are not used repeatedly.</p>
Maximum password age	Sets the maximum number of days that a password is valid. After this number of days, the user will have to change the password.	<p>The default value is 42 days. Setting the number of days too high provides hackers with an extended window of opportunity to crack or brute force the password. Setting the number of days too low frustrates users who have to change their passwords too frequently, and could result in more frequent calls to the IT help desk.</p>
Minimum password age	Sets the minimum number of days that must pass before a password can be changed.	<p>Set the minimum password age to at least 1 day. By doing so, you require that the user can only change their password once a day. This helps enforce other settings.</p> <p>For example, if the past five passwords are remembered, this ensures that at least five days must pass before the user can reuse the original password. If the minimum password age is set to 0, the user can change their password six times on the same day and begin reusing the original password on the same day.</p>

(continued from previous page)

Policy	Function	Best practice
Minimum password length	Specifies the fewest number of characters that a password can have.	Set the length to between 8 and 12 characters (provided that they also meet complexity requirements). A longer password is more difficult to crack than a shorter password, assuming the password is not a common word.
Store passwords by using reversible encryption	Provides support for apps that need to know a user password for authentication purposes.	Do not use this setting unless you use an app that requires it. Enabling this setting decreases the security of stored passwords.

Account Lockout Policy

Account Lockout Policies that you can configure are listed in the following table.

Policy	Function	Best practice
Account lockout threshold	Specifies the number of failed login attempts that are allowed before the account is locked. For example, if the threshold is set to 3, the account is locked out after a user enters incorrect login information three times.	A setting of 5 allows for reasonable user error, and limits repeated login attempts for malicious purposes. Note that the threshold can have a negative impact by allowing for a denial of service attack on user objects, especially from the Internet. Because of this, some organizations are going toward a higher threshold.
Account lockout duration	Allows you to specify a timeframe, in minutes, after which the account automatically unlocks and resumes normal operation. If you specify 0, then the account is locked indefinitely until an administrator manually unlocks it.	After the threshold has been reached and the account is locked out, the account should remain locked long enough to block or deter any potential attacks, but short enough not to interfere with productivity for legitimate users. A duration of 30 to 90 minutes works well in most situations.
Reset account lockout counter after	Defines a timeframe for counting the incorrect login attempts. If the policy is set for one hour, and the account lockout threshold is set for three attempts, a user can enter the incorrect login information three times within one hour. If they enter incorrect information twice, but get it correct the third time, the counter will reset after one hour has elapsed (from the first incorrect entry) so that future failed attempts will again start counting at one.	Using a timeframe between 30 and 60 minutes is usually sufficient to deter automated attacks and manual attempts by an attacker to guess a password.

Kerberos Policy

This policy is for domain user accounts, and determines Kerberos-related settings such as ticket lifetimes and enforcement. Kerberos policies do not exist in Local Computer Policy.

Using Security Compliance Manager

SCM provides the following core uses.

- Use SCM to maintain compliance and report on compliance. Many organizations adhere to specific industry or government regulations and must submit to periodic compliance tests. SCM can be used to validate that computers are configured for compliance, especially when paired with System Center 2012 R2 Configuration Manager, which has a Desired Configuration Management (DCM) component that can be used for gathering compliance information. In this scenario, a baseline from SCM can be exported and used in System Center 2012 R2 Configuration Manager.
- Use SCM to configure computers for compliance or security policies. SCM can be used to minimize the administrative effort when configuring computers for compliance or security policies. SCM can export a GPO which can then be linked to the appropriate containers in AD DS.
- Use SCM to maintain settings across two independent environments. Multiple GPOs can be imported into SCM and then used for comparing and/or merging settings across environments. This is useful when an organization has a production environment and a development environment (or multiple iterations of each environment).
- Use SCM to learn about Microsoft recommended security settings. The built-in security guides are in-depth and product-specific. They contain pertinent information and recommendations that will help an organization understand risks and mitigation. They can be used to formulate or update security policies and ensure that IT teams have the security knowledge to successfully deploy and maintain the environments.

What is Security Compliance Manager?

Overview

Security Compliance Manager (SCM) is a free tool from Microsoft that helps administrators secure computers whether the computers reside locally, remotely, or in the cloud. SCM is a Microsoft Solution Accelerator, currently in version 3.0, which automates some of the administrative tasks of securing computers. SCM works as a standalone tool or can be enhanced by combining it with System Center 2012 R2 Configuration Manager.

SCM is a free tool from Microsoft that helps administrators secure computers whether the computers reside locally, remotely, or in the cloud. It features:

- Baselines
- Security guides
- Support for standalone computers
- Import GPO backups

What does SCM do?

SCM provides the following core features:

- **Baselines.** Baselines are based on Microsoft security guides and best practices; they are a foundation for deploying new settings. The baseline settings are specific to an operating system version, a specific product version, or a specific component and can be downloaded or imported into SCM in the form of .cab files as new ones become available. The baselines can be compared to your existing settings, or to default settings. SCM provides an interface to view all of the settings. The settings can be customized and then exported as a GPO backup.
- **Security guides.** The security guides are Microsoft guides for the major operating system versions and product versions. They contain instructions and recommendations to help secure your environment. Some of the guides that are included with SCM are the guides for Windows® 7 SP1, Internet Explorer 10, Microsoft Exchange Server, and Windows Server 2012.
- **Support for deploying policies to standalone computers.** In addition to automating the deployment of settings for domain-joined computers by using Group Policy, SCM helps minimize the administrative overhead of securing non-domain-joined computers too.
- **Import backups of existing GPOs.** Existing backed-up GPOs can be imported into SCM for comparison with the baselines. You can then customize the settings before exporting the new settings to a GPO backup.

Requirements for SCM 3.0

SCM can be installed on a Windows client operating system or on a Windows server operating system. The following prerequisites are required for SCM 3.0.

- **Microsoft Visual C++ 2010 x86.** The installer comes prepackaged with SCM 3.0. If it is not already installed on the destination computer, then the SCM installer prompts to install it.
- **Microsoft SQL Server 2008 (including Express edition) or later** installed on the destination computer. If SQL is not installed then the SCM installer will install Microsoft SQL Server 2008 Express.
- **Microsoft Word and Microsoft Excel.** Although not required as part of the installation of SCM, some of the supporting material and guides require Word or Excel. In the case of text documents, WordPad (installed by default in Windows) can suffice but may be a degraded experience. However, all of the documents can be saved elsewhere and opened up from another computer that already has Word and Excel.

Lab A: Increasing Security for Server Resources

Scenario

Your manager has given you some security-related settings that need to be implemented on all member servers. You also need to implement file system auditing for a file share used by the Marketing department. Finally, you need to implement auditing for domain logons.

Objectives

After completing this lab, you will be able to:

- Use Group Policy to secure member servers.
- Audit who is accessing specific files.
- Audit domain logons.

Lab Setup

Estimated time: 60 minutes

Virtual machines	20410C-LON-DC1 20410C-LON-SVR1 20410C-LON-SVR2 20410C-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 **20410C-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**.
Wait until the virtual machine starts.
4. Sign in using the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
5. Repeat steps 2 through 4 for 20410C-LON-SVR1 and 20410C-LON-SVR2.
6. Repeat steps 2 and 3 for 20410C-LON-CL1. Do not sign in to LON-CL1 until directed to do so.

Exercise 1: Using Group Policy to Secure Member Servers

Scenario

A. Datum uses the Computer Administrators group to provide administrators with permissions to administer member servers. As part of the installation process for a new server, the Computer Administrators group from the domain is added to the local Administrators group on the new server. Recently, this important step was missed when configuring several new member servers.

To ensure that the Computer Administrators group is always given permission to manage member servers, your manager has asked you to create a GPO that sets the membership of the local Administrators group on member servers to include Computer Server Administrators. This GPO also needs to enable Admin Approval Mode for UAC.

The main tasks for this exercise are as follows:

1. Create a Member Servers organizational unit (OU) and move servers into it.
2. Create a Server Administrators group.
3. Create a Member Server Security Settings Group Policy Object (GPO) and link it to the Member Servers OU.
4. Configure group membership for local administrators to include Server Administrators and Domain Admins.
5. Verify that Computer Administrators has been added to the local Administrators group.
6. Modify the Member Server Security Settings GPO to remove Users from Allow Log On Locally.
7. Modify the Member Server Security Settings GPO to enable User Account Control: Admin Approval Mode for the Built-in Administrator account.
8. Verify that a non-administrative user cannot sign in to a member server.

► Task 1: Create a Member Servers organizational unit (OU) and move servers into it

1. On LON-DC1, open **Active Directory Users and Computers**.
2. Create a new OU named **Member Servers OU**.
3. Move servers **LON-SVR1** and **LON-SVR2** to **Member Servers OU**.

► Task 2: Create a Server Administrators group

- On LON-DC1, in **Member Servers OU**, create a new global security group called **Server Administrators**.

► Task 3: Create a Member Server Security Settings Group Policy Object (GPO) and link it to the Member Servers OU

1. On LON-DC1, open the **Group Policy Management Console**.
2. In the Group Policy Management Console, in the Group Policy Objects container, create a new GPO with a name **Member Server Security Settings**.
3. In the Group Policy Management Console, link the **Member Server Security Settings** to **Member Servers OU**.

► Task 4: Configure group membership for local administrators to include Server Administrators and Domain Admins

1. On LON-DC1, open the **Group Policy Management Console**.
2. For the Default Domain Policy, open the Group Policy Management Editor window.

3. In the Group Policy Management Editor window, go to **Computer Configuration\Policies\Windows Settings\Security Settings\Restricted Groups**.
4. Add the **Server Administrators** and **Domain Admins** groups to the **Administrators** group.
5. Close the Group Policy Management Editor window.

► **Task 5: Verify that Computer Administrators has been added to the local Administrators group**

1. Switch to LON-SVR1.
2. Open a Windows PowerShell® window, and at the Windows PowerShell prompt, type following command:

```
Gpupdate /force
```

3. Open **Server Manager**, open the Computer Management console, and then expand **Local Users and Groups**.
4. Confirm that the **Administrators** group contains both **ADATUM\Domain Admins** and **ADATUM\Server Administrators** as members.
5. Close the Computer Management console.

► **Task 6: Modify the Member Server Security Settings GPO to remove Users from Allow Log On Locally**

1. On LON-DC1, in the Group Policy Management Console, edit the **Member Server Security Settings** GPO.
2. In the Group Policy Management Editor window, go to **Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\User Rights Assignment**.
3. Configure **Allow log on locally** for **Domain Admins** and **Administrators** security groups.

► **Task 7: Modify the Member Server Security Settings GPO to enable User Account Control: Admin Approval Mode for the Built-in Administrator account**

1. On LON-DC1, in the Group Policy Management Editor window, go to **Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\Security Options**.
2. Enable **User Account Control: Admin Approval Mode for the Built-in Administrator account**.
3. Close the Group Policy Management Editor window.

► **Task 8: Verify that a non-administrative user cannot sign in to a member server**

1. Switch to LON-SVR1.
2. Open a Windows PowerShell window, and at the Windows PowerShell prompt, type following command:

```
Gpupdate /force
```

3. Sign out from LON-SVR1.
4. Try to sign in to LON-SVR1 as **Adatum\Adam** with the password **Pa\$\$w0rd**.
Verify that you cannot sign in to LON-SVR1.

5. To prepare for the next exercise, sign out of LON-SVR1, and sign back in to LON-SVR1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.

Results: After completing this exercise, you should have used Group Policy to secure Member servers.

Exercise 2: Auditing File System Access

Scenario

The manager of the Marketing department has concerns that there is no way to track who is accessing files that are on the departmental file share. Your manager has explained that only users with permissions are allowed to access the files. However, the manager of the Marketing department would like to try recording who is accessing specific files.

Your manager has asked you to enable auditing for the file system that is on the Marketing department file share, and to review the results with the manager of the Marketing department.

The main tasks for this exercise are as follows:

1. Modify the Member Server Security Settings GPO to enable object access auditing.
2. Create and share a folder.
3. Enable auditing on the Marketing folder for Domain Users.
4. Create a new file in the file share from LON-CL1.
5. View the results in the security log on the domain controller.

► Task 1: Modify the Member Server Security Settings GPO to enable object access auditing

1. Switch to LON-DC1.
2. In the Group Policy Management Console, edit the **Member Server Security Settings** GPO.
3. In the Group Policy Management Editor window, go to **Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\Audit Policy**.
4. Enable **Audit object access** with both **Success** and **Failure** settings.
5. Sign out of LON-DC1.

► Task 2: Create and share a folder

1. Switch to LON-SVR1.
2. On LON-SVR1, on drive C, create a new folder with the name **Marketing**.
3. Configure the Marketing folder with Read/Write sharing permissions for user **Adam**.

► Task 3: Enable auditing on the Marketing folder for Domain Users

1. On LON-SVR1, in the Local Disk (C:) window, configure auditing on the **Marketing** folder, with following settings:
 - Select a principal: **Domain Users**
 - Type: **All**
 - Permission: **Read & execute, List folder content, Read, Write**
 - Leave other settings with their default values

2. Refresh Group Policy by typing the following command at the Windows PowerShell prompt:

```
gpupdate /force
```

► **Task 4: Create a new file in the file share from LON-CL1**

1. Sign in to LON-CL1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Open the Command Prompt window and type the following command:

```
gpupdate /force
```
3. Close the Command Prompt window.
4. Sign out from LON-CL1, and then sign in again as **Adatum\Adam** with the password **Pa\$\$w0rd**.
5. Open the **Marketing** folder on LON-SVR1, by using the following Universal Naming Convention (UNC) path: **\\LON-SVR1\Marketing**.
6. Create a text document with a name **Employees**.
7. Sign out from LON-CL1.

► **Task 5: View the results in the security log on the domain controller**

1. Switch to LON-SVR1, and start **Event Viewer**.
2. In the Event Viewer window, expand **Windows Logs**, and then open **Security**.
3. Verify that following event and information is displayed:
 - Source: Microsoft Windows Security Auditing
 - Event ID: 4663
 - Task category: File System
 - An attempt was made to access an object

Results: After completing this exercise, you should have enabled file system access auditing.

Exercise 3: Auditing Domain Logons

Scenario

After a security review, the IT policy committee has decided to begin tracking all user logons to the domain. Your manager has asked you to enable auditing of domain logons and verify that they are working.

The main tasks for this exercise are as follows:

1. Modify the Default Domain Policy GPO.
2. Run GPOupdate.
3. Sign in to LON-CL1 with an incorrect password.
4. Review event logs on LON-DC1.
5. Sign in to LON-CL1 with the correct password.
6. Review event logs on LON-DC1.

► Task 1: Modify the Default Domain Policy GPO

1. Sign in to LON-DC1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On LON-DC1, start **Server Manager**, and then from Server Manager, start **GPMC**.
3. On LON-DC1, in the Group Policy Management Console, edit the **Default Domain Policy** GPO.
4. In the Group Policy Management Editor window, go to **Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\Audit Policy**.
5. Enable **Audit account logon events** with both **Success** and **Failure** settings.
6. Update Group Policy by using the **gpupdate /force** command.

► Task 2: Run GPOUpdate

1. Sign in to LON-CL1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Open the Command Prompt window and type the following command:

```
gpupdate /force
```

3. Close the Command Prompt window, and sign out from LON-CL1.

► Task 3: Sign in to LON-CL1 with an incorrect password

- Sign in to LON-CL1 as **Adatum\Adam** with the password **password**.



Note: This password is intentionally incorrect to generate a security log entry that shows that that an unsuccessful sign in attempt has been made.

► Task 4: Review event logs on LON-DC1

1. On LON-DC1, start **Event Viewer**.
2. In the Event Viewer window, expand **Windows Logs**, and then click **Security**.
3. Review the event logs for the following message: "**Event ID 4771 Kerberos pre-authentication failed. Account Information: Security ID: ADATUM\Adam.**"

► Task 5: Sign in to LON-CL1 with the correct password

1. Sign in to LON-CL1 as **Adatum\Adam** with the password **Pa\$\$w0rd**.



Note: This password is correct, and you should be able to sign in successfully as **Adam**.

2. Sign out of LON-CL1.

► Task 6: Review event logs on LON-DC1

1. On LON-DC1, start **Event Viewer**.
2. In the Event Viewer window, expand **Windows Logs**, and then click **Security**.
3. Review the event logs for the following message: "**Event ID 4624 An account was successfully logged on. New Logon: Security ID: ADATUM\Adam.**"

Results: After completing this exercise, you should have enabled domain logon auditing.

► **Prepare for the next lab**

- To prepare for the next lab, leave the virtual machines running.

Lab Review Questions

Question: What happens if you configure the Computer Administrators group, but not the Domain Admins group, to be a member of the Local Administrators group on all the computers in a domain?

Question: Why do you need to not allow local logon on some computers?

Question: What happens when an unauthorized user tries to access a folder that has auditing enabled for both successful and unsuccessful access attempts?

Question: What happens when you configure auditing for domain logons for both successful and unsuccessful logon attempts?

Lesson 3

Restricting Software

Users need to have access to the apps that help them do their jobs. However, unnecessary or unwanted apps often get installed on client computers, whether unintentionally or for malicious or non-business purposes. Unsupported or unused software is not maintained or secured by the administrators, and could be used as an entry point for attackers to gain unauthorized access or spread computer viruses. Consequently, it is of the utmost importance for you to ensure that only necessary software is installed on all the computers in your organization. It is also vital that you prevent software from running that is not allowed or is no longer used or supported.

Lesson Objectives

After completing this lesson, you will be able to:

- Explain how to use software restriction policies (SRPs) to restrict unauthorized software from running on servers and clients.
- Describe the purpose of AppLocker®.
- Describe AppLocker rules and how to use them to restrict unauthorized software from running on servers and clients.
- Describe how to create AppLocker rules.

What Are Software Restriction Policies?

Introduced in the Windows® XP operating system and the Windows Server 2003 operating system, software restriction policies (SRPs) give administrators tools that they can use to identify and specify which apps are permitted to run on client computers. You configure and deploy SRP settings to clients by using Group Policy.

SRPs are used in Windows Server 2012 to provide Windows Vista® compatibility. An SRP set is made up of rules and security levels.

- Software Restriction Policies (SRPs) allow administrators to identify which apps are allowed to run on client computers
- SRPs can be based on the following:
 - Hash
 - Certificate
 - Path
 - Zone
- SRPs are applied through Group Policy

Rules

Rules govern how SRP responds to an app that is being run or installed. Rules are the key constructs within an SRP, and a group of rules together determines how an SRP responds to apps that are being run. Rules can be based on one of the following criteria that apply to the primary executable file for the app in question:

- Hash. A cryptographic fingerprint of the file.
- Certificate. A software publisher certificate that is used to digitally sign a file.
- Path. The local or UNC path to where the file is stored.
- Zone. The Internet zone.

Security Levels

Each applied SRP is assigned a security level that governs the way that the operating system reacts when the app that is specified in the rule is run. The three available security levels are as follows:

- **Disallowed.** The software identified in the rule will not run, regardless of the access rights of the user.
- **Basic User.** Allows the software identified in the rule to run as a standard, non-administrative user.
- **Unrestricted.** Allows the software identified in the rule to run unrestricted by SRP.

Using these three settings, there are two primary ways to use SRPs:

- If an administrator has a comprehensive list of all the software that is allowed to run on clients, the Default Security Level can be set to Disallowed. All apps that are allowed to run can be identified in SRP rules that apply either the Basic User or Unrestricted security level to each individual app, depending on the security requirements.
- If an administrator does not have a comprehensive list of the software that is allowed to run on clients, the Default Security Level can be set to Unrestricted or Basic User, depending on security requirements. All apps that are not allowed to run can then be identified in SRP rules, which would use a security level setting of Disallowed.

You can configure settings for SRPs by accessing the following location from the GPMC:

- Computer Configuration\Policies\Windows Settings\Security Settings\Software Restriction Policies

What Is AppLocker?

AppLocker, which was introduced in the Windows 7 operating system and Windows Server 2008 R2, is a security setting feature that controls which apps users are allowed to run.

AppLocker provides administrators a variety of methods for determining quickly and concisely the identity of apps that they may want to restrict, or to which they may want to permit access. You apply AppLocker through Group Policy to computer objects within an OU. You can also apply Individual AppLocker rules to individual AD DS users or groups.

AppLocker applies Application Control Policies in Windows Server 2012 and Windows 8. AppLocker contains capabilities and extensions that:

- Reduce administrative overhead
- Helps administrators control how users can access and use files:

- .exe files
- Windows Installer files (.msi and .msp files)
- scripts
- Packaged apps (Windows Store apps)
- DLLs

Benefits of AppLocker:

- Controls how users can access and run all types of apps
- Allows the definition of rules based on a wide variety of variables
- Provides for importing and exporting entire AppLocker policies

AppLocker also contains options for monitoring or auditing the application of rules. AppLocker can help organizations prevent unlicensed or malicious software from running, and can selectively restrict ActiveX® controls from being installed. It can also reduce the total cost of ownership by ensuring that workstations are standardized across the enterprise, and that users are running only the software and apps that are approved by the enterprise.

Using AppLocker technology, companies can reduce administrative overhead and help administrators control how users can access and use files, such as .exe files, scripts, Windows Installer files (.msi and .msp files), DLLs, and packaged apps (Windows Store apps).

You can use AppLocker to restrict software that:

- Is not allowed to be used in the company. For example, software that can disrupt employees' business productivity, such as social networking software, or software that streams video files or pictures that can use large amounts of network bandwidth and disk space.
- Is no longer used or it has been replaced with a newer version. For example, software that is no longer maintained, or for which licenses have expired.
- Is no longer supported in the company. Software that is not updated with security updates might pose a security risk.
- Should be used only by specific departments.

You can configure settings for AppLocker by accessing the following location from the GPMC:

- Computer Configuration\Policies\Windows Settings\Security Settings\Application Control Policies



Note: AppLocker uses the Application Identity service to verify the attributes of a file. This service should be configured to start automatically on each computer where AppLocker will be applied. If the Application Identity service is not running, then AppLocker policies will not be enforced.



Additional Reading: For more information about AppLocker, see AppLocker overview at <http://go.microsoft.com/fwlink/?LinkID=266745>.

AppLocker Rules

AppLocker defines rules based on file attributes that are derived from the digital signature of the file. File attributes in the digital signature include:

- Publisher name
- Product name
- File name
- File version

AppLocker defines rules based on file attributes such as:

- Publisher name
- Product name
- File name
- File version

Rule actions

- Allow or Deny conditions
- Enforce or Audit Only policies

Default Configuration

By default, no AppLocker policies are defined, which means that no apps are blocked. You can configure default rules for each rule collection. The default rules ensure that apps in the Program Files and Windows directories are allowed to run, and all apps are allowed to run for the Administrators group. The default rules should be enabled if you are going to implement AppLocker policies, as these apps are necessary for Windows operating systems to run and operate normally.

Allow and Deny Rule Actions

Allow and Deny are rule actions that allow or deny execution of apps based on a list of apps that you configure. The Allow action on rules limits execution of apps to an allowed list of apps, and blocks everything else. The Deny action on rules takes the opposite approach and allows the execution of any app except those on a list of denied apps. These actions also provide a means to identify exceptions to those actions.

Enforce or Audit Only

When AppLocker policy is set to Enforce, rules are enforced and all events are audited. When AppLocker policy is set to Audit Only, rules are evaluated and events are written in to the AppLocker Log, but no enforcement takes place. By using the Audit Only setting, administrators can gather information about apps that are being run, understand which apps will not run when enforcement is used, and can see the ramifications of AppLocker enforcement on the end users.

Demonstration: Creating AppLocker Rules

In this demonstration, you will see how to:

- Create a GPO to enforce the default AppLocker Executable rules.
- Apply the GPO to the domain.
- Test the AppLocker rule.

Demonstration Steps

Create a GPO to enforce the default AppLocker Executable rules

1. On LON-DC1, open the **Group Policy Management Console**.
2. Create a new GPO named **WordPad Restriction Policy**.
3. Edit the **WordPad Restriction Policy's Security Settings** by using AppLocker to create a new **Executable Rule**.
4. Set the permission of the new rule to **Deny**, the condition to **Publisher**, and then select **wordpad.exe**. If prompted, click **OK** to create default rules.
5. In the Group Policy Management Editor window, go to **Computer Configuration\Policies\Windows Settings\Security Settings\Application Control Policies\AppLocker**.
6. In AppLocker, configure enforcement with **Enforce rules**.
7. In the Group Policy Management Editor window, go to **Computer Configuration\Policies\Windows Settings\Security Settings\System Services**.
8. Configure **Application Identity Properties** with **Define this policy setting**, and **Select service startup mode** with **Automatic**.

Apply the GPO to the domain

1. In the Group Policy Management Console, apply the **WordPad Restriction Policy** GPO to the **Adatum.com** domain.
2. Open the Command Prompt window, type **gpupdate /force**, and then press Enter.

Test the AppLocker rule

1. Sign in to LON-CL1 as **Adatum\Alan** with the password **Pa\$\$w0rd**.
2. In the Command Prompt window, type **gpupdate /force**, and then press Enter.
Wait for the policy to update.
3. Attempt to start WordPad, and verify that WordPad does not start.

Lesson 4

Configuring Windows Firewall with Advanced Security

Windows Firewall with Advanced Security is an important tool for enhancing the security of Windows Server 2012. This snap-in helps to prevent several different security issues such as port scanning or malware. Windows Firewall with Advanced Security has multiple firewall profiles, each of which applies unique settings to different types of networks. You can manually configure Windows Firewall rules on each server, or configure them centrally by using Group Policy.

Lesson Objectives

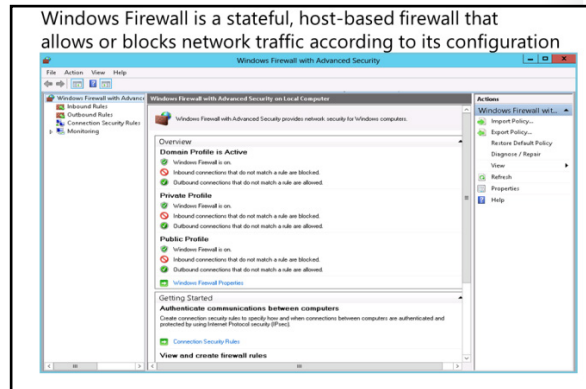
After completing this lesson, you will be able to:

- Describe the features of Windows Firewall with Advanced Security.
- Explain why a host-based firewall is important.
- Describe Firewall Profiles.
- Describe connection security rules.
- Explain how to deploy Windows Firewall rules.
- Secure network traffic by using Windows Firewall.

What Is Windows Firewall with Advanced Security?

Windows Firewall is a host-based firewall that is included in Windows Server 2012. This snap-in runs on the local computer and restricts network access to and from that computer.

Unlike a perimeter firewall, which provides protection only from threats on the Internet, a host-based firewall provides protection from threats wherever they originate. For example, Windows Firewall protects a host from a threat within the local area network (LAN).



Inbound and Outbound Rules

Inbound rules control communication that is initiated by another device or computer on the network, with the host computer. By default, all inbound communication is blocked except the traffic that is explicitly allowed by an inbound rule.

Outbound rules control communication that is initiated by the host computer, and is destined for a device or computer on the network. By default, all outbound communication is allowed except the traffic that is explicitly blocked by an outbound rule. If you choose to block all outbound communication except the traffic that is explicitly allowed, you must carefully catalog the software that is allowed to run on that computer and the network communication required by that software.

You can create inbound and outbound rules based on User Datagram Protocol (UDP) and TCP ports, as well as other protocols. You can also create inbound and outbound rules that allow a specific executable network access, regardless of the port number that is being used.

Connection Security Rules

You use Connection Security Rules to configure IPsec for Windows Server 2012. When these rules are configured, you can authenticate communication between computers, and then use that information to create firewall rules based on specific user and computer accounts.

Additional Configuration Options

Windows Firewall with Advanced Security is a Microsoft Management Consoles (MMC) snap-in that allows you to perform advanced configuration of Windows Firewall.

Windows Firewall in Windows 8 and Windows Server 2012 provides the following features:

- Supports filtering for both incoming and outgoing traffic.
- Integrates firewall filtering and IPsec protection settings.
- Enables you to configure rules to control network traffic.
- Provides network location-aware profiles.
- Enables you to import or export policies.

You can configure settings for Windows Firewall on each computer individually, or by accessing the following location from the GPMC:

- Computer Configuration\Policies\Windows Settings\Security Settings\Windows Firewall with Advanced Security



Note: Windows Server 2012 introduces the additional option for administering Windows Firewall by using the Windows PowerShell command-line interface.

Discussion: Why Is a Host-Based Firewall Important?

Review the discussion question and participate in a discussion to identify the benefits of using a host-based firewall such as Windows Firewall with Advanced Security.

Question: Why is it important to use a host-based firewall such as Windows Firewall with Advanced Security?

- Why is it important to use a host-based firewall such as Windows Firewall with Advanced Security?

10 minutes



Firewall Profiles

Windows Firewall with Advanced Security uses firewall profiles to provide a consistent configuration for networks of a specific type, and allows you to define a network as either a domain network, a public network, or a private network.

With Windows Firewall with Advanced Security, you can define a configuration set for each type of network; each configuration set is referred to as a *firewall profile*. Firewall rules are activated only for specific firewall profiles.

Windows Firewall with Advanced security includes the profiles in the following table.

- Firewall profiles are a set of configuration settings that apply to a particular network type
- The firewall profiles are:
 - Domain
 - Public
 - Private
- Windows Server 2012 includes the ability to have multiple active firewall profiles

Profile	Description
Public	Use when you are connected to an untrusted public network. Other than domain networks, all networks are categorized as Public. By default, the Public profile (which is the most restrictive) is used in Windows Vista, Windows 7, and Windows 8.
Private	Use when you are connected behind a firewall. A network is categorized as private only if an administrator or a program identifies the network as private. Networks marked as Home or Work in Windows Vista, Windows 7, and Windows 8 are added to the Private profile.
Domain	Use when your computer is part of a Windows operating system domain. Windows operating systems automatically identify networks on which it can authenticate access to the domain controller. The Domain profile is assigned to these networks, and this setting cannot be changed. No other networks can be placed in this category.

Windows Server 2012 allows multiple firewall profiles to be active on a server simultaneously. This means that a multi-homed server that is connected to both the internal network and the perimeter network can apply the domain firewall profile to the internal network, and the public or private firewall profile to the perimeter network.

Connection Security Rules

A connection security rule forces authentication between two peer computers before they can establish a connection and transmit secure information. They also secure that traffic by encrypting the data that is transmitted between computers. Windows Firewall with Advanced Security uses IPsec to enforce these rules.

- Connection security rules:
- Authenticate two computers before they begin communications
 - Secure information being sent between two computers
 - Use key exchange, authentication, data integrity, and data encryption (optionally)
- How firewall rules and connection rules are related:
- Firewall rules allow traffic through, but do not secure that traffic
 - Connection security rules can secure the traffic, but only if a firewall rule was previously configured

The configurable connection security rules are:

- **Isolation.** An isolation rule isolates computers by restricting connections that are based on credentials such as domain membership or health status. Isolation rules allow you to implement an isolation strategy for servers or domains.
- **Authentication Exemption.** You can use an authentication exemption to designate connections that do not require authentication. You can designate computers by a specific IP address, an IP address range, a subnet, or a predefined group such as a gateway.
- **Server-to-Server.** A server-to-server rule protects connections between specific computers. This type of rule usually protects connections between servers. When creating the rule, specify the network endpoints between which communications are protected. Then designate requirements and the authentication that you want to use.
- **Tunnel.** With a tunnel rule, you can protect connections between gateway computers. Typically, you use a tunnel rule when connecting across the Internet between two security gateways.
- **Custom.** Use a custom rule to authenticate connections between two endpoints when you cannot set up authentication rules that you need by using the other rules available in the new Connection Security Rule Wizard.

How Firewall Rules and Connection Security Rules Work Together

Firewall rules allow traffic through the firewall, but do not secure that traffic. To secure traffic with IPsec, you can create connection security rules. However, connection security rules do not allow traffic through a firewall. You must create a firewall rule to do this. Connection security rules are not applied to programs and services; instead, they are applied between the computers that make up the two endpoints.

Deploying Firewall Rules

How you deploy Windows Firewall rules is an important consideration. Choosing the appropriate method ensures that rules are deployed accurately and with minimum effort. You can deploy Windows Firewall rules in the following ways:

- **Manually.** You can individually configure firewall rules on each server. However, in an environment with more than a few servers, this is labor intensive and prone to error. This method is typically used only during testing and troubleshooting.
- **Using Group Policy.** The preferred way to distribute firewall rules is by using Group Policy. After creating and testing a GPO with the required firewall rules, you can deploy the firewall rules quickly and accurately to a large number of computers.
- **Exporting and importing firewall rules.** Windows Firewall with Advanced Security also gives you the option to import and export firewall rules. You can export firewall rules to create a backup before you manually configure firewall rules during troubleshooting.

You can deploy Windows Firewall rules in the following ways:

- **Manually.** Used during testing, troubleshooting, or for individual computers.
- **Using Group Policy.** The preferred way. Create and test the rules, and then deploy them to a large number of computers.
- **Exporting and importing.** Uses Windows Firewall with Advanced Security. When you import rules, they replace all current rules.

Always test firewall rules in an isolated, non-production environment before you deploy them in production.



Note: When you import firewall rules, they are treated as a complete set, and replace all currently-configured firewall rules.

Demonstration: Implementing secured network traffic with Windows Firewall

In this demonstration, you will see how to:

- Check to see if ICMP (Internet Control Message Protocol) v4 is blocked.
- Enable ICMP v4 from LON-CL2 to LON-SVR2.
- Create a connection security rule so that traffic is authenticated to the destination host.
- Validate ICMP v4 after the connection security rule is in place.

Demonstration Steps

Check to see if ICMP v4 is blocked

1. Sign in to LON-CL2 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On LON-CL2, ping **10.10.0.11** and notice that the ping times out.

Enable ICMP v4 from LON-CL2 to LON-SVR2

1. Sign in to LON-SVR2 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On LON-SVR2, create a firewall rule to allow ICMPv4 from LON-CL2.
3. On LON-CL2, ping **10.10.0.11**.

Notice that the ping goes through successfully.

Create a connection security rule

- On LON-SVR2, create an isolation-based connection security rule to authentication inbound traffic and request authentication for outbound traffic.

Validate ICMP v4

- On LON-CL2, ping **10.10.0.11**.

Notice that the ping goes through successfully.

Lab B: Configuring AppLocker and Windows Firewall

Scenario

Your manager has asked you to implement AppLocker to restrict non-standard apps from running. He also has asked you to create new Windows Firewall rules for any member servers running web-based apps.

Objectives

After completing this lab, you will be able to:

- Configure AppLocker Policies.
- Configure Windows Firewall.

Lab Setup

Estimated time: 60 minutes

Virtual machines	20410C-LON-DC1 20410C-LON-SVR1 20410C-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 **20410C-LON-DC1**, and in the Actions pane, click **Connect**.
Wait until the virtual machine starts.
3. If needed, sign in using the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
4. Repeat steps 2 and 3 for 20410C-LON-SVR1 and 20410C-LON-CL1.

Exercise 1: Configuring AppLocker Policies

Scenario

Your manager has asked you to configure new AppLocker policies to control the use of apps on user desktops. The new configuration should allow apps to be run only from approved locations. All users must be able to run apps from C:\Windows and C:\Program Files.

You also need to add an exception to run a custom-developed app that resides in a non-standard location.

The first stage of the implementation records from which locations apps are being run now. The second stage of implementation prevents unauthorized apps from running.

The main tasks for this exercise are as follows:

1. Create an OU for client computers.
2. Move LON-CL1 to the Client Computers OU.
3. Create a Software Control GPO and link it to the Client Computers OU.
4. Run GPUupdate.
5. Run app1.bat in the C:\CustomApp folder.
6. View AppLocker events in an event log.
7. Create a rule that allows software to run from a specific location.
8. Modify the Software Control GPO to enforce rules.
9. Verify that an application can still be run.
10. Verify that an app cannot be run.

► **Task 1: Create an OU for client computers**

1. Switch to LON-DC1.
2. Open **Active Directory Users and Computers**.
3. Create new OU called **Client Computers OU**.

► **Task 2: Move LON-CL1 to the Client Computers OU**

- On LON-DC1, in Active Directory Users and Computers, move LON-CL1 to Client Computers OU.

► **Task 3: Create a Software Control GPO and link it to the Client Computers OU**

1. On LON-DC1, open the Group Policy Management Console.
2. In the Group Policy Management Console, in the Group Policy Objects container, create a new GPO with the name **Software Control GPO**.
3. For the Software Control GPO, open the Group Policy Management Editor window.
4. In the Group Policy Management Editor window, go to **Computer Configuration/Policies/Windows Settings/Security Settings/Application Control Policies/AppLocker**.
5. Create default rules for the following:
 - **Executable Rules**
 - **Windows Installer Rules**
 - **Script Rules**
 - **Packaged app Rules**
6. Configure rule enforcement with the Audit only option for the following:
 - **Executable Rules**
 - **Windows Installer Rules**
 - **Script Rules**
 - **Packaged app Rules**
7. In the Group Policy Management Editor window, go to **Computer Configuration/Windows Settings/Security Settings**.
8. Click **System Services**, and then double-click **Application Identity**.
9. In the **Application Identity Properties** dialog box, click **Define this policy setting** and above **Select service startup mode**, click **Automatic**, and then click **OK**.

10. Close the Group Policy Management Editor window.
11. In the Group Policy Management Console, link the **Software Control GPO** to **Client Computers OU**.

► Task 4: Run GPUdate

1. Switch to LON-CL1.
2. Open the Command Prompt window, and type the following command:

```
gpupdate /force
```

3. Close the Command Prompt window, and restart LON-CL1.

► Task 5: Run app1.bat in the C:\CustomApp folder

1. Sign in to LON-CL1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. At the command prompt, type the following command and then press Enter:

```
gpresult /R
```

Review the result of the command and ensure that **Software Control GPO** is displayed under Computer Settings, Applied Group Policy Objects.

3. If **Software Control GPO** is not displayed, restart LON-CL1 and repeat steps 1 and 2.
4. At the command prompt, type the following command, and then press Enter:

```
C:\CustomApp\app1.bat
```

► Task 6: View AppLocker events in an event log

1. On LON-CL1, start **Event Viewer**.
2. In the Event Viewer window, browse to **Application and Services Logs/Microsoft/Windows/AppLocker**, and review the events.
3. Click **MSI and Scripts** and review event log 8005 that contains the following text:
%OSDRIVE%\CUSTOMAPP\APP1.BAT was allowed to run.



Note: If no events are displayed, ensure that the Application Identity service has started, and try again.

► Task 7: Create a rule that allows software to run from a specific location

1. On LON-DC1, edit the **Software Control GPO**.
2. In the Group Policy Management Editor window, go to **Computer Configuration/Policies/Windows Settings/Security Settings/Application Control Policies/AppLocker**.
3. Create a new script rule with the following configuration:
 - Permissions: **Allow**
 - Conditions: **Path**
 - Path: **%OSDRIVE%\CustomApp\app1.bat**
 - Name and Description: **Custom App Rule**

► Task 8: Modify the Software Control GPO to enforce rules

1. Use the **Enforce rules** option to configure rule enforcement for the following:
 - **Executable Rules**
 - **Windows Installer Rules**
 - **Script Rules**
 - **Packaged app Rules**
2. Close the Group Policy Management Editor window.

► Task 9: Verify that an application can still be run

1. Switch to LON-CL1.
2. Open the Command Prompt window, and type the following command:


```
gpupdate /force
```
3. Close the Command Prompt window and restart LON-CL1.
4. Sign in to LON-CL1 as **Adatum\Tony** with the password **Pa\$\$w0rd**.
5. Open the Command Prompt window and verify that you can run the **app1.bat** app, which is located in the C:\CustomApp folder.

► Task 10: Verify that an app cannot be run

1. On LON-CL1, from the CustomApp folder, copy **app1.bat** to the **Documents** folder.
2. Verify that app cannot be run from the Documents folder, and that the following message appears: "This program is blocked by Group Policy. For more information, contact your system administrator."

Results: After completing this exercise, you should have configured AppLocker policies for all users whose computer accounts are located in the Client Computers OU. The policies you configured should allow these users to run apps that are located in the folders C:\Windows and C:\Program Files, and run the custom-developed app app1.bat in the C:\CustomApp folder.

Exercise 2: Configuring Windows Firewall

Scenario

Your manager has asked you to configure Windows Firewall rules for a set of new app servers. These app servers have a web-based program that is listening on a non-standard port. You need to configure Windows Firewall to allow network communication through this port. You will use security filtering to ensure that the new Windows Firewall rules apply only to the app servers.

The main tasks for this exercise are as follows:

1. Create a group named Application Servers.
2. Add LON-SVR1 as a group member.
3. Create a new Application Servers GPO.
4. Link the Application Servers GPO to the Member Servers OU.
5. Use security filtering to limit the Application Server GPO to members of Application Server group.
6. Run GPUpdate on LON-SVR1.
7. View the firewall rules on LON-SVR1.

► **Task 1: Create a group named Application Servers**

- On LON-DC1, in Active Directory Users and Computers, in the **Member Servers OU**, create a new global security group named **Application Servers**.

► **Task 2: Add LON-SVR1 as a group member**

- In Active Directory Users and Computers, in the **Member Servers OU**, open **Application Servers Properties**, and then add LON-SVR1 as a group member.

► **Task 3: Create a new Application Servers GPO**

1. On LON-DC1, open the Group Policy Management Console.
2. In the Group Policy Management Console, in the Group Policy Objects container, create a new GPO with a name **Application Servers GPO**.
3. In the Group Policy Management Editor window, go to **Computer Configuration/Policies/Windows Settings/Security Settings/Windows Firewall with Advanced Security/Windows Firewall with Advanced Security - LDAP://{GUID}**.
4. Configure an inbound rule with the following settings:
 - Rule Type: **Custom**
 - Protocol type: **TCP**
 - Specific Ports: **8080**
 - Scope: **Any IP address**
 - Action: **Allow the connection**
 - Profile: **Domain** (clear both the **Private** and **Public** check boxes)
 - Name: **Application Server Department Firewall Rule**
5. Close the Group Policy Management Editor window.

► **Task 4: Link the Application Servers GPO to the Member Servers OU**

- In the Group Policy Management Console, link the **Application Servers GPO** to the **Member Servers OU**.

► **Task 5: Use security filtering to limit the Application Server GPO to members of Application Server group**

1. On LON-DC1, open the Group Policy Management Console.
2. Expand the **Member Servers OU**, and then click **Application Servers GPO**.
3. In the right-hand pane, under **Security Filtering**, remove **Authenticated Users**, and configure **Application Servers GPO** to apply only to the Application Servers security group.

► **Task 6: Run GPUpdate on LON-SVR1**

1. Switch to LON-SVR1.
2. Open the Command Prompt window, and type the following command:

```
gpupdate /force
```

3. Close the Command Prompt window.
4. Restart LON-SVR1, and then sign back in as **Adatum\Administrator** with the password of **Pa\$\$w0rd**.

► Task 7: View the firewall rules on LON-SVR1

1. Switch to LON-SVR1.
2. Start Windows Firewall with Advanced Security.
3. In the Windows Firewall with Advanced Security window, in Inbound rules, verify that the **Application Server Department Firewall Rule** that you created earlier using Group Policy is configured.
4. Verify that you cannot edit the **Application Server Department Firewall Rule**, because it is configured through Group Policy.

Results: After completing this exercise, you should have used Group Policy to configure Windows Firewall with Advanced Security to create rules for app servers.

► Prepare for the next module

When you finish the lab, revert the virtual machines to their initial state by performing the following steps:

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20410C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for 20410C-LON-SVR1 and 20410C-LON-CL1.

Lab Review Questions

Question: You configured an AppLocker rule based on a software path. How can you prevent users from moving the folder containing the software so that they can still run it?

Question: You would like to introduce a new app that requires the use of specific ports. What information do you need to configure Windows Firewall with Advanced Security, and from what source can you get it?

Module Review and Takeaways

Module Review Questions

Question: Does the defense-in-depth model prescribe specific technologies that you should use to protect Windows Server operating system servers?

Question: What setting must you configure to ensure that users are allowed only three invalid sign in attempts?

Question: You are creating a GPO with standardized firewall rules for the servers in your organization. You tested the rules on a standalone server in your test lab. The rules appear on the servers after the GPO is applied, but they are not taking effect. What is the most likely cause of this problem?

Question: Last year, your organization developed a security strategy that included all aspects of a defense-in-depth model. Based on that strategy, your organization implemented security settings and policies on the entire IT infrastructure environment. Yesterday, you read in an article that new security threats were detected on the Internet, but now you realize that your company strategy does not include a risk analysis and mitigation plan for those new threats. What should you do?

Best Practices

The following are best practices:

- Always make a detailed security risk assessment before planning which security features your organization should deploy.
- Create a separate GPO for security settings that apply to different type of users in your organization, because each department might have different security needs.
- Ensure that the security settings that you configure are reasonably easy to use so that employees accept them. Frequently, very strong security policies are too complex or difficult for employees to adopt.
- Always test security configurations that you plan to implement with a GPO in an isolated, non-production environment. Only deploy policies in your production environment after you complete this testing successfully.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
The user cannot sign in locally to a server.	
After configuring auditing, there are too many events logged in the Security Event Log in Event Viewer.	
Some users complain that their business apps can no longer access resources on the server.	

Tools

Tool	Use	Where to find it
Group Policy Management Console	A graphical tool that you use to create, edit, and apply GPOs	Server Manager/Tools
AppLocker	Applies security settings that control which apps are allowed to be run by users	Group Policy Management Editor snap-in
Windows Firewall with Advanced Security	A host-based firewall that is included as a feature in Windows Server 2008 and later	Server Manager/Tools if configured individually, or Group Policy Management Editor snap-in for deploying with Group Policy
Security Compliance Manager	Deploying security policies based on Microsoft Security Guide recommendations and industry best practices	Download from the Microsoft website at http://go.microsoft.com/fwlink/?LinkID=266746 .

Module 13

Implementing Server Virtualization with Hyper-V

Contents:

Module Overview	13-1
Lesson 1: Overview of Virtualization Technologies	13-2
Lesson 2: Implementing Hyper-V	13-8
Lesson 3: Managing Virtual Machine Storage	13-17
Lesson 4: Managing Virtual Networks	13-26
Lab: Implementing Server Virtualization with Hyper-V	13-32
Module Review and Takeaways	13-38

Module Overview

Server virtualization has only been a part of the Windows Server® operating system since the release of Windows Server 2008 and the introduction of the Hyper-V® role. By using Server virtualization, your organization can save money through server consolidation. However, to use Server Virtualization more efficiently, server administrators need to be able to decide which server workloads will run effectively in virtual machines, and which server workloads must remain deployed in a more traditional server environment.

This module introduces you to the Hyper-V role in Windows Server 2012 and Windows Server 2012 R2, the components of the role, how best to deploy the role, and the new features of the Hyper-V role that are introduced with Windows Server 2012 and Windows Server 2012 R2.

Objectives

After completing this module, you will be able to:

- Describe virtualization technologies.
- Implement Hyper-V.
- Manage virtual machine storage.
- Manage virtual networks.

Lesson 1

Overview of Virtualization Technologies

You can deploy many different types of virtualization technologies on networks where Windows® operating systems are deployed. The types of virtualization technologies that you select depend on what your organization needs to accomplish. Although this module is focused primarily on server virtualization, in this lesson you will learn about other types of virtualization technologies, and the situations in which it is appropriate to deploy them.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe server virtualization using Hyper-V.
- Describe Windows Azure™.
- Explain when you would use desktop virtualization.
- Determine the components required to implement presentation virtualization.
- Explain the advantages of using Microsoft application virtualization rather than traditional application deployment.

Server Virtualization

With server virtualization, you can create separate virtual machines and run them concurrently on a single server that is running Hyper-V. These virtual machines are known as *guests*. The computer that is running Hyper-V is known as the *virtualization server* or the *management operating system*.

Virtual machine guests function as normal computers. When users are signed into a guest virtual machine remotely using Remote Desktop Connection (RDC) or a Windows PowerShell® remote session, you would have to examine closely the properties of the computer on which the user is working to determine whether it is a virtual machine or a traditionally deployed physical machine. Virtual machines that are hosted on the same virtualization server are independent of one another. You can run multiple virtual machines that are using different operating systems on a virtualization server simultaneously, provided the virtualization server has enough resources.

Benefits of server virtualization with Hyper-V:

- Invisible to users
- Guest machines can use different operating systems
- More efficient use of hardware
- Service and application isolation
- Workload consolidation
- Simplifies server deployment by using:
 - Virtual machine templates
 - Virtual machine self-service portals

Implementing Virtual Machines to Maximize Hardware Usage

By implementing virtual machines, you use hardware more efficiently. In most cases, a service or application does not consume more than a fraction of the virtualization server's resources. This means that you can install multiple services and applications on the same virtualization server and then deploy them to multiple virtual machines. This ensures that the resources of that virtualization server are used more effectively. For example, assume that you have four separate services and applications, each of which consumes from 10 to 15 percent of a virtualization server's hardware resources. You can install these services and applications in virtual machines, and then place them on the same hardware where, on average, they consume a total of 40 to 60 percent of the virtualization server's hardware.

This is a simplified example. In real-world environments, you must make adequate preparations before co-locating virtual machines. You have to ensure that the hardware resource needs of all the virtual machines that are hosted on the virtualization server do not exceed the server's hardware resources.

Isolating Services and Applications

Keeping one particular service or application functioning reliably can be challenging and it becomes even more complicated when you deploy multiple services and applications on the same server. For example, you might need to deploy two separate operating systems at a branch office, but these operating systems conflict when running on the same computer. If you can afford only one server then you can solve this problem by running these applications within virtual machines on the same server.

Consolidating Servers

With server virtualization, you can consolidate servers that would otherwise need to run on separate hardware onto a single virtualization server. Because each virtual machine on a virtualization server is isolated from the other virtual machines on the same server, it is possible to deploy services and applications that are incompatible with one another on the same physical computer, provided that you host them within virtual machines. Examples of such services and applications include Microsoft® Exchange Server 2013, SQL Server® 2012, and Active Directory® Domain Services (AD DS). This means that an organization only needs to deploy one physical server in place of the three servers that they would have needed in the past.



Best Practice: We recommend that you do not deploy a Microsoft Exchange mailbox server on the same computer that contains a domain controller role. We also recommend that you do not deploy a SQL Server 2012 database engine instance on the same computer that contains the domain controller role. Instead, deploy each of these workloads on separate virtual machines and then run those virtual machines as guests on the same virtualization server; this is a supported configuration.

Simplifying Server Deployment

You can also use virtualization to simplify the process of server deployment:

- Virtual machine templates for common server configurations are included with products such as Microsoft System Center 2012 - Virtual Machine Manager (VMM). In these templates, many of the parameters have been preconfigured with common settings so you do not have to configure the setting of every parameter yourself.
- You can also create virtual machine self-service portals that enable end users to provision approved servers and applications automatically. This lessens the workload of the systems administration team. You create these virtual machine self-service portals with VMM and Microsoft System Center 2012 - Service Manager.

What Is Windows Azure?

Windows Azure is a cloud-based platform on which you can purchase capacity, either for virtual machines, for applications, or for services such as SQL Server databases on SQL Azure™. One of the advantages of using Windows Azure is that you pay only for the capacity you use, rather than paying a fixed rate. For example, rather than paying a monthly flat rate to rent a server on a rack at a hosting provider, you pay less when the server is less busy and you pay more when the server gets busier.

- Windows Azure is a cloud-based platform for hosting virtual machines and applications
- You pay only for the resources that you use
- You can increase and decrease capacity automatically and swiftly
- You can use Windows Azure to:
 - Host websites
 - Host production applications
 - Host virtual machines
 - Test proof-of-concept solutions

Cloud-based capacity is elastic, meaning it can grow or shrink quickly as required. For example, in a traditionally hosted solution, you might choose a specific server chassis, but then if your need for capacity or performance grows you have to switch to a bigger class of server hardware. All of this takes time and planning. Similarly, if your need for capacity or performance decreases, you need to decide whether migrating to a lower class of hardware is worth the cost, or if your organization should continue to pay for a class of hardware that you do not need right now—and may or may not need in the future. By using a hosting provider, capacity is scaled automatically and you do not have to spend the time or money that it takes to switch from one server to another.

Cloud-based virtual machines, applications, and services can be useful when you have to provide proof-of-concept solutions for proposed projects. Rather than purchase test hardware and then deploy a proof-of-concept solution to it, you can deploy a cloud-based virtual machine quickly, and then deploy the proof-of-concept solution to that. Then, once the proof-of-concept solution is validated, you can discard the virtual machine—or keep it, depending on operational concerns. In addition to being faster, this solution is less expensive than buying the hardware for the proof-of-concept solution, which you may opt to discard if the project is not approved.

Hosting Websites or Production Applications

On cloud-based platforms such as Windows Azure you can deploy applications without having to deploy the underlying server infrastructure. For example, if you need a database, instead of deploying both Windows Server 2012 and SQL Server 2012 and then deploying the specific database, you can rent the cloud-based database server and host the database there.

For a successful cloud-based strategy, you must determine which services and applications are more economical to deploy on a cloud-based platform, and which services and applications are more economical to host in a more traditional server environment on your own premises. Many factors that are unique to your organization are involved in making this determination, and a strategy that is best for one organization may not be appropriate for another.

Desktop Virtualization

Client Hyper-V

You can install the Hyper-V role on computers that are running the Windows 8 Pro, Windows 8 Enterprise Windows 8.1 Pro and Windows 8.1 Enterprise operating systems. This allows you to run virtual machine guests on client computers. Client Hyper-V, the Hyper-V feature in Windows 8 and Windows 8.1 operating systems, has slightly different processor requirements than Hyper-V on Windows Server 2012 or Windows Server 2012 R2. Specifically, with the Windows 8 and Windows 8.1 client operating systems, the computer must have an x64 platform that supports Second-Level Address Translation (SLAT), and have a minimum of 4 gigabytes (GB) of random access memory (RAM). This differs from Hyper-V on Windows Server 2012 and Windows Server 2012 R2 which does not require SLAT.

Desktop virtualization includes the following technologies:

- Client (Local) Hyper-V
- VDI

RemoteFX allow virtual machines to display rich graphics and video capabilities

RemoteFX requires:

- GPU that supports DirectX 9.0c or later
- CPU that supports SLAT

Client Hyper-V on Windows 8 and Windows 8.1

The Client Hyper-V role on Windows 8 and Windows 8.1 supports many of the features that are available with Hyper-V on Windows Server 2012, but does not support Windows Server 2012 features such as virtual machine migration. Client Hyper-V also does not support publishing applications that are installed on the virtual machine guest to the management operating system's **Start** menu. This was a feature of Windows XP Mode on Windows 7, which uses Windows Virtual PC. (Windows Virtual PC is the client virtualization feature available to some computers running specific editions of Windows 7.)

Client Hyper-V in Enterprise Environments

In enterprise environments, Client Hyper-V is often used for development purposes, or to allow specific users to run previous versions of the Windows operating system, thereby allowing them to access applications that are incompatible with Windows 8 or Windows 8.1

Virtual Desktop Infrastructure

In Virtual Desktop Infrastructure (VDI), client operating systems are hosted centrally as virtual machines, and clients connect to these virtual machines using client software such as RDC. You can configure a server to support VDI by selecting a Remote Desktop Services installation in the Add Roles and Features Wizard. When you configure a virtualization server to function as a VDI server, you can install the Remote Desktop Virtualization Host role feature in addition to the Hyper-V role.

VDI can simplify the management of client operating systems by:

- Ensuring that all client computers that are hosted on a single server are backed up regularly.
- Hosting the client virtual machines on a highly available virtualization server.
- Ensuring that users can still access their virtual machine using other RDC methods when a client computer fails.

You can also use VDI to implement a Bring Your Own Device (BYOD) policy. In this scenario, workers bring their own computer to the office and use RDC software to connect to the virtual machine to which they are assigned.

RemoteFX

RemoteFX is a technology that benefits VDI deployments by providing a set of enhancements to remote desktop connections. With RemoteFX, virtual machines can display rich graphics and video capabilities including media streaming. It also provides support for multi-touch. To use RemoteFX, the Hyper-V host must have at least one GPU that supports DirectX 9.0c or later, and a CPU that supports second-level address translation (SLAT). If multiple GPU are installed on the Hyper-V host, all GPUs must be identical.

Presentation Virtualization

Presentation virtualization differs from desktop virtualization in the following ways:

- In desktop virtualization, each user is assigned their own virtual machine that is running a client operating system. In presentation virtualization, users sign in and run separate sessions on a server or servers. For example, users Adam and Gavin might be signed in simultaneously to the same remote desktop server, yet be running different sessions using RDC.
- With desktop virtualization, the applications run within virtual machines. With presentation virtualization, the desktop and the applications run on the virtualization server.

Differences between desktop virtualization and presentation virtualization

Desktop virtualization:

- Users are assigned their own virtual machines that are running a client operating system
- The desktop and apps run within virtual machines

Presentation virtualization:

- Users sign in and run separate sessions on the server
- The desktop and apps run on the host server

Presentation virtualization technologies include:

- Remote Desktop Services
- Full Desktop with RDC
- Applications using RemoteApp
- Remote Access through RD Gateway

On networks that use Windows Server 2012, presentation virtualization is provided by the Remote Desktop Services server role. Clients can access presentation virtualization in the following ways:

- Full Desktop. Clients can use a remote desktop client such as RDC to access a full desktop session and run applications on the Windows Server 2012 virtualization server.
- RemoteApp programs. Rather than use a full desktop client such as RDC, the Windows Server feature RemoteApp makes it possible for applications that run on the Windows Server 2012 server to display on the client computer.
- Remote Desktop Web Access. Using Remote Desktop Web Access (RD Web Access), clients can access a website on a specially configured server, and then launch RemoteApp programs and Remote Desktop sessions from their browser.

Remote Desktop Gateway

Remote Desktop Gateway (RD Gateway) makes it possible for external clients to access Remote Desktop and RemoteApp without using virtual private network (VPN), or DirectAccess, a feature of the Windows 7 and Windows 8 operating systems. RD Gateway is a role service that you can install on a computer running Windows Server 2012. You deploy RD Gateway servers on perimeter networks. You then configure the RDC client with the address of RD Gateway servers. When you do this, the client checks to see if the target remote desktop server is on the organizational network. If the remote desktop server is on the network, the client makes a direct connection to it. If the remote desktop server is not on the network, the client routes the connection to it through the RD Gateway server.

What is Microsoft Application Virtualization?

With application virtualization, you do not install apps permanently on client computers. Instead, apps are deployed from a server to clients when users want to use the app. Microsoft Application Virtualization (App-V) uses the Microsoft Application Virtualization Desktop Client (which is installed on the client). App-V is available as part of the Microsoft Desktop Optimization Pack, and is not a native Windows Server 2012 role or feature.

Benefits of App-V

- Application isolation
- Incompatible applications can run on the same server
- Application streaming
- Application deployment is quicker
- Application portability
- Applications can follow users across multiple computers

UE-V

- Application and operating system settings follow users across multiple computers

App-V Features and Benefits

There are three main benefits of App-V:

application isolation application, application streaming, and application portability.

- **Application Isolation.** App-V isolates the application from the operating system and runs it in a separate virtual environment. This means that you can run applications that might be incompatible when run together on the same computer. For example, you can use App-V to deploy and run different versions of Microsoft Office Word simultaneously.
- **Application Streaming.** Another useful feature of App-V is application streaming. When an application is streamed, only those parts of the application that are being used are transmitted to the client computer. This speeds up application deployment because only part of the application must be transmitted across the network to the client computer.
- **Application Portability.** When App-V is deployed with Microsoft System Center 2012 Configuration Manager, users can use the same applications on multiple client computers, without requiring a traditional installation on those client computers. For example, a user can sign in to a colleague's computer and then have App-V stream an application to them so that they can use it on that computer. The application is not installed locally, and when the user signs out, the application is no longer available to other users on that computer.

User Experience Virtualization

Just as App-V allows users to access their applications from different client computers, Microsoft User Experience Virtualization (UE-V) allows users to have the same operating system and application settings on multiple devices that are running Windows 7 and Windows 8. For example, when a user configures a setting for an application that is delivered through App-V on one computer (such as configuring a custom tab on a ribbon in a Microsoft Office product), that setting is available automatically when that application is delivered through App-V to another computer.

Lesson 2

Implementing Hyper-V

Understanding how Hyper-V works and how virtual machines function is critical to deploying server virtualization effectively in a Windows Server 2012 network environment. This lesson discusses Hyper-V, and the hardware requirements for deploying Hyper-V on a computer running Windows Server 2012. This lesson also discusses the components of a virtual machine (with an emphasis on Dynamic Memory), and the benefits of virtual machine integration services. Finally, it discusses how to measure virtual machine resource use with Windows PowerShell cmdlets.

Lesson Objectives

After completing this lesson, you will be able to:

- Install the Hyper-V role onto a server.
- Describe the appropriate hardware for Hyper-V deployment.
- Describe virtual machine hardware components.
- Configure Dynamic Memory.
- Configure virtual machine integration services.
- Configure virtual machine start and stop actions.
- Perform Hyper-V resource metering tasks.
- Describe the new features of Hyper-V in Windows Server 2012 R2.

What is Hyper-V?

Hyper-V is the hardware virtualization role that is available in Windows Server 2012. Hardware virtualization provides a hypervisor layer that has direct access to the host server's hardware. The host operating system and all virtual machines running on the host access the hardware through the hypervisor layer. This is in contrast to software virtualization products such as Microsoft Virtual Server 2005 R2, that use the virtualization server's operating system to provide indirect access to the server's hardware.

You can deploy Hyper-V to a computer that is running Windows Server 2012 by using the Add Roles and Features Wizard. You can configure Windows Server 2012 as a virtualization server by using the Hyper-V role. Windows Server 2012 can then host virtual machine guests that are running supported operating systems. You can manage virtual machine administration locally through Windows PowerShell, or you can manage it remotely through the Hyper-V Manager console.

You can install the Hyper-V role on both the Server Core installation of Windows Server 2012, and Windows Server 2012 in the non-server core configuration. There is also a Microsoft Hyper-V Server 2012 edition, which includes only the components necessary to host virtual machines.

Hyper-V:

- Is the hardware virtualization role in Windows Server 2012
- Gives virtual machine guests direct access to the host's hardware
- Compatible Windows Server operating systems:
 - Windows Server 2012
 - Microsoft Hyper-V Server 2012



Note: In some documentation, the virtualization server (for example, the Windows Server 2012 computer that is running Hyper-V) is called the *parent partition*, and a virtual machine that is running on the server is called the *child partition*.

Hardware Requirements for Hyper-V

The server on which you plan to install the Hyper-V role must meet the following hardware requirements:

- The server must have an x64 platform that supports hardware assisted virtualization and Data Execution Prevention.
- The server must have enough CPU capacity to meet the requirements of the guest virtual machines.

A virtual machine hosted on Hyper-V in Windows Server 2012 can support up to 64 virtual processors.

- The server must have enough memory to support all of the virtual machines that must run concurrently, plus enough memory to run the host Windows Server 2012 operating system.
 - The server must have at least 4 GB of RAM.
 - A virtual machine hosted on Hyper-V in Windows Server 2012 can support a maximum of 1 terabyte (TB) of RAM.
- The storage subsystem performance must meet the input/output (I/O) needs of the guest virtual machines. Whether deployed locally or on storage area networks (SANs), you may have to place different virtual machines on separate physical disks, or you may have to deploy a high performance redundant array of independent disks (RAID), solid-state drives (SSD), hybrid-SSD, or a combination of all three.
- The virtualization server's network adapters must be able to support the network throughput needs of the guest virtual machines. You can improve network performance by installing multiple network adapters and using multiple Network Interface Cards (NICs).

Factors to consider when planning hardware for servers running Hyper-V:

- Processor characteristics
 - Must have an x64 platform that supports hardware assisted virtualization and Data Execution Protection
- Processing capacity
- Memory
- Storage subsystem performance
- Network throughput (typically multiple NICs)

Virtual Machine Hardware

Virtual machines use virtual (or *simulated*) hardware. The management operating system, Windows Server 2012 with Hyper-V, uses the virtual hardware to mediate access to actual hardware. For example, you can map a virtual network adapter to a virtual network that you map in turn to an actual network interface.

Virtual machines have the following simulated hardware by default:

- BIOS
- Memory
- Processor
- IDE Controller 0 and 1
- SCSI Controller
- Synthetic Network Adapter
- COM 1 and 2
- Diskette Drive

You can add the following hardware to a virtual machine:

- SCSI Controller (up to 4)
- Network Adapter
- Legacy Network Adapter
- Fibre Channel adapter
- RemoteFX 3D video adapter

Virtual machines have the following simulated hardware by default:

- BIOS. Simulates the computer's BIOS. On a standalone computer you can configure various BIOS-related parameters, and similarly, on a virtual machine you can configure some of the same parameters. These include:
 - The boot order for the virtual machine's virtual hardware.
 - From which device the virtual machine boots (for example, from a DVD drive, Integrated Drive Electronics (IDE), legacy network adapter, or floppy disk).
 - Whether Num Lock is enabled at boot.

- Memory. You can allocate up to 1 TB of memory resources to an individual virtual machine.
- Processor. You can allocate up to 64 virtual processors to a single virtual machine.
- IDE controller 0. A virtual machine can support only two IDE controllers and, by default, two are allocated to each virtual machine. Each IDE controller can support two devices.

You can connect virtual hard drives or virtual DVD drives to an IDE controller. You can use IDE controllers to connect virtual hard disks and DVD drives to virtual machines that use any operating system that does not support integration services.

- IDE controller 1. Allows additional virtual hard drives and DVD drives to be deployed to the virtual machine.
- SCSI controller. You can use a small computer system interface (SCSI) controller only on virtual machines that have 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 1. Allows you to configure a connection through a named pipe.
- COM 2. Allows you to configure an additional connection through a named pipe.
- Disk drive. Allows you to map a virtual floppy disk image to a virtual disk drive.

You can add the following hardware to a virtual machine by editing the virtual machine's properties and then clicking **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. You can use legacy network adapters with any operating systems that do not support integration services. You can also use legacy network adapters to deploy operating system images throughout the network. A single virtual machine can have up to four legacy network adapters.
- Fibre Channel adapter. If you add a Fibre Channel adapter to a virtual machine, the virtual machine can then connect directly to a Fibre Channel SAN. You can only add a Fibre Channel adapter to a virtual machine if the virtualization server has a Fibre Channel host bus adapter (HBA) that also has a Windows Server 2012 driver that supports virtual Fibre Channel.
- RemoteFX 3D video adapter. If you add a RemoteFX 3D video adapter to a virtual machine, the virtual machine can then display high performance graphics by leveraging Microsoft DirectX® and graphics processing power on the host Windows Server 2012 server.



Additional Reading: For more information about virtual Fibre channel adapters, see Hyper-V Virtual Fibre Channel Overview at <http://go.microsoft.com/fwlink/?LinkId=269712>.

What is Dynamic Memory?

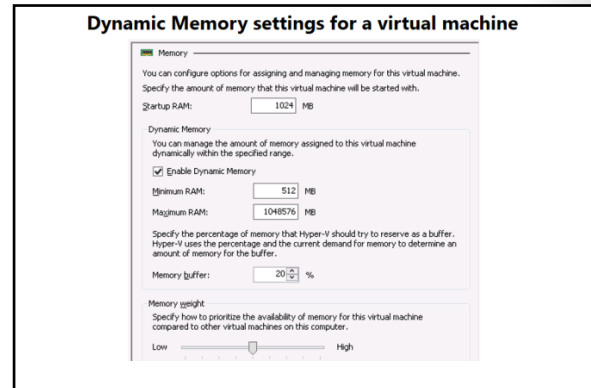
In the first release of Hyper-V with Windows Server 2008, you could only assign a static amount of memory to virtual machines. Unless you took special precautions to measure the precise amount of memory that a virtual machine required, you were likely to either under-allocate or over-allocate memory.

Dynamic Memory was introduced with Windows Server 2008 R2 Service Pack 1 (SP1). With Dynamic Memory you can:

- Allocate a minimum amount of memory to a virtual machine.
- Allow the virtual machine to request additional memory as the virtual machine needs it.
- Configure a maximum amount of memory to a virtual machine.

Thus, by using Dynamic Memory you no longer have to guess how much memory a virtual machine requires; instead, you can configure Hyper-V so that the virtual machine is allocated as much memory as it needs.


With Windows Server 2012, you can modify some of the Dynamic Memory minimum and maximum memory values while the virtual machine is running. This was not possible with Windows Server 2008 R2 SP1. You can perform this task from a Virtual Machine's **Settings** dialog box.



 **Note:** Virtual machines must support Hyper-V integration services to use Dynamic Memory.

Smart Paging

Virtual machines may need more memory during startup than they need during normal operation. Smart Paging, which is a new feature in Windows Server 2012, assigns additional temporary memory to a virtual machine when you restart the virtual machine. This means that you can allocate memory based on what the virtual machine needs when it is operating normally, rather than the amount that it needs during startup. Smart Paging uses disk paging to assign additional temporary memory to a virtual machine while it is restarting. However, using Smart Paging may result in lower performance because it uses disk resources that would otherwise be used by the host server and other virtual machines.

 **Note:** You can configure virtual machine memory using the **Set-VMemory** Windows PowerShell cmdlet.

 **Additional Reading:** For more information about Hyper-V Dynamic Memory, see Hyper-V Dynamic Memory Overview at <http://go.microsoft.com/fwlink/?LinkId=269713>.

Configuring Virtual Machine Integration Services

You must install Virtual Machine Integration Services if you want to use features such as operating system shutdown, time synchronization, and if you want to install virtual hardware components, such as SCSI adapters and synthetic network adapters, onto the virtual machines.

Virtual machine guest operating systems that are supported by Hyper-V and that can use Integration Services 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 2012
- Windows Small Business Server 2011
- Windows Server 2003 with SP2
- 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 SP1 or SP2
- SUSE Linux Enterprise Server 10 with Service Pack 4 (SP4)
- Windows 7 with SP1
- Windows Vista® with SP2
- Windows XP with Service Pack 3 (SP3)

Possible integration services:

- Operating system shutdown
- Time synchronization
- Data exchange
- Heartbeat
- Backup (volume snapshot)



Note: Support for the Windows XP operating system expires in April 2014. Support for Windows Server 2003 and Windows Server 2003 R2 expires in July 2015.

You can install the Hyper-V integration services components on an operating system by accessing the Virtual Machine Connection window, and then in the **Action** menu, clicking the **Insert Integration Services Setup Disk** item. You can then install the relevant operating system drivers either manually or automatically. You can also enable the following virtual machine integration components:

- Operating system shutdown. Allows the server running Hyper-V to initiate a graceful shutdown of the guest virtual machine.
- Time synchronization. Allows the virtual machine to use the virtualization server's processor for the purpose of time synchronization.
- Data exchange. Allows the server running Hyper-V to write data to the registry of the virtual machine.

- Heartbeat. Allows Hyper-V to determine if the virtual machine has become unresponsive.
- Backup (volume checkpoint). Allows the Volume Shadow Copy Service (VSS) provider to create checkpoints of the virtual machine for the purposes of backup operation, without interrupting the virtual machine's normal operations.

Configuring Virtual Machine Start and Stop Actions

You can use virtual machine start and stop actions to ensure that critical virtual machines always start automatically whenever a server running Hyper-V is restarted, and that they are shut down gracefully if the server receives a shutdown command. When you configure the virtual machine start and stop actions, you select the steps that the server running Hyper-V will perform on specific virtual machines when the physical server is started or shut down. You configure startup and shutdown settings for each virtual machine by editing the properties of the virtual machine.

Possible automatic start actions:

- Nothing
- Automatically start if it was running when the service stopped
- Always start this virtual machine automatically

Possible automatic stop actions:

- Save the virtual machine state
- Turn off the virtual machine
- Shut down the guest operating system

Automatic Start Options

You can configure the following options in the Automatic Start Actions window:

- Nothing. The virtual machine does not start automatically when the server running Hyper-V starts, even if the virtual machine was in a running state when the server was shut down.
- Automatically start if it was running when the service stopped. The virtual machine restarts if it was running when the server running Hyper-V received the command to shut down, or if the virtual machine was running when the server suffered a failure that caused it to be powered off.
- Always start this virtual machine automatically. The virtual machine always starts when the server running Hyper-V starts. You can configure a startup delay to ensure that multiple virtual machines do not attempt to all start up at the same time.

Automatic Stop Options

You can configure the following options in the Automatic Stop Actions window:

- Save the virtual machine state. This option saves the active state of the virtual machine to disk, including memory, when the server receives a shutdown command. This makes it possible for the virtual machine to restart when the server that is running Hyper-V restarts.
- Turn off the virtual machine. The virtual machine is turned off when the server receives a shutdown command. Data may be lost when this happens.
- Shut down the guest operating system. The virtual machine is shut down in a graceful manner when the server receives a shutdown command. This option is available only if integration services components are installed on the virtual machine.



Note: You can also configure virtual machine automatic start and stop actions by using the Windows PowerShell cmdlet **Set-VM** with the **AutomaticStartAction** and **AutomaticStopAction** parameters.

Hyper-V Resource Metering

Resource metering allows you to track the resource use of virtual machines that are hosted on Windows Server 2012 servers with the Hyper-V role installed.

With resource metering, you can measure the following parameters on individual Hyper-V virtual machines:

- Average CPU use.
- Average physical memory use, including:
 - Minimum memory use.
 - Maximum memory use.
- Maximum disk space allocation.
- Incoming network traffic for a network adapter.
- Outgoing network traffic for a network adapter.

Parameters that you can measure with resource metering:

- Average CPU use
- Average physical memory use, including:
 - Minimum memory use
 - Maximum memory use
- Maximum disk space allocation
- Incoming network traffic for a network adapter
- Outgoing network traffic for a network adapter

By measuring how much of these resources each virtual machine uses, an organization can bill departments or customers based on how much resources their virtual machines use, rather than charging a flat fee per virtual machine. An organization with only internal customers can also use these measurements to see patterns of use and plan future expansions. You perform resource metering tasks from a Windows PowerShell command-line interface using the following cmdlets:

- **Enable-VMResourceMetering.** Starts collecting data on a per virtual machine basis.
- **Disable-VMResourceMetering.** Disables resource metering on a per virtual machine basis.
- **Reset-VMResourceMetering.** Resets virtual machine resource metering counters.
- **Measure-VM.** Displays resource metering statistics for a specific virtual machine.



Note: There is no GUI tool that you can use to perform resource metering.



Additional Reading: For more information about resource metering for Hyper-V, see Hyper-V Resource Metering Overview at <http://go.microsoft.com/fwlink/?LinkId=269714>.

What's New with Hyper-V in Windows Server 2012 R2

The Hyper-V role in Windows Server 2012 R2 includes a large number of improvements and new features that were not available in Windows Server 2012.

New features in Windows Server 2012 R2 Hyper-V

The Hyper-V role in Windows Server 2012 R2 includes a large number of improvements and new features that were not available in Windows Server 2012. The following table lists new features in Windows Server 2012 R2 Hyper-V.

New or Improved	Feature
New to Windows Server 2012 R2	<ul style="list-style-type: none"> Shared virtual hard disk Automatic virtual machine activation Enhanced session mode Storage quality of service Virtual machine generation
Improved in Windows Server 2012 R2	<ul style="list-style-type: none"> Resize virtual hard disk Live migration Failover Clustering Integration services Export Replica Linux support Management

Feature	Description
Shared virtual hard disk	You can use this feature to cluster virtual machines by using shared virtual hard disk (VHDX format) files.
Automatic virtual machine activation	You can configure this feature to automatically activate virtual machines on computers running the Datacenter edition of Windows Server 2012 R2.
Enhanced session mode	You can use this feature to provide support for redirection of an increased number of local resources including audio, printers, clipboard, display configuration, smart cards, USB devices and supported Plug and Play devices
Storage quality of service	You can use this feature to specify maximum and minimum I/O loads in terms of I/O operations per second on a per virtual hard disk basis.
Virtual machine generation	You can use this feature to provide support for Generation 1 and Generation 2 virtual machines.

Improved features in Windows Server 2012 R2 Hyper-V

The following table lists improved features in Windows Server 2012 R2 Hyper-V.

Feature	Improvement
Resize virtual hard disk	This feature allows you to resize virtual hard disks while the virtual machine is running.
Live migration	This feature provides improved performance, including compression of virtual machine RAM and cross-version live migration between Windows Server 2012 and Windows Server 2012 R2 Hyper-V.
Failover Clustering	This feature provides virtual network adapter protection and virtual machine storage protection.
Integration Services	This feature provides the ability to copy files to a virtual machine without using a network connection or having to shut down the virtual machine.
Export	This feature allows you to export a virtual machine with all checkpoints or a single virtual machine checkpoint while the virtual machine is running.

(continued from previous page)

Feature	Improvement
Replica	This feature supports extended replication and configurable replication frequency.
Linux support	This feature provides support for Linux virtual machine backup and for VMs running Linux to support dynamic memory.
Management	This feature provides support for managing Hyper-V on Windows Server 2012 R2 from computers running Windows® 8 or Windows Server 2012.



Additional Reading: For more information, see "What's New in Hyper-V in Windows Server 2012 R2" at <http://go.microsoft.com/fwlink/?LinkID=331078>.

Lesson 3

Managing Virtual Machine Storage

Hyper-V provides many different virtual machine storage options. By knowing which option is appropriate for a given situation, you can help ensure that a virtual machine performs well. However, 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 virtualization server.

In this lesson, you will learn about different virtual hard disk types, different virtual hard disk formats, and the benefits and limitations of using virtual machine checkpoints.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the purpose of virtual hard disk.
- Describe SMB 3.0.
- Explain how to create a virtual hard disk type.
- Explain how to manage virtual hard disks.
- Explain how to deploy differencing virtual hard disks to reduce storage needs.
- Explain how to use virtual machine checkpoints.

What Is a Virtual Hard Disk?

A virtual hard disk is a file that represents a traditional hard disk drive that you can configure as a virtual hard disk with partitions and an operating system. You can use virtual hard disks on virtual machines, and you can mount virtual hard disks as local volumes using the Windows Server 2008 R2, Windows Server 2012, and Windows 8, and Windows 7 operating systems. Windows Server 2012 supports boot from virtual hard disk. This makes it possible for you to configure a computer to boot into a Windows Server 2012 operating system that is deployed on a virtual hard disk, or into certain editions of the Windows 8 operating system that are deployed on a virtual hard disk. You can create a virtual hard disk using:

- The Hyper-V Manager console.
- The Disk Management console.
- The DiskPart (diskpart.exe) command-line tool.
- The Windows PowerShell cmdlet **New-VHD**.

VHDX format has the following benefits over the VHD format:

- The disks can be larger (64 TB versus 2 TB)
- The disk is less likely to become corrupted
- The format supports better alignment when deployed to a large sector disk
- The format supports larger block size for dynamic and differencing disks
- Shared Virtual Hard Disks can be used by multiple virtual machines
- Storage QoS allows you to limit virtual hard disk IOPS



Note: Some editions of Windows 7 and Windows Server 2008 R2 also support boot from virtual hard disk.

Virtual hard disks in .vhd format vs. virtual hard disks in .vhdx format

Virtual hard disks traditionally use the .vhd extension. Windows Server 2012 introduces a new type of virtual hard disk which uses the .vhdx extension. Virtual hard disks with the .vhdx format have the following benefits over virtual hard disks that were used in Hyper-V on Windows Server 2008 and Windows Server 2008 R2:

- Virtual hard disks with the .vhdx format can be as large as 64 TB, whereas virtual hard disks with the .vhd format are limited to 2 TB.
- Virtual hard disks with the .vhdx format are less likely to be corrupted if the virtualization server suffers an unexpected power outage.
- The .vhdx format supports better alignment when deployed to a large sector disk.
- Virtual hard disks with the .vhdx format can hold larger dynamic and differencing virtual hard disks, which means that the dynamic and differencing virtual hard disks perform better.

You can convert a virtual hard disk with the .vhd format to the .vhdx format by using the Edit Virtual Hard Disk Wizard; you may want to do this if you have upgraded a Windows Server 2008 or Windows Server 2008 R2 virtualization server to Windows Server 2012 or Windows Server 2012 R2. You can also convert a virtual hard disk with the .vhdx format to the .vhd format.

SMB Share Support

Windows Server 2012 supports storing all virtual machine files, including virtual hard disks on SMB 3.0 file shares. This is an alternative to storing these files on Internet SCSI (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 you choose the virtual hard disk location or when you attach an existing virtual hard disk. The file share must support SMB 3.0. This means that you must place virtual hard disks on file shares that are hosted on file servers with Windows Server 2012. Older versions of Windows Server do not support SMB 3.0.



Additional Reading: For more information about virtual hard disk formats, see Hyper-V Virtual Hard Disk Format Overview at <http://go.microsoft.com/fwlink/?LinkId=269715>.

IDE vs. SCSI adapters

Virtual hard disks can be connected to virtual machines using two different virtual storage controller types: IDE or SCSI. When connected to an IDE controller, the virtual disk is accessed as an ATA device. When connected to a SCSI controller, the virtual disk is accessed as a SCSI device. The difference between the two is described in the following table.

IDE Controllers	SCSI Controllers
<ul style="list-style-type: none"> • Available only in Generation 1 virtual machines. • A virtual machine can have two IDE controllers. • Each IDE controller supports a maximum of two connected IDE devices (disks or virtual DVD drives). • You cannot add or remove devices from an IDE controller when a virtual machine is running. • Generation 1 virtual machines can only be booted locally off a device that is connected to an IDE controller. 	<ul style="list-style-type: none"> • Available in both Generation 1 and Generation 2 virtual machines. • A virtual machine can have up to 4 SCSI controllers. • Each SCSI controller supports up to 64 attached devices. • Can add or remove SCSI devices while a virtual machine is running. • Generation 2 virtual machines can only boot off a device that is attached to a SCSI controller.

Although there are differences in performance when you use an IDE or SCSI controller in a host virtual machine, these differences are not apparent when you use virtualized IDE or SCSI controllers.

Shared Virtual Hard Disks with Windows Server 2012 R2

Windows Server 2012 R2 allows you to configure shared virtual hard disks. A shared virtual hard disk is a virtual hard disk that is connected to multiple virtual machines. You can only use shared virtual hard disks for hard disks that are in .vhdx format and that are connected to virtual SCSI controllers. Shared virtual hard disk files are stored on failover clusters, either on a Cluster Shared Volume (CSV) on block storage, or on an SMB 3.0 scale-out file server.

QoS Management

Virtual hard disks in Windows Server 2012 R2 support the configuration of quality of service (QoS) parameters. When you configure the QoS parameters, you can specify the maximum number of input/output operations (IOPS) for the virtual disk, which minimizes the chance that a single virtual hard disk will consume the majority of the IOPS capacity of the underlying storage. You also can configure a virtual hard disk to trigger an alert if the number of IOPS falls below a threshold value. IOPS are measured in 8 KB increments. You cannot configure storage QoS when you are using shared virtual hard disks.



Additional Reading: For more information about virtual hard disk sharing, see <http://go.microsoft.com/fwlink/?LinkID=331079>.



Additional Reading: For more information about the storage quality of service for Hyper-V, see <http://go.microsoft.com/fwlink/?LinkID=331080>.

What is SMB 3.0

SMB 3.0 is a high-performance network file sharing protocol that became available with the release of Windows Server 2012. Previous versions of the Windows Server operating system used earlier versions of the SMB protocol for network file sharing. When you use Hyper-V in Windows Server 2012, you can store virtual machine configuration files and virtual hard disk files on SMB 3.0 shares.

The following table lists several features of SMB 3.0 that make it useful to support virtual machine storage.

SMB 3.0 can be used as shared storage to host highly available virtual machines

The following features of SMB 3.0 are useful when supporting highly-available virtual machines:

- SMB Transparent Failover
- SMB Scale Out
- SMB Multichannel
- SMB Direct (SMB over RDMA)
- SMB Encryption
- VSS for SMB File Shares

Feature	Function
SMB Transparent Failover	Enables clients that are accessing an SMB 3.0 share that is hosted on a failover cluster to continue to do so without interruption when the node hosting the file share fails.
SMB Scale Out	Provides simultaneous access to files by using direct I/O through all nodes that make up the file server cluster.
SMB Multichannel	Enables the share to use multiple network interfaces to aggregate bandwidth and provide fault tolerance.

(continued from previous page)

Feature	Function
SMB Direct	Enables shares to use network adapters that support Remote Direct Memory Access (RDMA).
SMB Encryption	Enables files that are transferred across the network to remain in an encrypted state without requiring a certificate services deployment.
VSS for SMB File Shares	Enables files that are hosted on SMB shares to be backed up in a consistent manner while they are open.



Additional Reading: For more information about updated links on Windows Server 2012 file server and SMB 3.0, see <http://go.microsoft.com/fwlink/?LinkID=331081>.


Creating Virtual Disk Types

When you configure a virtual hard disk, you can choose between several different disk types, including fixed, dynamic, and direct-attached storage.

Creating fixed virtual hard disks

When you create fixed virtual hard disks, all of the hard disk space that you have specified is allocated during the creation process. This minimizes fragmentation, which improves virtual hard disk performance if the disk is on a traditional (i.e., non-solid state) storage device. Allocating all of the specified hard disk space

during the creation process has a disadvantage though—in many situations, you do not know precisely how much disk space a virtual machine needs and you might allocate space that is not actually required.



Note: Disk fragmentation is less of an issue when virtual hard disks are hosted on RAID volumes, or on SSDs. Improvements in Hyper-V since it was first introduced with Windows Server 2008 also minimize the performance differences between dynamic and fixed virtual hard disks.

To create a fixed virtual hard disk, perform the following procedure:

1. Open the Hyper-V Manager console.
2. On 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. In the New Virtual Hard Disk Wizard, on the **Choose Disk Format** page, click either **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 in which to host the virtual hard disk file.

7. On the **Configure Disk** page, choose one of the following options:
 - **Create a new blank virtual hard disk of the specified size.**
 - **Copy the contents of a specified physical disk.** With this option you can replicate an existing physical disk on the server as a virtual hard disk. The fixed virtual hard disk will be the same size as the physical disk. Replicating an existing physical hard disk does not change the data on that disk.
 - **Copy the contents of a specified virtual hard disk.** With this option, you can create a new fixed hard disk based on the contents of an existing virtual hard disk.



Note: You can create a new fixed hard disk using the Windows PowerShell cmdlet **New-VHD** with the **-Fixed** parameter.

Dynamically expanding virtual hard disks

When you create a dynamically expanding VHD, you specify a maximum size, but the disk uses only the space that it needs and grows as necessary. A dynamically expanding virtual hard disk can be created with the .vhd format or the .vhdx format. A new dynamically expanding virtual hard disk with the .vhd format is allocated approximately 260 kilobytes (KB). A new dynamically expanding virtual hard disk with the .vhdx format is allocated approximately 4,096 KB.

As you save files to a dynamically expanding virtual hard disk, it grows. However, if you delete files from a dynamically expanding virtual hard disk, it does not shrink. You can only shrink a dynamically expanding virtual hard disk file by performing a compact operation.

To create a dynamically expanding virtual hard disk you follow the steps for creating a fixed virtual hard disk shown above, with the exception that, on the Choose Disk Type page (in step 5), you click **Dynamically Expanding** instead of **Fixed Size**.



Note: You can create a new dynamic hard disk using the Windows PowerShell cmdlet **New-VHD** with the **-Dynamic** parameter.

Direct-attached Storage

Direct-attached storage makes it possible for a virtual machine to access a physical disk drive. You can use direct-attached storage to connect a virtual machine directly to an iSCSI logical unit number (LUN). When you use direct-attached storage, the virtual machine must have exclusive access to the target disk. To ensure this, you must take the disk offline.

You can attach direct-attached storage by performing the following procedure:

1. Ensure that the target hard disk is offline. If it is not, then use the Disk Management console on the virtualization server to take it offline.
2. Use the Hyper-V Manager console to edit the existing virtual machine's properties.
3. Click an IDE or SCSI controller, click **Add**, and then click **Hard Drive**.
4. In the **Hard Drive** dialog box, click **Physical Hard Disk**. From the drop-down menu, select the disk that you want to use as direct-attached storage.



Note: If you connect direct-attached storage to a virtual machine's SCSI controller, then you do not have to shut down the virtual machine. If you want to connect to a virtual machine's IDE controller, then you must first shut down the virtual machine.

Question: Why might you consider using fixed virtual hard disks instead of dynamically expanding virtual hard disks?

Question: In what situations might you encounter difficulties if you use dynamically expanding disks?

Managing Virtual Hard Disks

From time to time, you need to perform maintenance operations on virtual hard disks. For example, you might want to convert a virtual hard disk to another format as your needs change, or you might want to compact a virtual hard disk to free up space. 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 format
- Convert a virtual hard disk in .vhdx format to .vhd format
- Compact a dynamically expanding virtual hard disk
- Expand a dynamically expanding virtual hard disk
- Expand a fixed virtual hard disk

The following are maintenance operations you can perform on virtual hard disks:

- Convert from fixed to dynamic
- Convert from dynamic to fixed
- Convert from VHD to VHDX format
- Convert from VHDX to VHD format
- Shrink a dynamic virtual hard disk
- Expand a dynamic or fixed virtual hard disk

Converting a Disk

When you convert a virtual hard disk, the contents of the existing virtual hard disk are copied to a newly-created virtual hard disk. For example, when you convert a fixed virtual hard disk to a dynamically expanding virtual hard disk, a new dynamic disk is created, the contents of the fixed disk are copied to the new dynamic disk, and then the fixed disk is deleted.

To convert a virtual hard disk from fixed to dynamic or from dynamic to fixed, perform the following steps:

1. In the Hyper-V Manager console, from the Actions pane, click **Edit Disk**.
2. In the Edit Virtual Hard Disk Wizard, on the **Before You Begin** page, click **Next**.
3. On the **Local Virtual Hard Disk** page, click **Browse**, and then select the virtual hard disk that you want to convert.
4. On the **Choose Action** page, click **Convert**, and then click **Next**.
5. On the **Convert Virtual Hard Disk** page, choose between the **VHD** and the **VHDX** formats.
6. On the **Convert Virtual Hard Disk** page, choose between **Fixed Size** and **Dynamically Expanding**. If you also want to convert the hard disk type, choose the appropriate type, and then click **Next**.
7. On the **Configure Disk** page, choose the destination location for the disk.

Changing the Size of a Disk

You can compact a dynamically expanding virtual hard disk that is not using all of the space that it is allocated. However, you cannot compact a fixed virtual hard disk without first converting it to a dynamically expanding virtual hard disk. You can expand both dynamically expanding virtual hard disks and fixed virtual hard disks.

You can use one of two methods to change the size of a virtual hard disk. They are:

- Use the Windows PowerShell cmdlets **resize-partition** and **resize-vhd**.
- In the Edit Virtual Hard Disk Wizard, select either the **Compact** or the **Expand** option.

Windows Server 2012 R2 is the first version in which you can resize a virtual hard disk while the virtual machine is still active.

Reducing Storage Needs with Differencing Virtual Hard Disks

Differencing virtual hard disks are separate virtual hard disks that record the changes made to a parent disk. You can use differencing virtual hard disks to reduce the amount of hard disk space consumed by virtual hard disks; this increases disk performance by reducing the space used. Differencing virtual hard disks work well with SSDs. They also work well where the available space on the parent volume is limited and the disk performance compensates for the performance drawbacks of using a differencing virtual hard disk.

When using differencing disks, you:

- Can reduce space that is used by storage, but at the cost of performance
- Can link multiple differencing disks to a single parent disk
- Cannot modify the parent disk
- Can use the Inspect Disk tool to reconnect a differencing disk to a missing parent

You can link multiple differencing virtual hard disks to a single parent disk. However, if you modify the parent disk, the links to all of the differencing virtual hard disks fail.

You can reconnect a differencing virtual hard disk to the parent using the Inspect Disk tool, which is available in the Actions pane of the Hyper-V Manager console. You can also use the Inspect Disk tool to locate the parent disk of a differencing virtual hard disk.

You can create a differencing virtual hard disk using the Hyper-V Manager console or using the Windows PowerShell cmdlet **New-VHD**.

To create a differencing virtual hard disk using the Hyper-V Manager console, perform the following steps:

1. Open the Hyper-V Manager console.
2. In the Actions pane, click **New**, and then click **Hard Disk**.
3. In the New Virtual Hard Disk Wizard, on the **Before You Begin** page, click **Next**.
4. On the **Choose Disk Format** page, click **VHD**, and then click **Next**.
5. On the **Choose Disk Type** page, click **Differencing**, and then click **Next**.
6. On the **Specify Name and Location** page, provide the location of the parent hard disk.

To create a differencing virtual hard disk using the **New-VHD** Windows PowerShell cmdlet, follow the pattern of the following example. To create a new differencing virtual hard disk named `c:\diff-disk.vhd` that uses the virtual hard disk `c:\parent.vhd`, use the following Windows PowerShell command:

```
New-VHD c:\diff-disk.vhd -ParentPath C:\parent.vhd
```

Using Checkpoints

A *checkpoint* is a static image of the data on a virtual machine at a given moment. Checkpoints are stored in either .avhd or .avhdx format depending on the virtual hard disk format. You can create a checkpoint of a virtual machine from the **Action** menu of the Virtual Machine Connection window, or from the Hyper-V Manager console. Each virtual machine can have a maximum of 50 checkpoints. Prior to Windows Server 2012 R2, checkpoints were known as "snapshots."

A *checkpoint* is a static image of the data on a virtual machine at a given moment
Checkpoints are not replacements for backups
You can perform a virtual machine export of a checkpoint
When you create a checkpoint, Hyper-V writes differencing virtual hard disk
When you apply a checkpoint, the virtual machine reverts to the configuration as it existed at the time the checkpoint was created

You can create checkpoints at any time, even when a virtual machine is shut down. When you create a checkpoint of a running virtual machine, the checkpoint includes the contents of the virtual machine's memory.

When taking checkpoints of multiple virtual machines that are part of the same group, for example a virtual domain controller and virtual member server, you should create these checkpoints simultaneously. This ensures that items such as computer account passwords are the same on all of the checkpoints.

Remember that when you revert to a checkpoint, you are reverting to a computer's state at that point in time. If you revert a virtual machine back to a point before it had performed a computer password change with a domain controller, you need to rejoin that computer to the domain or run the **netdom resetpwd** command.

Checkpoints vs. Backups

Checkpoints are not a replacement for backups. Checkpoints are stored on the same volume as the virtual hard disks. If that volume fails, both the checkpoints and the virtual hard disk file are lost.

Exporting Checkpoints

You can perform a virtual machine export of a checkpoint. When you do this, Hyper-V creates full virtual hard disks that represent the state of the virtual machine at the time the checkpoint was created. If you choose to export an entire virtual machine, all checkpoints associated with the virtual machine are also exported.

Differencing virtual hard disk Files

When you create a checkpoint, Hyper-V writes differencing virtual hard disk (.avhd, or .avhdx) files, which store the data that differentiates the checkpoint from the previous checkpoint, or from the parent virtual hard disk. When you delete checkpoints, this data is either discarded or merged back into the previous checkpoint or parent virtual hard disk. For example:

- If you delete the most recent checkpoint, the data is discarded. With Hyper-V in Windows Server 2012, this space is reclaimed immediately rather than when the virtual machine is shut down.
- If you delete the second-most recent checkpoint, the data is merged so that the earlier and latter checkpoint states of the virtual machine retain their integrity.

Managing Checkpoints

When you apply a checkpoint, the virtual machine reverts to the configuration as it existed at the time the checkpoint was created. Reverting to a checkpoint does not delete existing checkpoints. If you apply a checkpoint after making a configuration change in a different checkpoint, you are prompted to create another checkpoint. It is only necessary to create a new checkpoint if you want to return to that current configuration.

It is possible to create checkpoint trees that have different branches. For example, consider the following scenario:

You create a checkpoint of a virtual machine on Monday, Tuesday, and Wednesday. On Thursday you apply the Tuesday checkpoint. Immediately after applying the Tuesday checkpoint, you make changes to the virtual machine's configuration.

In this scenario, the original branch is the series of checkpoints created on Monday, Tuesday, and Wednesday. You create a new branch by applying the Tuesday checkpoint and making changes to the virtual machine. You can have multiple branches as long as you do not exceed the 50 checkpoints per virtual machine limit.

Lesson 4

Managing Virtual Networks

Hyper-V provides several different options for network communication between virtual machines. You can configure virtual machines that communicate with an external network in a manner similar to traditionally-deployed physical hosts. You can configure virtual machines to communicate only with a limited number of other virtual machines that are hosted on the same server. Knowing the options available for Hyper-V virtual networks ensures that you can use those options to best meet your organization's needs.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe virtual switches.
- Describe Virtual Local Area Networks.
- Describe Virtual Switch Extensions.
- Explain how to manage a virtual machine MAC address pool.
- Explain how to configure virtual network adapters.
- Describe advanced features of virtual network adapters.
- Describe NIC teaming.

What Is a Virtual Switch?

A *virtual switch* is a virtual version of a network switch. (The term "virtual network," which was used in Windows Server 2008, has been replaced by the term "virtual switch" in Windows Server 2012.) Virtual switches control how network traffic flows between multiple virtual machines that are hosted on the virtualization server, and between virtual machines and the rest of the organizational network. You can manage virtual switches through the Virtual Switch Manager, which is accessible through the Actions pane of the Hyper-V Manager console.

Hyper-V on Windows Server 2012 supports the following three types of virtual switches:

- | | |
|-----------------|---|
| External | Used to map a network to a specific network adapter or network adapter team |
| Internal | Used to communicate between the virtual machines on the host and between the virtual machines and the host itself |
| Private | Used to communicate between virtual machines, but not between the virtual machines and the host itself |

VLAN IDs are used to extend VLANs within the host's network switch to VLANs on the external network

Hyper-V on Windows Server 2012 supports the following three different types of virtual switches:

- **External.** This type of switch maps 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 local area network (LAN) service on the virtualization server, and if the virtualization server has a compatible adapter.
- **Internal.** Internal virtual switches communicate between multiple virtual machines on the virtualization server, and between the virtual machines and the virtualization server.
- **Private.** Private switches communicate only between multiple virtual machines on the virtualization server. You cannot use private switches to communicate between the virtual machines and the virtualization server.

Configuring VLAN IDs

When configuring a virtual network, you can also configure a virtual LAN (VLAN) ID that is associated with the network. By doing this, you can extend existing VLANs on the external network to VLANs within the virtualization server's network switch. VLANs enable you to partition network traffic, and they function as separate logical networks. Note that traffic can only pass from one VLAN to another if it passes through a router.

You can configure the following extensions for each virtual switch type:

- Microsoft NDIS Capture. This extension allows data that is traversing across the virtual switch to be captured.
- Microsoft Windows Filtering Platform. This extension allows data that is traversing across the virtual switch to be filtered.



Additional Reading: For more information about virtual switches, see Hyper-V Virtual Switch Overview at <http://go.microsoft.com/fwlink/?LinkId=269716>.

What are Virtual Local Area Networks

Virtual Local Area Networks (VLANs) allow you to logically segment network traffic that is running on the same physical network. VLANs function as separate broadcast domains. This means that broadcast traffic on one VLAN is not intercepted and processed by hosts on another VLAN, even when those hosts are connected to the same hardware switch. Each VLAN has an ID that is encapsulated within an Ethernet frame.

You can assign VLAN IDs to Hyper-V virtual switches and network adapters. To use VLANs with Hyper-V virtual switches and network adapters, the host's physical network adapters must support VLAN tagging. This feature must be enabled on the network adapter. When configuring VLAN IDs for Hyper-V, configure the VLAN ID on the virtual switch or on each individual virtual machine's network adapter. You do not need to configure the VLAN ID on the host's physical network adapter.

When you use Virtual Local Area Networks, you can:

- Logically segment network traffic running on the same physical and virtual networks
- Configure tagging on each virtual switch
- Configure tagging on each virtual network adapter

Note that host NICs must support VLAN tagging

Hyper-V supports the 802.1q specification for VLAN trunking. VLAN IDs that are used with virtual switches and virtual network adapters can also be used with networking equipment that supports these standards. This means that you can have the same VLAN span multiple Hyper-V hosts when connected to compatible network equipment.

You might implement VLANs with Hyper-V switches and virtual network adapters to support the following scenarios:

- Isolate network storage traffic. You can isolate network storage traffic such as iSCSI traffic from other traffic. Using VLANs means that a separate storage network might not be required.
- Isolate cluster traffic. You can isolate intra-node cluster traffic from other traffic.
- Security isolation. You can isolate hosts from each other for security reasons. For example, to make some virtual hosts available to Network Access Protection Clients that have been placed on an isolated VLAN so that they are able to remediate their configuration to a healthy state.


Virtual Switch Extensions

Virtual Switch (vSwitch) Extensions enable third-party vendors to create virtual switches that you can add to Hyper-V that add monitoring, filtering, and forwarding functionality. Some vendors have created virtual editions of their hardware switches. These virtual switches can be managed using the same management suite that is used to manage the physical switches. This enables an organization's networking team to extend switch management across both the virtual and the physical infrastructure.

The virtual switch extensions are listed in the following table.

Extension	Purpose
Network packet inspection	Examine network packets while they traverse the virtual switch.
Network packet filter	Create, filter, and modify packets that traverse the virtual switch.
Network forwarding	Create a forwarding extension for each virtual switch.
Intrusion detection or firewall	Filter and modify TCP/IP packets, monitor or authorize connections, filter IPsec traffic, and filter remote procedure calls.

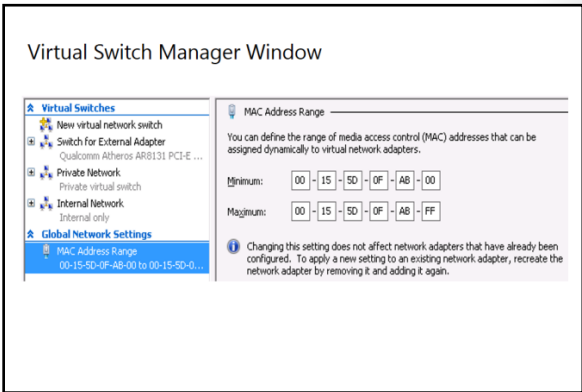
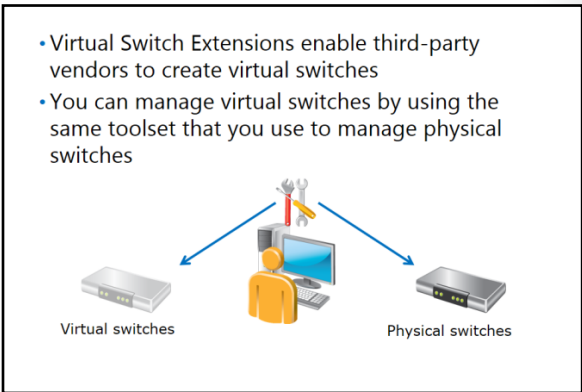
Consult third-party vendor catalogs to determine which virtual switches are available to run on the Hyper-V platform.

 **Additional Reading:** For more information about virtual switch extensions, see "Hyper-V Virtual Switch Overview" at <http://go.microsoft.com/fwlink/?LinkID=331084>.

Managing Virtual Machine MAC Addresses

Unless you specify a static media access control (MAC) address, Hyper-V dynamically allocates an address to each virtual machine network adapter from a pool of MAC addresses. You can configure the address range of this pool from the **MAC Address Range** setting of the Virtual Switch Manager console. By default, a server running Hyper-V has a pool of 255 MAC addresses.

When virtual machines use private or internal networks, the MAC address that is allocated to network adapters is not likely to be of concern, because the server running Hyper-V ensures that duplicate MAC addresses are not assigned to different virtual machines. However, when you have multiple servers that are running Hyper-V and are hosting virtual machines that use adapters connected to



external networks, you should ensure that each server uses a different pool of MAC addresses. This ensures that separate servers that connect to the same network do not assign the same MAC addresses to the virtual machines that they host.

When virtual machines are allocated IP addresses through a Dynamic Host Configuration Protocol (DHCP) reservation, you should consider using static MAC addresses. A DHCP reservation ensures that a particular IP address is always allocated to a specific MAC address.

You can configure the MAC address range by performing the following steps:

1. Open the Hyper-V Manager console.
2. Select the Hyper-V host that you wish to configure.
3. On the Actions pane, click **Virtual Switch Manager**.
4. Under **Global Network Settings**, click **MAC Address Range**.
5. Specify a minimum and a maximum range for the MAC address.

MAC addresses are in hexadecimal format. When configuring ranges for multiple Hyper-V hosts, you should consider changing the values of the second from the last pair of digits. The following table displays examples of ranges for multiple Hyper-V hosts.

Hyper-V host	MAC address range
Host 1	Minimum: 00-15-5D-0F-AB-00 Maximum: 00-15-5D-0F-AB-FF
Host 2	Minimum: 00-15-5D-0F-AC-00 Maximum: 00-15-5D-0F-AC-FF
Host 3	Minimum: 00-15-5D-0F-AD-00 Maximum: 00-15-5D-0F-AD-FF

Configuring Virtual Network Adapters

Virtual network adapters allow the virtual machine to communicate using the virtual switches that you configure in the Virtual Switch Manager console. You can edit the properties of a virtual machine to modify the properties of a network adapter. From the Network Adapter pane on the virtual machine's **Settings** dialog box, you can configure the following:

- Virtual Switch. You configure to which virtual switch the network adapter connects.
- VLAN ID. You specify a VLAN ID that the virtual machine uses for communication that passes through this adapter.
- Bandwidth Management. You allocate a minimum and a maximum bandwidth for the adapter. The minimum bandwidth allocation is reserved by Hyper-V for the network adapter, even when virtual network adapters on other virtual machines are working at capacity.

- Properties of a network adapter:
 - Virtual Switch
 - VLAN ID
 - Bandwidth Management
- Features of a virtual network adapter:
 - MAC address allocation
 - DHCP Guard
 - Router Guard
 - Port Mirroring
 - NIC Teaming

Both synthetic network adapters and legacy network adapters support the following advanced features:

- **MAC address allocation.** You can configure a MAC address to be assigned from the MAC address pool, or you can configure the network adapter to use a fixed MAC address. You can also configure MAC address spoofing. This is useful when the virtual machine needs to provide specific network access, such as when the virtual machine is running a mobile device emulator that requires network access.
- **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 server running Hyper-V that hosts virtual machines for others, but does not have direct control over the configuration of those virtual machines.
- **Router Guard.** This feature drops router advertisement and redirection messages from virtual machines that are 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.** This feature allows you to copy incoming and outgoing packets from a network adapter to another virtual machine that you have configured for monitoring.
- **NIC Teaming.** This feature allows you to add the virtual network adapter to an existing team on the server running Hyper-V.

Legacy network adapters emulate common network adapter hardware. You use legacy network adapters in the following situations:

- You want to support network boot installation scenarios for virtual machines. For example, you want to deploy an operating system image from a Windows Deployment Services (Windows DS) server or through Configuration Manager.
- You need to support operating systems that do not support integration services and do not have a driver for the synthetic network adapter.

Legacy network adapters do not support the hardware acceleration features that synthetic network adapters support. You cannot configure a virtual machine queue, IPsec task offloading, or single root I/O virtualization (SR-IOV) for legacy network adapters. These advanced features are covered in the next topic.

Network Adapter Advanced Features

In addition to the features described earlier, synthetic network adapters support the following advanced features:

- **Virtual Machine Queue.** This feature 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 management operating system to the virtual machine. Virtual Machine Queue requires that the host computer has a network adapter that supports this feature.
- **IPsec task offloading.** This feature enables the host's network adapter to perform calculation-intensive security association tasks. In the event that sufficient hardware resources are not available, the guest operating system performs these tasks. You can configure a maximum number of offloaded security

- Virtual Machine Queue delivers network traffic directly to the guest
- IPsec task offloading enables the host's network adapter to perform calculation-intensive security association tasks
- SR-IOV enables multiple virtual machines to share the same Peripheral Component Interconnect (PCI) Express physical hardware resources
- vRSS balances the network processing across multiple virtual processor cores in a virtual machine

associations between 1 and 4,096. IP security (IPsec) task offloading requires guest operating system support and network adapter support.

- SR-IOV. Single-root I/O virtualization (SR-IOV) enables multiple virtual machines to share the same Peripheral Component Interconnect (PCI) Express physical hardware resources. If sufficient resources are not available, then network connectivity falls back and is provided through the virtual switch. SR-IOV requires specific hardware and special drivers to be installed on the guest operating system, and might also need to be enabled in the computer BIOS.
- Virtual Receive Side Scaling (vRSS). vRSS enables network adapters to balance network processing load across the processor cores assigned to a virtual machine. vRSS enables a virtual machine to process higher amounts of network traffic than it could process if only a single CPU core was responsible for processing traffic. You can implement vRSS by allocating a virtual machine multiple cores through the advanced network. To use vRSS, the host's processor must support Receive Side Scaling (RSS) and the host's network adapters must support Virtual Machine Queue (VMQ).

What is NIC teaming?

NIC teaming allows you to combine up to 32 network adapters and then use them as a single network interface. NIC teaming provides redundancy, allowing network communication to occur over the combined network interface even when one or more of the network adapters fail. The combination of network adapters also increases the bandwidth available to the combined network interface. NIC teaming is a feature available in the Windows Server 2012 operating system that can be used by both the Hyper-V host and Hyper-V virtual machines.

NIC Teaming:

- Provides redundancy and aggregates bandwidth
- Is supported at the host and virtual machine level

NIC Teaming in virtual machines:

- Requires multiple virtual network adapters
- Must be enabled on virtual network adapters
- Can then be implemented in virtual machine OS (if supported)

When used with virtual machines, NIC Teaming allows virtual machines to team virtual network adapters that connect to separate virtual switches.

To get the benefit of NIC Teaming, the host must have at least two external virtual switches. When you have multiple virtual network adapters attached to the same switch, if the physical network adapter that the virtual switch is connected to fails, those virtual network adapters will lose connectivity. When configuring NIC teaming for virtual machines, network adapters connected to virtual switches can use SR-IOV.

Enable virtual machine NIC Teaming for virtual machines on the Advanced Features page of the virtual network adapter in Hyper-V manager. You can also enable NIC Teaming for virtual machines by using the **Set-VMNetworkAdapter** Windows PowerShell cmdlet. To enable NIC Teaming within the virtual machine operating system, NIC Teaming must be enabled on the virtual network adapter, or, the virtual network adapter must be configured to allow MAC address spoofing. Once you have enabled virtual NIC Teaming on the virtual network adapter (or enabled MAC address spoofing), you can configure NIC teaming within the virtual machine.

A new feature of Windows Server 2012 R2 is dynamic NIC Teaming. In Windows Server 2012, new traffic is assigned to a particular NIC and the traffic flow remains with that NIC throughout the session. In dynamic NIC Teaming, traffic flow is balanced across all available NICs in a team.

Lab: Implementing Server Virtualization with Hyper-V

Scenario

Your assignment is to configure the infrastructure service for a new branch office.

To use the server hardware that is currently available at branch offices more effectively, your manager has decided that all branch office servers will run as virtual machines. You must now configure a virtual network and a new virtual machine for these branch offices.

Objectives

After performing this lab, you will be able to:

- Install the Hyper-V role onto a server.
- Configure virtual networking.
- Create and configure a virtual machine.
- Use virtual machine checkpoints.

Lab Setup

Estimated Time: 75 minutes

Virtual machine	20410C-LON-HOST1
User name	Adatum\Administrator
Password	Pa\$\$w0rd

Before beginning the lab, you must complete the following steps:

1. Reboot the classroom computer and from the Windows Boot Manager, select **20410C-LON-HOST1**.
2. Sign in to **LON-HOST1** with the **Administrator** account and the password **Pa\$\$w0rd**.

Exercise 1: Installing the Hyper-V Role onto a Server

Scenario

The first step in migrating to a virtualized environment for the branch office is installing the Hyper-V role on a new Windows Server 2012 server.

The main tasks for this exercise are as follows:

1. Install the Hyper-V role onto a server.
2. Complete the Hyper-V role installation and verify the settings.

► Task 1: Install the Hyper-V role onto a server

1. In Server Manager, click **Local Server**, and then configure the following network settings:
 - IP Address: **172.16.0.31**
 - Subnet mask: **255.255.0.0**
 - Default gateway: **172.16.0.1**
 - Preferred DNS server: **172.16.0.10**

2. Use the Add Roles and Features Wizard to add the Hyper-V role to LON-HOST1 with the following options:
 - **Do not create a virtual switch.**
 - **Use the Default stores locations.**
 - **Allow the server to restart automatically if required.**
3. After a few minutes, the server restarts automatically. Ensure that you restart the machine from the boot menu as **20410C-LON-HOST1**. The computer will restart several times.

► **Task 2: Complete the Hyper-V role installation and verify the settings**

1. Sign in to **LON-HOST1** using the account **Administrator** with the password **Pa\$\$word**.
2. When the installation of the Hyper-V tools completes, click **Close**.
3. Open the Hyper-V Manager console, and then click **LON-HOST1**.
4. Edit the Hyper-V settings of LON-HOST1, and configure the following settings:
 - Keyboard: **Use on the virtual machine**
 - Virtual Hard Disks: **C:\Users\Public\Documents\Hyper-V\Virtual Hard Disks**

Results: After completing this exercise, you should have installed the Hyper-V role onto a physical server.

Exercise 2: Configuring Virtual Networking

Scenario

After installing the Hyper-V role on the new server, you need to configure the virtual network. You need to create both a network that is connected to the physical network, and a private network that can be used only for communication between virtual machines. The private network will be used when the virtual machines are configured for high availability. You also need to configure a specific range of 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.
4. Configure the Media Access Control (MAC) address range.

► **Task 1: Configure the external network**

1. Open the Hyper-V Manager console, and then click **LON-HOST1**.
2. Use the **Virtual Switch Manager** to create a new External virtual network switch with the following properties:
 - Name: **Switch for External Adapter**
 - External Network: Mapped to the host computer's physical network adapter. (This varies depending on the host computer.)

► Task 2: Create a private network

1. On LON-HOST1, open the Hyper-V Manager console.
2. 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

1. On LON-HOST1, open the Hyper-V Manager console.
2. Use the Virtual Switch Manager to create a new virtual switch with the following properties.
 - Name: **Internal Network**
 - Connection type: **Internal network**

► Task 4: Configure the Media Access Control (MAC) address range

1. On LON-HOST1, open the Hyper-V Manager console.
2. Use the Virtual Switch Manager to configure the following MAC Address Range settings:
 - Minimum: **00-15-5D-0F-AB-A0**
 - Maximum: **00-15-5D-0F-AB-EF**

Results: After completing this exercise, you should have configured virtual switch options on a physically-deployed Windows Server 2012 server running the Hyper-V role.

Exercise 3: Creating and Configuring a Virtual Machine

Scenario

You have been asked to deploy two virtual machines to LON-HOST1. You have copied a sysprepped virtual hard disk file that hosts a Windows Server 2012 installation.

To minimize disk space use at the cost of performance, you are going to create two differencing virtual hard disk files based on the sysprepped virtual hard disk. You will then use these differencing virtual hard disk files as the virtual hard disk files for the new virtual machines.

The main tasks for this exercise are as follows:

1. Create differencing virtual hard disks.
2. Create virtual machines.
3. Enable resource metering.

► Task 1: Create differencing virtual hard disks

1. Use File Explorer to create the following two folders:
 - **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\ Base14A-WS12R2.vhd**
3. Open Windows PowerShell and execute the following command:

```
New-VHD "E:\Program Files\Microsoft Learning\Base\LON-GUEST2\LON-GUEST2.vhd"
-ParentPath "E:\Program Files\Microsoft Learning\Base\ Base14A-WS12R2.vhd"
```

4. Inspect the disk at **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\ Base14A-WS12R2.vhd** as a parent.

► Task 2: Create virtual machines

1. On LON-HOST1, in the Hyper-V Manager console, in the Actions pane, click **New**, and then click **Virtual Machine**.
2. Create a virtual machine with the following properties:
 - Name: **LON-GUEST1**
 - Location: **E:\Program Files\Microsoft Learning\Base\LON-GUEST1**
 - Generation: **Generation 1**
 - 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**
3. Open Windows PowerShell and execute 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"
```

4. Use the Hyper-V Manager console to edit the settings of LON-GUEST2 by configuring the following:
 - Automatic Start Action: **Nothing**
 - Automatic Stop Action: **Shut down the guest operating system**

► Task 3: Enable resource metering

- At the Windows PowerShell prompt, enter the following commands:

```
Enable-VMResourceMetering LON-GUEST1
Enable-VMResourceMetering LON-GUEST2
```

Results: After completing this exercise, you should have deployed two separate virtual machines using a sysprepped virtual hard disk file as a parent disk for two differencing virtual hard disks.

Exercise 4: Using Virtual Machine Checkpoints

Scenario

You are in the process of developing a strategy to mitigate the impact of incorrectly applied change requests. As a part of this strategy development, you are testing the speed and functionality of virtual machine checkpoints to roll back to a previously existing stable configuration.

In this exercise, you will deploy Windows Server 2012 in a virtual machine. You will then create a stable configuration for that virtual machine, and create a virtual machine checkpoint. Finally, you will modify the configuration, and roll back to the checkpoint.

The main tasks for this exercise are as follows:

1. Deploy Windows Server 2012 in a virtual machine.
2. Create a virtual machine checkpoint.
3. Modify the virtual machine.
4. Revert to the existing virtual machine checkpoint.
5. View resource metering data.

► Task 1: Deploy Windows Server 2012 in a virtual machine

1. Use the Hyper-V Manager console to start LON-GUEST1.
2. Open the Virtual Machine Connection Window, and perform the following steps to deploy Windows Server 2012 on the virtual machine:
 - On the **Settings** page, click **Next** to accept the **Region and Language** settings.
 - On the **Settings** page, click **I Accept**.
 - On the **Settings** page, enter the password **Pa\$\$w0rd** twice, and then click **Finish**.
3. Sign in to the virtual machine using the account **Administrator** and the password **Pa\$\$w0rd**.
4. Reset the name of the virtual machine to **LON-GUEST1**, and then restart the virtual machine.

► Task 2: Create a virtual machine checkpoint

1. Sign in to the LON-GUEST1 virtual machine, and verify that the name of the computer is set to **LON-GUEST1**.
2. Create a checkpoint of LON-GUEST1, and name the checkpoint **Before Change**.

► Task 3: Modify the virtual machine

1. Sign in to the LON-GUEST1 virtual machine, and use the Server Manager console to change the computer's name to **LON-Computer1**.
2. Reboot the virtual machine.
3. Sign in to the LON-GUEST1 virtual machine, and verify that the server name is set to **LON-Computer1**.

► Task 4: Revert to the existing virtual machine checkpoint

1. Use the Virtual Machine Connection window to revert the virtual machine.
2. Verify that the **Computer Name** of the virtual machine is now set to **LON-GUEST1**.

► Task 5: View resource metering data

1. On LON-HOST1, issue the following command:

```
Measure-VM LON-GUEST1
```

2. Note the average CPU, average RAM, and total disk use figures, and then close Windows PowerShell.

Results: After completing this exercise, you should have used virtual machine checkpoints to recover from a virtual machine misconfiguration.

► Revert the virtual machines

After you finish the lab, restart the computer in Windows Server 2012.

1. On the taskbar, click the **Windows PowerShell** icon.
2. In the Windows PowerShell window, enter the following command, and then press Enter:

```
Shutdown /r /t 5
```

3. From the Windows Boot Manager, select **Windows Server 2012**.

Lab Review Questions

Question: What type of virtual network switch would you create if you wanted to allow the virtual machine to communicate with the LAN that is connected to the Hyper-V virtualization server?

Question: How can you ensure that no one single virtual machine uses all available bandwidth provided by the Hyper-V virtualization server?

Question: What Dynamic Memory configuration task was not possible on previous versions of Hyper-V, but which you can now perform on a virtual machine that is hosted on the Hyper-V role on a Windows Server 2012 server?

Module Review and Takeaways

Module Review Questions

Question: In which situations should you use a fixed memory allocation instead of Dynamic Memory?

Question: In which situations must you use virtual hard disks with the new .vhdx format instead of virtual hard disks with the old .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?

Best Practices

When implementing server virtualization with Hyper-V, use the following best practices:

- Ensure that the processor on the computer that will run Hyper-V supports hardware assisted virtualization.
- Ensure that a virtualization server is provisioned with adequate RAM. Having multiple virtual machines paging the hard disk drive because they have inadequate memory decreases performance for all virtual machines on the server.
- Monitor virtual machine performance carefully. A virtual machine that uses a disproportionate amount of server resources can reduce the performance of all other virtual machines that are hosted on the same virtualization server.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Cannot deploy Hyper-V on an x64 platform.	
Virtual machine does not use Dynamic Memory.	

Tools

You can use the following tools with Hyper-V to deploy and manage virtual machines.

Tool	Use	Where to find it
Sysinternals disk2vhd tool	Use to convert physical hard disks to VHD format.	Microsoft TechNet website

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.




Module 1: Deploying and Managing Windows Server 2012

Lab: Deploying and Managing Windows Server 2012 R2

Exercise 1: Deploying Windows Server 2012 R2

► Task 1: Install the Windows Server® 2012 server

1. Open the Hyper-V® Manager console.
 2. Click **20410C-LON-SVR3**.
 3. In the Actions pane, click **Settings**.
 4. Under Hardware, click **DVD Drive**.
 5. Click **Image file**, and then click **Browse**.
 6. Browse to **D:\Program Files\Microsoft Learning\20410\Drives**, and then click **Windows2012R2RTM.iso**.
 7. Click **Open**, and then click **OK**.
 8. In the Hyper-V Manager console, double-click **20410C-LON-SVR3**.
 9. In the Virtual Machine Connection Window, in the **Action** menu, click **Start**.
 10. In the Windows Setup Wizard, on the **Windows Server 2012 R2** page, 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**
 11. On the **Windows Server 2012 R2** page, click **Install now**.
 12. On the **Select the operating system you want to install** page, select **Windows Server 2012 R2 Datacenter Evaluation (Server with a GUI)**, and then click **Next**.
 13. On the **License terms** page, review the operating system license terms, select the **I accept the license terms** check box, and then click **Next**.
 14. On the **Which type of installation do you want?** page, click **Custom: Install Windows only (advanced)**.
 15. On the **Where do you want to install Windows?** page, verify that **Drive 0 Unallocated Space** has enough space for the Windows Server 2012 R2 operating system, and then click **Next**.
-  **Note:** Depending on the speed of the equipment, the installation takes approximately 20 minutes. The virtual machine will restart several times during this process.
16. On the **Settings** page, in both the **Password** and **Reenter password** boxes, enter the password **Pa\$\$w0rd**, and then click **Finish**.

► Task 2: Change the server name

1. Sign in to LON-SVR3 as **Administrator** with the password **Pa\$\$w0rd**.
2. In Server Manager, click **Local Server**.
3. Click the randomly-generated name next to **Computer name**.
4. In the **System Properties** dialog box, on the **Computer Name** tab, click **Change**.
5. In the **Computer Name/Domain Changes** dialog box, in the **Computer name** text box, enter the name **LON-SVR3**, and then click **OK**.
6. In the **Computer Name/Domain Changes** dialog box, click **OK**.
7. Close the **System Properties** dialog box.
8. In the **Microsoft Windows** dialog box, click **Restart Now**.

► Task 3: Change the date and time

1. Sign in to server LON-SVR3 as **Administrator** with the password **Pa\$\$w0rd**.
2. On the taskbar, click the time display. A pop-up window with a calendar and a clock appears.
3. In the pop-up window, click **Change date and time settings**.
4. In the **Date and Time** dialog box, click **Change Time Zone**.
5. In the **Time Zone Settings** dialog box, set the time zone to your current time zone, and then click **OK**.
6. In the **Date and Time** dialog box, click **Change Date and Time**.
7. Verify that the date and time that display in the **Date and Time Settings** dialog box match those in your classroom, and then click **OK**.
8. To close the **Date and Time** dialog box, click **OK**.

► Task 4: Configure the network

1. On LON-SVR3, in the Server Manager console, click **Local Server**.
2. In the **Server Manager** console, next to **Ethernet**, click **IPv4 Address Assigned by DHCP, IPv6 Enabled**.
3. In the **Network Connections** dialog box, right-click **Ethernet**, and then click **Properties**.
4. In the **Ethernet Properties** dialog box, click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
5. In the **Internet Protocol Version 4 (TCP/IPv4) Properties** dialog box, click **Use the following IP address**, enter the following IP address information, and then click **OK**:
 - IP address: **172.16.0.101**
 - Subnet Mask: **255.255.0.0**
 - Default Gateway: **172.16.0.1**
 - Preferred DNS server: **172.16.0.10**
6. Click **Close** to close the **Ethernet Properties** dialog box.
7. Close the **Network Connections** dialog box.

► Task 5: Add the server to the domain

1. On LON-SVR3, in the Server Manager console, click **Local Server**.
2. Next to Workgroup, click **WORKGROUP**.
3. In the **System Properties** dialog box, on the **Computer Name** tab, click **Change**.
4. In the **Computer Name/Domain Changes** dialog box, in the **Member Of** area, click the **Domain** option.
5. In the **Domain** box, type **adatum.com**, and then click **OK**.
6. In the **Windows Security** dialog box, enter the following details, and then click **OK**:
 - Username: **Administrator**
 - Password: **Pa\$\$w0rd**
7. In the **Computer Name/Domain Changes** dialog box, click **OK**.
8. When informed that you must restart the computer to apply changes, click **OK**.
9. In the **System Properties** dialog box, click **Close**.
10. In the Microsoft Windows dialog box, click **Restart Now**.
11. After LON-SVR3 restarts, sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.

Results: After completing this exercise, you should have deployed Windows Server 2012 R2 on LON-SVR3. You also should have configured LON-SVR3, including name change, date and time, networking.

Exercise 2: Configuring Windows Server 2012 R2 Server Core

► Task 1: Set computer name

1. Sign in to LON-CORE as **Administrator** with the password **Pa\$\$w0rd**.
2. At the command prompt, type **sconfig.cmd** and press Enter.
3. To select Computer Name, type **2**, and then press Enter.
4. Enter the computer name **LON-CORE**, and then press Enter.
5. In the **Restart** dialog box, click **Yes**.
6. Sign in to server LON-CORE using the **Administrator** account with the password **Pa\$\$w0rd**.
7. At the command prompt, type **hostname**, and then press Enter to verify the computer's name.

► Task 2: Change the computer's date and time

1. Ensure you are signed in to server LON-CORE as **Administrator** with the password **Pa\$\$w0rd**.
2. At the command prompt, type **sconfig.cmd**, and then press Enter.
3. To select **Date and Time**, type **9**, and then press Enter.
4. In the **Date and Time** dialog box, click **Change time zone**. Set the time zone to the same time zone that your classroom uses, and then click **OK**.
5. In the **Date and Time** dialog box, click **Change Date and Time**, and verify that the date and time match those in your location. To dismiss the dialog boxes, click **OK** two times.
6. In the Command Prompt window, type **15**, and then press Enter to exit **Server Configuration**.

► Task 3: Configure the network

1. Ensure that you are signed in to server LON-CORE using the account **Administrator** and password **Pa\$\$w0rd**.
2. At the command prompt, type **sconfig.cmd**, and then press Enter.
3. To configure Network Settings, type **8**, and then press Enter.
4. Type the index number of the network adapter that you want to configure, and then press Enter.
5. On the **Network Adapter Settings** page, type **1**, and then press Enter. This sets the Network Adapter Address.
6. To select static IP address configuration, type **S**, and then press Enter.
7. At the Enter static IP address: prompt, type **172.16.0.111**, and then press Enter.
8. At the Enter subnet mask prompt, Type **255.255.0.0**, and then press Enter.
9. At the Enter default gateway prompt, type **172.16.0.1**, and then press Enter.
10. On the **Network Adapter Settings** page, type **2**, and then press Enter.

This configures the DNS server address.

11. At the Enter new preferred DNS server prompt, type **172.16.0.10**, and then press Enter.
12. In the **Network Settings** dialog box, click **OK**.
13. Press Enter to not configure an alternate DNS server address.
14. Type **4**, and then press Enter to return to the main menu.
15. Type **15**, and then press Enter to exit sconfig.cmd.
16. At the command prompt, type **ping lon-dc1.adatum.com** to verify connectivity to the domain controller from LON-CORE.

► Task 4: Add the server to the domain

1. Ensure that you are signed in to server LON-CORE using the account **Administrator** with password **Pa\$\$w0rd**.
2. At the command prompt, type **sconfig.cmd**, and then press Enter.
3. To switch to configure Domain/Workgroup, type **1**, and then press Enter.
4. To join a domain, type **D**, and then 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 then press Enter.
7. At the Type the password associated with the domain user prompt, type **Pa\$\$w0rd** and then press Enter.
8. At the Change Computer Name prompt, click **No**.
9. In the **Restart** dialog box, click **Yes**.
10. Sign in to server LON-CORE with the **Adatum\Administrator** account and the password **Pa\$\$w0rd**.

Results: After completing this exercise, you should have configured a Windows Server 2012 R2 Server Core deployment and verified the server's name.

Exercise 3: Managing Servers

► Task 1: Create a server group

1. Sign in to LON-DC1 with the **Administrator** account and the password **Pa\$\$w0rd**.
2. In the Server Manager console, click **Dashboard**, and then click **Create a server group**.
3. In the **Create Server Group** dialog box, click the **Active Directory** tab, and then click **Find Now**.
4. In the **Server group name** box, type **LAB-1**.
5. Use the arrow to add **LON-CORE** and **LON-SVR3** to the server group. Click **OK** to close the **Create Server Group** dialog box.
6. In the Server Manager console, click **LAB-1**. Press and hold the **Ctrl** key, and then select both **LON-CORE** and **LON-SVR3**.
7. Scroll down, and under the **Performance** section, select both **LON-CORE** and **LON-SVR3**.
8. Right-click **LON-CORE**, and then click **Start Performance Counters**.

► Task 2: Deploy features and roles to both servers

1. In Server Manager on LON-DC1, click **LAB-1**.
2. Scroll to the top of the pane, right-click **LON-CORE**, and then click **Add Roles and Features**.
3. In the Add Roles and Features Wizard, click **Next**.
4. On the **Select installation type** page, click **Role-based or feature-based installation**, and then click **Next**.
5. On the **Select destination server** page, verify that **LON-CORE.Adatum.com** is selected, and then click **Next**.
6. On the **Select server roles** page, select **Web Server (IIS)**, and then click **Next**.
7. On the **Features** page, select **Windows Server Backup**, and then click **Next**.
8. On the **Web Server Role (IIS)** page, click **Next**.
9. On the **Select Role Services** page, add the **Windows Authentication** role service, and then click **Next**.
10. On the **Confirm installation selections** page, select the **Restart the destination server automatically if required** check box, and then click **Install**.
11. Click **Close** to close the Add Roles and Features Wizard.
12. In Server Manager, right-click **LON-SVR3**, and then click **Add Roles and Features**.
13. In the Add Roles and Features Wizard, on the **Before you begin** page, click **Next**.
14. On the **Select installation type** page, click **Role-based or feature-based installation**. Click **Next**.
15. On the **Select destination server** page, verify that **LON-SVR3.Adatum.com** is selected, and then click **Next**.
16. On the **Server Roles** page, click **Next**.
17. On the **Select features** page, click **Windows Server Backup**, and then click **Next**.
18. On the **Confirm installation selections** page, select the **Restart the destination server automatically if required** check box, and then click **Install**.

19. Once the install commences, click **Close**.
20. In Server Manager, refresh the view, click the **IIS** node, and then verify that LON-CORE is listed.

► **Task 3: Review services and change a service setting**

1. Sign in to LON-CORE with the **Adatum\Administrator** account and the password **Pa\$\$w0rd**.
2. In the Command Prompt window, type the following command, and then press Enter:

```
netsh.exe advfirewall firewall set rule group="remote desktop" new enable=yes
```

3. Sign in to LON-DC1 with the **Adatum\Administrator** account and the password **Pa\$\$w0rd**.
4. In Server Manager, click **LAB-1**.
5. Right-click **LON-CORE**, and then click **Computer Management**.
6. In the Computer Management console, expand **Services and Applications**, and then click **Services**.
7. Right-click the **World Wide Web Publishing** service, and then click **Properties**. Verify that the **Startup type** is set to **Automatic**.
8. In the **World Wide Web Publishing Service** dialog box, on the **Log On** tab, verify that the service is configured to use the **Local System account**.
9. On the **Recovery** tab, configure the following settings, and then click the **Restart Computer Options** button:
 - First failure: **Restart the Service**
 - Second failure: **Restart the Service**
 - Subsequent failures: **Restart the Computer**.
 - Reset fail count after: **1** days
 - Restart service after: **1** minute
10. In the **Restart Computer Options** dialog box, in the **Restart Computer After** box, type **2**, and then click **OK**.
11. Click **OK** to close the **World Wide Web Publishing Services Properties** dialog box.
12. Close the Computer Management console.

Results: After completing this exercise, you should have created a server group, deployed roles and features, and configured the properties of a service.

Exercise 4: Using Windows PowerShell to Manage Servers

► **Task 1: Use Windows PowerShell to connect remotely to servers and view information**

1. Sign in to LON-DC1 with the **Adatum\Administrator** account and the password **Pa\$\$w0rd**.
2. In the Server Manager console, click **LAB-1**.
3. Right-click **LON-CORE**, and then click **Windows PowerShell**.

- At the command prompt, type the following, and then press Enter:

```
Import-Module ServerManager
```

- To review the roles and features installed on LON-CORE, at the command prompt, type the following, and then press Enter:

```
Get-WindowsFeature
```

- To review the running services on LON-CORE, at the command prompt, type the following, and then press Enter:

```
Get-service | where-object {$_.status -eq "Running"}
```

- To view a list of processes on LON-CORE, at the command prompt, type the following, and then press Enter:

```
Get-process
```

- To review the IP addresses assigned to the server, at the command prompt, type the following, and then press Enter:

```
Get-NetIPAddress | Format-table
```

- To review the most recent 10 items in the security log, at the command prompt, type the following, and then press Enter:

```
Get-EventLog Security -Newest 10
```

- Close Windows PowerShell.

► Task 2: Use Windows PowerShell to remotely install new features

- On LON-DC1, on the taskbar, click the **Windows PowerShell** icon.
- To verify that the XPS Viewer feature has not been installed on LON-SVR3, type the following command, and then press Enter:

```
Get-WindowsFeature -ComputerName LON-SVR3
```

- To deploy the XPS Viewer feature on LON-SVR3, type the following command, and then press Enter:

```
Install-WindowsFeature XPS-Viewer -ComputerName LON-SVR3
```

- To verify that the XPS Viewer feature has now been deployed on LON-SVR3, type the following command and then press Enter:

```
Get-WindowsFeature -ComputerName LON-SVR3
```

- In the Server Manager console, from the **Tools** drop-down menu, click **Windows PowerShell ISE**.

- In the Windows PowerShell ISE window, in the Untitled1.ps1 script pane, type the following, pressing Enter after each line:

```
Import-Module ServerManager
Install-WindowsFeature WINS -ComputerName LON-SVR3
Install-WindowsFeature WINS -ComputerName LON-CORE
```

- Click the **Save** icon.

8. Select the root of **Local Disk (C:)**.
9. Create a new folder named **Scripts**, and then save the script in that folder as **InstallWins.ps1**.
10. To run the script, press the F5 key.

Results: After completing this exercise, you should have used Windows PowerShell to perform a remote installation of features on multiple servers.

► Prepare for the next module

After you complete the lab, revert the virtual machines back to their initial state. To do this, complete the following steps:

1. On the host computer, switch to the Hyper-V Manager console.
2. In the **Virtual Machines** list, right click **20410C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for 20410C-LON-CORE and 20410C-LON-SVR3.

Module 2: Introduction to Active Directory Domain Services

Lab: Installing Domain Controllers

Exercise 1: Installing a Domain Controller

► Task 1: Add an Active Directory Domain Services (AD DS) role to a member server

1. On LON-DC1, in Server Manager, in the left column, click **All Servers**.
2. Right-click **All Servers**, and then click **Add Servers**.
3. In the **Add Servers** dialog box, in the **Name (CN)** box, type **LON-SVR1**, and then click **Find Now**.
4. Under Name, click **LON-SVR1**, and then click the arrow to add the server to the Selected column.
5. Click **OK** to close the **Add Servers** dialog box.
6. In Server Manager, in the Servers pane, right-click **LON-SVR1**, and then select **Add Roles and Features**.
7. In the Add Roles and Features Wizard, click **Next**.
8. On the **Select installation type** page, ensure that **Role-based or feature-based installation** is selected, and then click **Next**.
9. On the **Select destination server** page, ensure that **Select a server from the server pool** is selected.
10. Under **Server Pool**, verify that **LON-SVR1.Adatum.com** is highlighted, and then click **Next**.
11. On the **Select server roles** page, select the **Active Directory Domain Services** check box, click **Add Features**, and then click **Next**.
12. On the **Select features** page, click **Next**.
13. On the **Active Directory Domain Services** page, click **Next**.
14. On the **Confirm installation selections** page, select the **Restart the destination server automatically if required** check box, and then click **Install**.
Installation will take several minutes.
15. When the installation completes, click **Close** to close the Add Roles and Features Wizard.

► Task 2: Configure a server as a domain controller

1. On LON-DC1, in Server Manager, on the command bar, click the **Notifications** icon—it looks like a flag.
2. Under Post-deployment Configuration, click **Promote this server to a domain controller**.
The Active Directory Domain Services Configuration Wizard will open.
3. In the Active Directory Domain Services Configuration Wizard, on the **Deployment Configuration** page, ensure that **Add a domain controller to an existing domain** is selected, and then, beside the Domain line, click **Select**.
4. In the **Windows Security** dialog box, in the **Username** box type **Administrator**, in the **Password** box type **Pa\$\$w0rd** and then click **OK**.
5. In the **Select a domain from the forest** dialog box, click **adatum.com**, and then click **OK**.
6. Beside the **Supply the credentials to perform this operation** line, click **Change**.

7. In the **Windows Security** dialog box, in the **Username** box, type **Adatum\Administrator**, in the **Password** box, type **Pa\$\$w0rd**, and then click **OK**.
8. On the **Deployment Configuration** page, click **Next**.
9. On the **Domain Controller Options** page, ensure that **Domain Name System (DNS) server** is selected, and then deselect **Global Catalog (GC)**.



Note: You would usually also want to enable the global catalog, but for the purpose of this lab, this is done in the next lab task.

10. In the **Type the Directory Services Restore Mode (DSRM) password** section, type **Pa\$\$w0rd** in both text boxes, and then click **Next**.
11. On the **DNS Options** page, click **Next**.
12. On the **Additional Options** page, Click **Next**.
13. On the **Paths** page, accept the default folders, and then click **Next**.
14. On the **Review Options** page, click **View Script**, examine the Windows PowerShell script that the wizard generates.
15. Close the Notepad window.
16. On the **Review Options** page, click **Next**.
17. On the **Prerequisites Check** page, read any warning messages, and then click **Install**.
18. When the task completes successfully, click **Close**.
19. Wait for LON-SVR1 to restart.

► **Task 3: Configure a server as a global catalog server**

1. Sign in to LON-SVR1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. In Server Manager, click **Tools**, and then click **Active Directory Sites and Services**.
3. When Active Directory Sites and Services opens, expand **Sites**, expand **Default-First-Site-Name**, expand **Servers**, and then expand **LON-SVR1**.
4. In the left column, right-click **NTDS Settings**, and then click **Properties**.
5. In the **NTDS Settings Properties** dialog box, select **Global Catalog (GC)**, and then click **OK**.
6. Close Active Directory Sites and Services.

Results: After completing this exercise, you will have explored Server Manager and promoted a member server to be a domain controller.

Exercise 2: Installing a Domain Controller by Using IFM

► Task 1: Use the Ntdsutil tool to generate IFM

1. On LON-DC1, in the lower-left corner of the screen, click the **Start** button.
2. On the Start screen, type **CMD**, right click **Command Prompt** and then click **Run as administrator**.
3. At a command prompt, type the following, pressing Enter after each line:

```
Ntdsutil  
Activate instance ntds  
Ifm  
Create sysvol full c:\ifm
```

4. Wait for the IFM command to complete and then close the command prompt.

► Task 2: Add the AD DS role to the member server

1. Switch to **LON-SVR2**, and, if required, sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. In the lower-left corner of the screen, click the **Start** button.
3. On the Start screen, type **CMD**, and then press Enter.
4. Type the following command, and then press Enter:

```
Net use k: \\LON-DC1\c$\IFM
```

5. Switch to **Server Manager**.
6. From the list on the left, click **Local Server**.
7. In the toolbar, click **Manage**, and then click **Add Roles and Features**.
8. On the **Before you begin** page, click **Next**.
9. On the **Select installation type** page, ensure that **Role-based or feature-based installation** is selected, and then click **Next**.
10. On the **Select destination server** page, verify that **LON-SVR2.Adatum.com** is highlighted, and then click **Next**.
11. On the **Select server roles** page, click **Active Directory Domain Services**.
12. In the Add Roles and Features Wizard, click **Add Features**, and then click **Next**.
13. On the **Select Features** page, click **Next**.
14. On the **Active Directory Domain Services** page, click **Next**.
15. On the **Confirm installation selections** page, click **Restart the destination server automatically if required**. Click **Yes** at the message box.
16. Click **Install**.
17. After the installation completes, click **Close**.



Note: If you see a message stating that a delegation for the DNS server cannot be created, click **OK**.

► Task 3: Use IFM to configure a member server as a new domain controller

1. On LON-SVR2, At the command prompt, type the following command, and then press Enter:

```
Robocopy k: c:\ifm /copyall /s
```

2. Close the Command Prompt window.
3. In Server Manager, on the command bar, click the **Notifications** icon.
4. Under Post-deployment Configuration, click **Promote this server to a domain controller**.
The Active Directory Domain Services Configuration Wizard will open.
5. On the **Deployment Configuration** page, ensure that **Add a domain controller to an existing domain** is selected, and confirm that **adatum.com** is the target domain. Click **Next**.
6. On the **Domain Controller Options** page, ensure that both **Domain Name System (DNS) server** and **Global Catalog (GC)** are selected. For the **DSRM** password, type **Pa\$\$w0rd** in both boxes, and then click **Next**.
7. On the **DNS Options** page, click **Next**.
8. On the **Additional Options** page, select **Install from media**, in the **Install from media path** box, type **C:\ifm**, and then click **verify**.
9. When the path has been verified, click **Next**.
10. On the **Paths** page, click **Next**.
11. On the **Review Options** page, click **Next**, and then observe the Active Directory Domain Services Configuration Wizard as it performs a check for prerequisites.
12. Click **Install**, and wait while AD DS is configured.
While this task is running, read the information messages that display on the screen.
13. Wait for the server to restart.

Results: After completing this exercise, you will have installed an additional domain controller for the branch office by using IFM.

► Prepare for the next module

When you have completed the lab, revert the virtual machines back 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 **20410C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for 20410C-LON-SVR1, 20410C-LON-RTR, and 20410C-LON-SVR2.

Module 3: Managing Active Directory Domain Services Objects

Lab: Managing Active Directory Domain Services Objects

Exercise 1: Delegating Administration for a Branch Office

► Task 1: Delegate administration for Branch Administrators

1. Switch to LON-DC1.
2. In Server Manager, click **Tools**, and then click **Active Directory Users and Computers**.
3. In Active Directory Users and Computers, click **Adatum.com**.
4. Right-click **Adatum.com**, point to **New**, and then click **Organizational Unit**.
5. In the **New Object – Organizational Unit** dialog box, in **Name**, type **Branch Office 1**, and then click **OK**.
6. Right-click **Branch Office 1**, point to **New**, and then click **Group**.
7. In the **New Object – Group** dialog box, in **Group name**, type **Branch 1 Help Desk**, and then click **OK**.
8. Repeat steps 6 and 7 using **Branch 1 Administrators** as the new group name.
9. Repeat steps 6 and 7 using **Branch 1 Users** as the new group name.
10. In the navigation pane, click **IT**.
11. In the details pane, right-click **Holly Dickson**, and then click **Move**.
12. In the **Move** dialog box, click **Branch Office 1**, and then click **OK**.
13. Repeat steps 10 through 12 for the following OU's and users:
 - **Development** and the user **Bart Duncan**
 - **Managers** and the user **Ed Meadows**
 - **Marketing** and the user **Connie Vrettos**
 - **Research** and the user **Barbara Zighetti**
 - **Sales** and the user **Arlene Huff**
14. In the navigation pane, click **Computers**.
15. In the details pane, right-click **LON-CL1**, and then click **Move**.
16. In the **Move** dialog box, click **Branch Office 1**, and then click **OK**.
17. Switch to LON-CL1.
18. Point the mouse at the lower-right corner of the screen, and then click **Settings**.
19. Click **Power**, and then click **Restart**.
20. When the computer has restarted, sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
21. Switch to LON-DC1.
22. If necessary, switch to **Active Directory Users and Computers**.
23. In the navigation pane, right-click **Branch Office 1**, click **Delegate Control**, and then click **Next**.

24. On the **Users or Groups** page, click **Add**.
25. In the **Select Users, Computers, or Groups** dialog box, in **Enter the object names to select (examples)**, type **Branch 1 Administrators**, and then click **OK**.
26. On the **Users or Groups** page, click **Next**.
27. On the **Tasks to Delegate** page, in the **Delegate the following common tasks** list, select the following check boxes, and then click **Next**:
 - **Create, delete, and manage user accounts**
 - **Reset user passwords and force password change at next logon**
 - **Read all user information**
 - **Create, delete and manage groups**
 - **Modify the membership of a group**
 - **Manage Group Policy links**
28. On the **Completing the Delegation of Control Wizard** page, click **Finish**.
29. In the navigation pane, right-click **Branch Office 1**, click **Delegate Control**, and then click **Next**.
30. On the **Users or Groups** page, click **Add**.
31. In the **Select Users, Computers, or Groups** dialog box, in **Enter the object names to select (examples)**, type **Branch 1 Administrators**, and then click **OK**.
32. On the **Users or Groups** page, click **Next**.
33. On the **Tasks to Delegate** page, click **Create a custom task to delegate**, and then click **Next**.
34. On the **Active Directory Object Type** page, select **Only the following objects in the folder**, select the following check boxes, and then click **Next**:
 - **Computer objects**
 - **Create selected objects in this folder**
 - **Delete selected objects in this folder**
35. On the **Permissions** page, select both **General** and **Full Control**, and then click **Next**.
36. On the **Completing the Delegation of Control Wizard** page, click **Finish**.

► **Task 2: Delegate a user administrator for the Branch Office Help Desk**

1. On LON-DC1, in the navigation pane, right-click **Branch Office 1**, click **Delegate Control**, and then click **Next**.
2. On the **Users or Groups** page, click **Add**.
3. In the **Select Users, Computers, or Groups** dialog box, in **Enter the object names to select (examples)**, type **Branch 1 Help Desk**, and then click **OK**.
4. On the **Users or Groups** page, click **Next**.
5. On the **Tasks to Delegate** page, in the **Delegate the following common tasks** list, select the following check boxes, and then click **Next**:
 - **Reset user passwords and force password change at next logon**
 - **Read all user information**
 - **Modify the membership of a group**
6. On the **Completing the Delegation of Control Wizard** page, click **Finish**.

► Task 3: Add a member to the Branch Administrators

1. On LON-DC1, in the navigation pane, click **Branch Office 1**.
2. In the details pane, right-click **Holly Dickson**, and then click **Add to a group**.
3. In the **Select Groups** dialog box, in **Enter the object names to select (examples)**, type **Branch 1 Administrators**, and then click **OK**.
4. In the **Active Directory Domain Services** dialog box, click **OK**.
5. In the details pane, right-click **Branch 1 Administrators**, and then click **Add to a group**.
6. In the **Select Groups** dialog box, in **Enter the object names to select (examples)**, type **Server Operators**, and then click **OK**.
7. In the **Active Directory Domain Services** dialog box, click **OK**.
8. On your host computer, in the 20410C-LON-DC1 window, on the **Action** menu, click **Ctrl+Alt+Delete**.
9. On LON-DC1, click **Sign out**.
10. Sign in to LON-DC1 as **Adatum\Holly** with the password **Pa\$\$w0rd**.
You can sign in locally at a domain controller because Holly belongs indirectly to the Server Operators domain local group.
11. On the taskbar, click the **Server Manager** icon.
12. In the **User Account Control** dialog box, in **User name**, type **Holly**. In **Password**, type **Pa\$\$w0rd**, and then click **Yes**.
13. In Server Manager, click **Tools**, and then click **Active Directory Users and Computers**.
14. In Active Directory Users and Computers, expand **Adatum.com**.
15. In the navigation pane, click **Sales**.
16. In the details pane, right-click **Aaren Ekelund**, and then click **Delete**.
17. Click **Yes** to confirm.
18. Click **OK** to acknowledge that you do not have permissions to perform this task.
19. In the navigation pane, click **Branch Office 1**.
20. In the details pane, right-click **Ed Meadows**, and then click **Delete**.
21. Click **Yes** to confirm.

You are successful because you have the required permissions.

► Task 4: Add a member to the Branch Help Desk group

1. On LON-DC1, in the details pane, right-click **Bart Duncan**, and then click **Add to a group**.
2. In the **Select Groups** dialog box, in **Enter the object names to select (examples)**, type **Branch 1 Help Desk**, and then click **OK**.
3. In the **Active Directory Domain Services** dialog box, click **OK**.
4. Close Active Directory Users and Computers.
5. Close Server Manager.
6. On the desktop, click **Server Manager**. In the **User Account Control** dialog box, in **User name**, type **Adatum\Administrator**.

7. In **Password**, type **Pa\$\$w0rd**, and then click **Yes**.



Note: To modify the Server Operators membership list, you must have permissions beyond those available to the Branch 1 Administrators group.

8. In Server Manager, click **Tools**.
9. In the Tools list, click **Active Directory Users and Computers**.
10. In Active Directory Users and Computers, expand **Adatum.com**.
11. In the navigation pane, click **Branch Office 1**.
12. In the details pane, right-click **Branch 1 Help Desk**, and then click **Add to a group**.
13. In the **Select Groups** dialog box, in **Enter the object names to select (examples)**, type **Server Operators**, and then click **OK**.
14. In the **Active Directory Domain Services** dialog box, click **OK**.
15. On your host computer, in the 20410C-LON-DC1 window, on the **Action** menu, click **Ctrl+Alt+Delete**.
16. On LON-DC1, click **Sign out**.
17. Sign in as **Adatum\Bart** with the password **Pa\$\$w0rd**.
You can sign in locally at a domain controller because Bart belongs indirectly to the Server Operators domain local group.
18. On the desktop, click **Server Manager**.
19. In the **User Account Control** dialog box, in **User name**, type **Bart**. In **Password**, type **Pa\$\$w0rd**, and then click **Yes**.
20. In Server Manager, click **Tools**.
21. Click **Active Directory Users and Computers**.
22. In Active Directory Users and Computers, expand **Adatum.com**.
23. In the navigation pane, click **Branch Office 1**.
24. In the details pane, right-click **Connie Vrettos**, and then click **Delete**.
25. Click **Yes** to confirm.
You are unsuccessful because Bart lacks the required permissions.
26. Click **OK**.
27. Right-click **Connie Vrettos**, and then click **Reset Password**.
28. In the **Reset Password** dialog box, in **New password** and **Confirm password**, type **Pa\$\$w0rd**, and then click **OK**.
29. Click **OK** to confirm the successful password reset.
30. On your host computer, in the 20410C-LON-DC1 window, on the **Action** menu, click **Ctrl+Alt+Delete**.
31. On LON-DC1, click **Sign out**.
32. Sign in to LON-DC1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.

Results: After completing this exercise, you will have successfully created an OU, and delegated administration of it to the appropriate group.

Exercise 2: Creating and Configuring User Accounts in AD DS

► Task 1: Create a user template for the branch office

1. On LON-DC1, on the taskbar, click the **File Explorer** icon.
2. Double-click **Local Disk (C:)**.
3. On the menu, click **Home**, and then click **New folder**.
4. Type **branch1-userdata**, and then press **Enter**.
5. Right-click **branch1-userdata**, and then click **Properties**.
6. In the **branch1-userdata Properties** dialog box, on the **Sharing** tab, click **Advanced Sharing**.
7. Select **Share this folder**, and then click **Permissions**.
8. In the **Permissions for branch1-userdata** dialog box, for the **Full Control** permission select the **Allow** check box, and then click **OK**.
9. In the **Advanced Sharing** dialog box, click **OK**, and then in the **branch1-userdata Properties** dialog box, click **Close**.
10. In Server Manager, click **Tools**, and then click **Active Directory Users and Computers**, and then expand **Adatum.com**.
11. Right-click **Branch Office1**, point to **New**, and then click **User**.
12. In the **New Object – User** dialog box, in **Full name**, type **_Branch_template**.
13. In **User logon name**, type **_Branch_template**, and then click **Next**.
14. In **Password** and **Confirm password**, type **Pa\$\$w0rd**.
15. Select the **Account is disabled** check box, and then click **Next**.
16. Click **Finish**.

► Task 2: Configure the template settings

1. On LON-DC1, from within the **Branch Office 1** OU, right-click **_Branch_template**, and then click **Properties**.
2. In the **_Branch_template Properties** dialog box, on the **Address** tab, in **City**, type **Slough**.
3. Click the **Member Of** tab, and then click **Add**.
4. In the **Select Groups** dialog box, in **Enter the object names to select (examples)**, type **Branch 1 Users**, and then click **OK**.
5. Click the **Profile** tab.
6. Under Home folder, click **Connect**, and in the **To** box, type **\\lon-dc1\branch1-userdata\%username%**.
7. Click **Apply**, and then click **OK**.

► Task 3: Create a new user for the branch office, based on the template

1. On LON-DC1, right-click **_Branch_template**, and then click **Copy**.
2. In the **Copy Object – User** dialog box, in **First name**, type **Ed**.
3. In **Last name**, type **Meadows**.
4. In **User logon name**, type **Ed**, and then click **Next**.

5. In **Password** and **Confirm password**, type **Pa\$\$w0rd**.
6. Clear the **User must change password at next login** check box.
7. Clear the **Account is disabled** check box, and then click **Next**.
8. Click **Finish**.
9. Right-click **Ed Meadows**, and then click **Properties**.
10. In the **Ed Meadows Properties** dialog box, on the **Address** tab, notice that the City is already configured.
11. Click the **Profile** tab.
Notice that the home folder location is already configured.
12. Click the **Member Of** tab.
Notice that Ed belongs to the Branch 1 Users group. Click **OK**.
13. On your host computer, in the 20410C-LON-DC1 window, on the **Action** menu, click **Ctrl+Alt+Delete**.
14. On LON-DC1, click **Sign out**.

► Task 4: Sign in as a user to test account settings

1. Switch to LON-CL1.
2. On your host computer, in the 20410C-LON-CL1 window, on the menu, click **Ctrl+Alt+Delete**.
3. On LON-CL1, click **Switch User**.
4. Sign in to LON-CL1 as **Adatum\Ed** with the password **Pa\$\$w0rd**.
5. On the Start screen, type **File Explorer** and then press **Enter**.
6. Verify that drive Z is present.
7. Double-click **Ed (\\lon-dc1\branch1-userdata) (Z:)**.
8. If you receive no errors, you have been successful.
9. On your host computer, in the 20410C-LON-CL1 window, on the **Action** menu, click **Ctrl+Alt+Delete**.
10. On LON-CL1, click **Sign out**.

Results: After completing this exercise, you will have successfully created and tested a user account created from a template.

Exercise 3: Managing Computer Objects in AD DS

► Task 1: Reset a computer account

1. On LON-DC1, sign in as **Adatum\Holly** with the password **Pa\$\$w0rd**.
2. On the taskbar, click the **Server Manager** icon.
3. In the **User Account Control** dialog box, in **User name**, type **Holly**.
4. In **Password**, type **Pa\$\$w0rd**, and then click **Yes**.
5. In Server Manager, click **Tools**, and then click **Active Directory Users and Computers**.

6. In Active Directory Users and Computers, expand **Adatum.com**.
7. In the navigation pane, click **Branch Office 1**.
8. In the details pane, right-click **LON-CL1**, and then click **Reset Account**.
9. In the **Active Directory Domain Services** dialog box, click **Yes**, and then click **OK**.

► **Task 2: Observe the behavior when a client logs on**

1. Switch to LON-CL1.
2. Sign in as **Adatum\Ed** with the password **Pa\$\$w0rd**.
A message appears stating that **The trust relationship between this workstation and the primary domain failed**.
3. Click **OK**.

► **Task 3: Rejoin the domain to reconnect the computer account**

1. On LON-CL1 click the back arrow and switch to **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On the Start screen, right-click the display, click **All apps**, and in the Apps list, click **Control Panel**.
3. In Control Panel, in the **View by** list, click **Large icons**, and then click **System**.
4. In the navigation list, click **Advanced system settings**.
5. In System Properties, click the **Computer Name** tab, and then click **Network ID**.
6. On the **Select the option that describes your network** page, click **Next**.
7. On the **Is your company network on a domain?** page, click **Next**.
8. On the **You will need the following information** page, click **Next**.
9. On the **Type your user name, password, and domain name for your domain account** page, in **Password**, type **Pa\$\$w0rd**. Leave the other boxes completed, and then click **Next**.
10. In the **User Account and Domain Information** dialog box, click **Yes**.
11. On the **Do you want to enable a domain user account on this computer?** page, click **Do not add a domain user account**, and then click **Next**.
12. Click **Finish**, and then click **OK**.
13. In the **Microsoft Windows** dialog box, click **Restart Now**.
14. Sign in as **Adatum\Ed** with the password **Pa\$\$w0rd**.
You are successful because the computer had been successfully rejoined.

Results: After completing this exercise, you will have successfully reset a trust relationship.

► **Prepare for the next module**

When you have completed the lab, revert the virtual machines back 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 **20410C-LON-CL1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for 20410C-LON-DC1.

Module 4: Automating Active Directory Domain Services Administration

Lab: Automating AD DS Administration by Using Windows PowerShell

Exercise 1: Creating User Accounts and Groups by Using Windows PowerShell

► Task 1: Create a user account by using Windows PowerShell

1. On LON-DC1, on the taskbar, click the **Windows PowerShell** icon.
2. At the Windows PowerShell prompt, type the following command, and then press Enter:

```
New-ADOrganizationalUnit LondonBranch
```

3. Type the following command, and then press Enter:

```
New-ADUser -Name Ty -DisplayName "Ty Carlson" -GivenName Ty -Surname Carlson -Path "ou=LondonBranch,dc=adatum,dc=com"
```

4. Type the following command, and then press Enter:

```
Set-ADAccountPassword Ty
```

5. When prompted for the current password, press Enter.
6. When prompted for the desired password, type **Pa\$\$w0rd**, and then press Enter.
7. When prompted to repeat the password, type **Pa\$\$w0rd**, and then press Enter.
8. At the Windows PowerShell prompt, type **Enable-ADAccount Ty**, and then press Enter.
9. On LON-CL1, sign in as **Ty** with the password **Pa\$\$w0rd**.
10. Verify that the sign-in is successful and then sign out of LON-CL1.

► Task 2: Create a group by using Windows PowerShell

1. To create a new global security group for users in the London branch office, on LON-DC1, at the Windows PowerShell prompt, type the following command, and then press Enter:

```
New-ADGroup LondonBranchUsers -Path "ou=LondonBranch,dc=adatum,dc=com" -GroupScope Global -GroupCategory Security
```

2. To add **Ty** as a member of LondonBranchUsers, type the following command, and then press Enter:

```
Add-ADGroupMember LondonBranchUsers -Members Ty
```

3. To confirm that Ty has been added as a member of LondonBranchUsers, type the following command, and then press Enter:

```
Get-ADGroupMember LondonBranchUsers
```

Results: After completing this exercise, you will have created user accounts and groups by using Windows PowerShell.

Exercise 2: Using Windows PowerShell to Create User Accounts in Bulk

► Task 1: Prepare the .csv file

1. On LON-DC1, on the taskbar, click the **File Explorer** icon.
2. In File Explorer, expand drive **E:**, expand **Labfiles**, and then click **Mod04**.
3. Right-click **LabUsers.ps1**, and then click **Edit**.
4. In Windows PowerShell Integrated Scripting Environment (ISE), read the comments at the top of the script, and then identify the requirements for the header in the .csv file.
5. Close **Windows PowerShell ISE**.
6. In File Explorer, double-click **LabUsers.csv**.
7. In the **How do you want to open this type of file (.csv)?** message, click **Notepad**.
8. In Notepad, type the following line at the top of the file:
FirstName,LastName,Department,DefaultPassword
9. Click **File**, and then click **Save**.
10. Close Notepad.

► Task 2: Prepare the script

1. On LON-DC1, in File Explorer, right-click **LabUsers.ps1**, and then click **Edit**.
2. In Windows PowerShell ISE, under Variables, replace **C:\path\file.csv** with **E:\Labfiles\Mod04\LabUsers.csv**.
3. Under Variables, replace **"ou=orgunit,dc=domain,dc=com"** with **"ou=LondonBranch,dc=adatum,dc=com"**.
4. Click **File**, and then click **Save**.
5. Scroll down and review the contents of the script.
6. Close Windows PowerShell ISE.

► Task 3: Run the script

1. On LON-DC1, on the taskbar, click the **Windows PowerShell** icon.
2. At the Windows PowerShell prompt, type **cd E:\Labfiles\Mod04**, and then press Enter.
3. Type **.\LabUsers.ps1**, and then press Enter.
4. Type the following command, and then press Enter:

```
Get-ADUser -Filter * -SearchBase "ou=LondonBranch,dc=adatum,dc=com"
```

5. Close Windows PowerShell.
6. On LON-CL1, sign in as **Luka** with the password **Pa\$\$w0rd**.

Results: After completing this exercise, you will have used Windows PowerShell to create user accounts in bulk.

Exercise 3: Using Windows PowerShell to Modify User Accounts in Bulk

► Task 1: Force all user accounts in LondonBranch to change their passwords at next sign in

1. On LON-DC1, on the taskbar, click the **Windows PowerShell** icon.
2. To create a query for user accounts in the LondonBranch OU, at the Windows PowerShell Prompt, type the following command, and then press Enter:

```
Get-ADUser -Filter * -SearchBase "ou=LondonBranch,dc=adatum,dc=com" | Format-Wide DistinguishedName
```

3. Verify that only users from the LondonBranch organizational unit (OU) are listed.
4. To modify the previous command to force all user to change their password the next time they sign in, at the Windows PowerShell prompt, type the following command, and then press Enter:

```
Get-ADUser -Filter * -SearchBase "ou=LondonBranch,dc=adatum,dc=com" | Set-ADUser -ChangePasswordAtLogon $true
```

5. Close **Windows PowerShell**.

► Task 2: Configure the address for user accounts in LondonBranch

1. On LON-DC1, in Server Manager, click **Tools**, and then click **Active Directory Administrative Center**.
2. In Active Directory Administrative Center, in the navigation pane, expand **Adatum (local)** and double-click **LondonBranch**.
3. Click the **Type** column header to sort based on the object type.
4. Select all user accounts, right-click the user accounts, and then click **Properties**.
5. In the Multiple Users pane, under Organization, select the **Address** check box.
6. In the **Street** box, type **Branch Office**.
7. In the **City** box, type **London**.
8. In the **Country/Region** box, click **United Kingdom**, and then click **OK**.
9. Close **Active Directory Administrative Center**.

Results: After completing this exercise, you will have modified user accounts in bulk.

► Prepare for the next module

When you finish the lab, revert all virtual machines back to their initial state by performing the following steps:

1. On the host computer, start Hyper-V® Manager.
2. In the **Virtual Machines** list, right-click **20410C-LON-CL1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 to 3 for **20410C-LON-DC1**.

Module 5: Implementing IPv4

Lab: Implementing IPv4

Exercise 1: Identifying Appropriate Subnets

► Task 1: Calculate the bits required to support the hosts on each subnet

- How many bits are required to support 100 hosts on the client subnet?

Answer: Seven bits are required to support 100 hosts on the client subnet ($2^7-2=126$, $2^6-2=62$).

- How many bits are required to support 10 hosts on the server subnet?

Answer: Four bits are required to support 10 hosts on the server subnet ($2^4-2=14$, $2^3-2=6$).

- How many bits are required to support 40 hosts on the future expansion subnet?

Answer: Six bits are required to support 40 hosts on the future expansion subnet ($2^6-2=62$, $2^5-2=30$).

- If all subnets are the same size, can they be accommodated?

Answer: No. If all subnets are the same size, then all subnets must use 7 bits to support 126 hosts. Only a single class C-sized address with 254 hosts has been allocated. Three subnets of 126 hosts would not fit.

- Which feature allows a single network to be divided into subnets of varying sizes?

Answer: Variable length subnet masking allows you to define different subnet masks when subnetting. Therefore, variable length subnet masking allows you to have subnets of varying sizes.

- How many host bits will you use for each subnet? Use the simplest allocation possible, which is one large subnet and two equal-sized smaller subnets.

Answer: The client subnet is 7 host bits. This allocation can accommodate up to 126 hosts and uses half of the allocated address pool.

The server and future expansion subnets are 6-host bits. This can accommodate up to 62 hosts on each subnet and uses the other half of the address pool.

► Task 2: Calculate subnet masks and network IDs

- Given the number of host bits allocated, what is the subnet mask that you will use for the client subnet? Calculate the subnet mask in binary and decimal.

- The client subnet is using 7 bits for the host ID. Therefore, you can use 25 bits for the subnet mask.

Binary	Decimal
11111111.11111111.11111111.10000000	255.255.255.128

2. Given the number of host bits allocated, what is the subnet mask that you can use for the server subnet? Calculate the subnet mask in binary and decimal.
- The server subnet is using 6 bits for the host ID. Therefore, you can use 26 bits for the subnet mask.

Binary	Decimal
11111111.11111111.11111111.11000000	255.255.255.192

3. Given the number of host bits allocated, what is the subnet mask that you can use for the future expansion subnet? Calculate the subnet mask in binary and decimal.
- The future expansion subnet is using 6 bits for the host ID. Therefore, you can use 26 bits for the subnet mask.

Binary	Decimal
11111111.11111111.11111111.11000000	255.255.255.192

4. For the client subnet, define the network ID, first available host, last available host, and broadcast address. Assume that the client subnet is the first subnet allocated from the available address pool. Calculate the binary and decimal versions of each address.

In the following table, the bits in bold are part of the network ID.

Description	Binary	Decimal
Network ID	11000000.10101000.1100010.00000000	192.168.98.0
First host	11000000.10101000.1100010.00000001	192.168.98.1
Last host	11000000.10101000.1100010.01111110	192.168.98.126
Broadcast	11000000.10101000.1100010.01111111	192.168.98.127

5. For the server subnet, define the network ID, first available host, last available host, and broadcast address. Assume that the server subnet is the second subnet allocated from the available address pool. Calculate the binary and decimal versions of each address.

In the following table, the bits in bold are part of the network ID.

Description	Binary	Decimal
Network ID	11000000.10101000.1100010.10000000	192.168.98.128
First host	11000000.10101000.1100010.10000001	192.168.98.129
Last host	11000000.10101000.1100010.10111110	192.168.98.190
Broadcast	11000000.10101000.1100010.10111111	192.168.98.191

6. For the future allocation subnet, define the network ID, first available host, last available host, and broadcast address. Assume that the future allocation subnet is the third subnet allocated from the available address pool. Calculate the binary and decimal versions of each address.

In the following table, the bits in bold are part of the network ID.

Description	Binary	Decimal
Network ID	11000000.10101000.1100010.11 000000	192.168.98.192
First host	11000000.10101000.1100010.11 000001	192.168.98.193
Last host	11000000.10101000.1100010.11 111110	192.168.98.254
Broadcast	11000000.10101000.1100010.11 111111	192.168.98.255

Results: After completing this exercise, you should have identified the subnets required to meet the requirements of the lab scenario.

Exercise 2: Troubleshooting IPv4

► Task 1: Prepare for troubleshooting

1. On LON-SVR2, on the taskbar, click the **Windows PowerShell** icon.
2. At the Windows PowerShell prompt, type the following cmdlet, and then press Enter:

```
Test-NetConnection LON-DC1
```

3. Verify that you receive a reply that contains **PingSucceeded:True** from **LON-DC1**.
4. Open a **File Explorer** window, and browse to **\\LON-DC1\E\$\Labfiles\Mod05**.
5. Right-click **Break2.ps1**, and then click **Run with PowerShell**.



Note: This script creates the problem that you will troubleshoot and repair in the next task.

6. Close File Explorer.

► Task 2: Troubleshoot IPv4 connectivity between LON-SVR2 and LON-DC1

1. On LON-SVR2, at the Windows PowerShell prompt, type the following, and then press Enter:

```
Test-NetConnection LON-DC1
```

2. Verify that you receive a reply that contains **PingSucceeded:False** from **LON-DC1**.
3. At the Windows PowerShell Prompt, type the following, and then press Enter:

```
Test-NetConnection -TraceRoute LON-DC1
```

Notice that the host is unable to find the default gateway, and that the following warning message appears: **"Name resolution of lon-dc1 failed – Status: HostNotFound."**

- At the Windows PowerShell Prompt, type the following, and then press Enter:

```
Get-NetRoute
```

Notice that the default route and the default gateway information is missing in the routing table.



Note: You should not be able to locate **DestinationPrefix 0.0.0.0/0** and **NextHop 10.10.0.1**.

- At the Windows PowerShell Prompt, type the following, and then press Enter:

```
Test-NetConnection 10.10.0.1
```

- Notice that the default gateway is responding by verifying that you receive a reply that contains **PingSucceeded:True** from **10.10.0.1**.
- At the Windows PowerShell Prompt, type the following, and then press Enter:

```
New-NetRoute -InterfaceAlias "Ethernet" -DestinationPrefix 0.0.0.0/0 -NextHop 10.10.0.1
```



Note: The **New-NetRoute** cmdlet will create the default route and the default gateway information that was missing.

- At the Windows PowerShell Prompt, type the following, and then press Enter:

```
Get-NetRoute
```

- Notice that the default route and the default gateway information is present in the routing table by locating **DestinationPrefix 0.0.0.0/0** and **NextHop 10.10.0.1**.
- At the Windows PowerShell prompt, type the following, and then press Enter:

```
Test-NetConnection LON-DC1
```

- Verify that you receive a reply that contains **PingSucceeded:True** from LON-DC1.

Results: After completing this lab, you should have resolved an IPv4 connectivity problem.

► Prepare for the next module

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

- On the host computer, start **Hyper-V Manager**.
- In the **Virtual Machines** list, right-click **20410C-LON-DC1**, and then click **Revert**.
- In the **Revert Virtual Machine** dialog box, click **Revert**.
- Repeat steps 2 and 3 for **20410C-LON-RTR** and **20410C-LON-SVR2**.

Module 6: Implementing Dynamic Host Configuration Protocol

Lab: Implementing DHCP

Exercise 1: Implementing DHCP

► Task 1: Install the Dynamic Host Configuration Protocol (DHCP) server role

1. Sign in to LON-SVR1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. In Server Manager, 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, select the **DHCP Server** check box.
7. In the Add Roles and Features Wizard, click **Add Features**, and then click **Next**.
8. On the **Select features** page, click **Next**.
9. On the **DHCP Server** page, click **Next**.
10. On the **Confirm installation selections** page, click **Install**.
11. On the **Installation progress** page, wait until the "Installation succeeded on lon-svr1.adatum.com" message appears, and then click **Close**.

► Task 2: Configure the DHCP scope and options

1. In the Server Manager Dashboard, click **Tools**, and then click **DHCP**.
2. In the DHCP console, expand and then right-click **lon-svr1.adatum.com**, and then click **Authorize**.
3. In the DHCP console, right-click **lon-svr1.adatum.com**, and then click **Refresh**.

Notice that the icons next to IPv4 IPv6 changes color from red to green, which means that the DHCP server has been authorized in Active Directory® Domain Services (AD DS).

4. In the DHCP console, in the navigation pane, click **lon-svr1.adatum.com**, expand and right-click **IPv4**, and then click **New Scope**.
5. In the New Scope Wizard, click **Next**.
6. On the **Scope Name** page, in the **Name** box, type **Branch Office**, and then click **Next**.
7. On the **IP Address Range** page, complete the page using the following information, and then click **Next**:
 - Start IP address: **172.16.0.100**
 - End IP address: **172.16.0.200**
 - Length: **16**
 - Subnet mask: **255.255.0.0**
8. On the **Add Exclusions and Delay** page, complete the page using the following information:
 - Start IP address: **172.16.0.190**
 - End IP address: **172.16.0.200**

9. Click **Add**, and then click **Next**.
10. On the **Lease Duration** page, click **Next**.
11. On the **Configure DHCP Options** page, click **Next**.
12. On the **Router (Default Gateway)** page, in the **IP address** box, type **172.16.0.1**, click **Add**, and then click **Next**.
13. On the **Domain Name and DNS Servers** page, click **Next**.
14. On the **WINS Servers** page, click **Next**.
15. On the **Activate Scope** page, click **Next**.
16. On the **Completing the New Scope Wizard** page, click **Finish**.

► **Task 3: Configure the client to use DHCP, and then test the configuration**

1. Sign in to **20410C-LON-CL1** as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On the **Start** page, type **Control Panel**, and then press Enter.
3. In Control Panel, under Network and Internet, click **View Network Status and Tasks**.
4. In the Network and Sharing Center window, click **Change adapter settings**.
5. In the Network Connections window, right-click **Ethernet**, and then click **Properties**.
6. In the Ethernet Properties window, click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
7. In the **Internet Protocol Version 4 (TCP/IPv4) Properties** dialog box, select the **Obtain an IP address automatically** radio button, select the **Obtain DNS server address automatically** radio button, click **OK**, and then click **Close**.
8. Right-click the **Start** button, and then click **Command Prompt**.
9. In the Command Prompt window, at the command prompt, type the following, and then press Enter:

```
ipconfig /renew
```

10. To test the configuration and verify that LON-CL1 has received an IP address from the DHCP scope, at a command prompt, type the following, and then press Enter:

```
ipconfig /all
```



Note: This command returns information such as IP address, subnet mask, and DHCP enabled status, which should be **Yes**.

► **Task 4: Configure a lease as a reservation**

1. In the Command Prompt window, at a command prompt, type the following, and then press Enter:

```
ipconfig /all
```

2. Write down the Physical Address of LON-CL1 network adapter.
3. Switch to LON-SVR1.
4. In the Server Manager dashboard, click **Tools**, and then click **DHCP**.

5. In the DHCP console, expand **lon-svr1.adatum.com**, expand **IPv4**, expand **Scope [172.16.0.0] Branch Office**, select and then right-click **Reservations**, and then click **New Reservation**.
6. In the New Reservation window:
 - In the Reservation Name field, type **LON-CL1**.
 - In the IP address field, type **172.16.0.155**.
 - In the MAC address field, type the physical address you wrote down in step 2.
 - Click **Add**, and then click **Close**.
7. Switch to **LON-CL1**.
8. In the Command Prompt window, at a command prompt, type the following, and then press Enter:

```
ipconfig /release
```

This causes LON-CL1 to release any currently leased IP addresses.
9. At a command prompt, type the following, and then press Enter:

```
ipconfig /renew
```

This causes LON-CL1 to lease any reserved IP addresses.
10. Verify that the IP address of LON-CL1 is now **172.16.0.155**.

Results: After completing this exercise, you should have implemented DHCP, configured DHCP scope and options, and configured a DHCP reservation.

► Prepare for the optional exercise

If you are going to complete the optional lab, revert the 20410C-LON-CL1 and 20410C-LON-SVR1 virtual machines by performing the following steps:

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20410C-LON-CL1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 1 through 3 for 20410C-LON-SVR1.
5. Start 20410C-LON-SVR1.

Exercise 2: Implementing a DHCP Relay Agent (Optional Exercise)

► Task 1: Install a DHCP relay agent

1. Sign in to LON-RTR as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. In Server Manager, click **Tools**, and then click **Routing and Remote Access**.
3. Add the DHCP relay agent to the router by performing the following steps:
 - a. In the navigation pane, expand **LON-RTR (local)**, expand **IPv4**, right-click **General**, and then click **New Routing Protocol**.
 - b. In the Routing protocols list, click **DHCP Relay Agent**, and then click **OK**.

► Task 2: Configure a DHCP relay agent

1. In the navigation pane, right-click **DHCP Relay Agent**, and then click **New Interface**.
2. In the New Interface for DHCP Relay Agent dialog box, click **Ethernet 2**, and then click **OK**.
3. In the DHCP Relay Agent Properties – Ethernet 2 Properties dialog box, click **OK**.
4. Right-click **DHCP Relay Agent**, and then click **Properties**.
5. In the DHCP Relay Agent Properties dialog box, in the Server address box, type **172.16.0.11**, click **Add**, and then click **OK**.
6. Close **Routing and Remote Access**.

► Task 3: Test the DHCP relay agent with a client



Note: To test how a client receives an IP address from the DHCP relay agent in another subnet, you need to create another DHCP scope.

1. Sign in to LON-SVR1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. In the **Desktop**, right-click the **PowerShell** icon and select **Run as administrator**.
3. At a Windows PowerShell command prompt, type the following, pressing Enter after each line:

```
Add-WindowsFeature -IncludeManagementTools dhcp
netsh dhcp add securitygroups
Restart-service dhcpserver
Add-DhcpServerInDC LON-SVR1 172.16.0.11
Add-DhcpServerv4Scope -Name "Branch Office 2" -StartRange 10.10.0.100 -EndRange 10.10.0.200 -SubnetMask 255.255.0.0
Add-DhcpServerv4ExclusionRange -ScopeID 10.10.0.0 -StartRange 10.10.0.190 -EndRange 10.10.0.200
Set-DhcpServerv4OptionValue -Router 10.10.0.1
Set-DhcpServerv4Scope -ScopeID 10.10.0.0 -State Active
```

4. To test the client, switch to **LON-CL2**.
5. On the Start screen, type **Control Panel**, and then press Enter.
6. Under Network and Internet, click **View network status and tasks**.
7. In the Network and Sharing Center window, click **Change Adapter Settings**, right-click **Ethernet**, and then click **Properties**.

8. In the Ethernet Properties window, click **Internet Protocol Version 4 (TCP/IPv4)** and then click **Properties**.
9. In the Internet Protocol Version 4 (TCP/IPv4) Properties dialog box, click **Obtain an IP address automatically**, click **Obtain DNS server address automatically**, click **OK**, and then click **Close**.
10. Right-click the Start button and then click Command Prompt.
11. In the Command Prompt window, at a command prompt, type the following, and then press Enter:

```
ipconfig /renew
```
12. Verify that IP address and DNS server settings on LON-CL2 are obtained from DHCP Server scope **Branch Office 2**, installed on **LON-SVR1**.



Note: The IP address should be in the following range: **10.10.0.100/16** to **10.10.0.200/16**.

Results: After completing this exercise, you should have implemented a DHCP relay agent.

► Prepare for the next module

After you finish the lab, revert the virtual machines back 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 **20410C-LON-DC1**, and then click **Revert**.
3. In the Revert Virtual Machine dialog box, click **Revert**.
4. Repeat steps 2 and 3 for 20410C-LON-SVR1, 20410C-LON-RTR, and 20410C-LON-CL2.

Module 7: Implementing DNS

Lab: Implementing DNS

Exercise 1: Installing and Configuring DNS

► Task 1: Configure LON-SVR1 as a domain controller without installing the Domain Name System (DNS) server role

1. On LON-SVR1, in the Server Manager console, 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, ensure that **LON-SVR1.Adatum.com** is selected, and then click **Next**.
5. On the **Select server roles** page, select **Active Directory Domain Services**.
6. When Add Roles and Features Wizard appears, click **Add Features**, and then click **Next**.
7. On the **Select features** page, click **Next**.
8. On the **Active Directory Domain Services** page, click **Next**.
9. On the **Confirm installation selections** page, click **Install**.
10. On the **Installation progress** page, when the **Installation succeeded** message appears, click **Close**.
11. In the Server Manager console, on the navigation page, click **AD DS**.
12. On the title bar where **Configuration required for Active Directory Domain Services at LON-SVR1** is visible, click **More**.
13. On the **All Server Task Details and Notifications** page, click **Promote this server to a domain controller**.
14. In the Active Directory Domain Services Configuration Wizard, on the **Deployment Configuration** page, ensure that **Add a domain controller to an existing domain** is selected, and then click **Next**.
15. On the **Domain Controller Options** page, deselect the **Domain Name System (DNS) server** check box, and leave the **Global Catalog (GC)** check box selected.
16. Type **Pa\$\$w0rd** in both text fields, and then click **Next**.
17. On the **Additional Options** page, click **Next**.
18. On the **Paths** page, click **Next**.
19. On the **Review Options** page, click **Next**.
20. On the **Prerequisites Check** page, click **Install**.
21. On the **You're about to be signed out** blue bar, click **Close**.



Note: The LON-SVR1 server automatically restarts as part of the procedure.

22. After LON-SVR1 restarts, sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.

► Task 2: Create and configure Contoso.com zone on LON-DC1

1. On the LON-DC1 virtual machine, in the Server Manager console, click **Tools**, and then click **DNS**.
2. Expand **LON-DC1**, right-click **Forward Lookup Zones**, and then select **New Zone**.
3. In the New Zone Wizard, on the **Welcome to the New Zone Wizard** page, click **Next**.
4. On the **Zone Type** page, deselect the **Store the zone in Active Directory** check box, and then click **Next**.
5. On the **Zone Name** page, type **Contoso.com**, and then click **Next**.
6. On the **Zone File** page, click **Next**.
7. On the **Dynamic Update** page, click **Next**.
8. On the **Completing the New Zone Wizard** page, click **Finish**.
9. Expand **Forward Lookup Zones**, and then select and right-click **contoso.com** zone and click **New Host (A or AAAA)**.
10. In the New Host window, in the **Name** textbox type **www**.
11. In the **IP address** box type **172.16.0.100**.
12. Click **Add Host**.
13. Click **OK** and then click **Done**.
14. Leave **DNS Manager** console open.

► Task 3: Review configuration settings on the existing DNS server to confirm root hints

1. On LON-DC1, in the DNS Manager console, click and then right-click **LON-DC1**, and then click **Properties**.
2. In the **LON-DC1 Properties** dialog box, click the **Root hints** tab. Ensure that root hints servers display.
3. Click the **Forwarders** tab. Ensure that the list displays no entries, and that the **Use root hints if no forwarders are available** option is selected.
4. Click **Cancel**.
5. Close the DNS Manager console.
6. In the taskbar, click the **Windows PowerShell** icon.
7. In Windows PowerShell, type the following cmdlets, pressing Enter after each, and observe the output returned:

```
Get-DnsServerRootHint  
Get-DnsServerForwarder
```

Note that both cmdlets are the respective Windows PowerShell equivalents of the DNS Console actions performed in steps 2 and 3 above.

► Task 4: Add the DNS server role for the branch office on the domain controller

1. On LON-SVR1, in the Server Manager console, 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, ensure that **LON-SVR1.Adatum.com** is selected, and then click **Next**.
5. On the **Select server roles** page, select **DNS Server**.
6. When the Add Roles and Features Wizard appears, click **Add Features**, and then click **Next**.
7. On the **Select Features** page, click **Next**.
8. On the **DNS Server** page, click **Next**.
9. On the **Confirm installation selections** page, click **Install**.
10. On the **Installation progress** page, when the "Installation succeeded" message appears, click **Close**.

► **Task 5: Verify replication of the Adatum.com Active Directory®-integrated zone**

1. On LON-SVR1, in the Server Manager console, click **Tools**.
2. On the list of tools, click **DNS**.
3. In the DNS Manager console, expand **LON-SVR1**, and then expand **Forward Lookup Zones**.
This container is probably empty.
4. Switch back to **Server Manager**, click **Tools**, and then click **Active Directory Sites and Services**.
5. In the **Active Directory Sites and Services** console, expand **Sites**, expand **Default-First-Site-Name**, expand **Servers**, expand **LON-DC1**, and then click **NTDS Settings**.
6. In the right pane, right-click the **LON-SVR1** replication connection, and select **Replicate Now**.



Note: If you receive an error message, proceed to the next step, and then retry this step after 3-4 minutes. If this retry fails, wait a few more minutes, and then try again.

7. In the navigation pane, expand **LON-SVR1**, and then click **NTDS Settings**.
8. In the right pane, right-click the **LON-DC1** replication connection, click **Replicate Now**, and then click **OK**.
9. Switch back to the DNS Manager console, right-click **Forward Lookup Zones**, and then click **Refresh**.
10. Ensure that both the **_msdcs.Adatum.com** and **Adatum.com** containers display.
11. Close DNS Manager.

► **Task 6: Use Windows PowerShell commands to test non-local resolution**

1. On LON-SVR1, on the taskbar, click the **Windows PowerShell** icon.
2. In Windows PowerShell, type the following cmdlet, and then press Enter:

```
Get-DnsClient
```

3. Note the entries labeled **Ethernet** in the InterfaceAlias column. In the Interface Index column, note the Interface Index number that is in the same row as Ethernet and IPv4. Write this number here:

4. In Windows PowerShell, type the following cmdlet, where *X* is the specific Interface Index number you wrote down in the last step, and then press Enter:

```
Set-DnsClientServerAddress -InterfaceIndex X -ServerAddress 0.0.0.0
```

5. In Windows PowerShell, type the following, and then press Enter:

```
Resolve-DNSName www.contoso.com
```

You should see an error message.

6. In Windows PowerShell, type the following, and then press Enter:

```
nslookup
```

7. At the nslookup > prompt, type the following and then press Enter:

```
www.contoso.com
```

You should see the following reply: ****** Unknown can't find www.contoso.com. No response from server.**

8. Type the following, and then press Enter:

```
Exit
```

Leave the Windows PowerShell window open.

► Task 7: Configure Internet name resolution to forward to the head office

1. At the Windows PowerShell prompt, type the following cmdlet, and then press Enter:

```
Set-DnsServerForwarder -IPAddress '172.16.0.10' -PassThru
```

2. At the Windows PowerShell prompt, type the following cmdlet, and then press Enter:

```
Restart-Computer
```

► Task 8: Use Windows PowerShell to confirm name resolution

1. Sign in to **LON-SVR1** as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On LON-SVR1, switch to a Windows PowerShell window.
3. Type the following cmdlet, and then press Enter:

```
nslookup www.contoso.com
```

Ensure that you receive an IP address for this host as a non-authoritative answer.

4. Close Windows PowerShell.

Results: After completing this exercise, you should have installed and configured DNS on 20410C-LON-SVR1.

Exercise 2: Creating Host Records in DNS

► Task 1: Configure a client to use LON-SVR1 as a DNS server

1. On LON-CL1, sign in as **Adatum\Administrator** using the password **Pa\$\$w0rd**.
2. On the Start screen, type **Control Panel**, and then press Enter.
3. In Control Panel, click **View network status and tasks**.

4. Click **Change adapter settings**.
5. Right-click **Ethernet**, and then click **Properties**.
6. In the **Ethernet Properties** dialog box, click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
7. In the **preferred DNS server** box, overwrite the IP address for **preferred DNS server** with **172.16.0.11**, click **OK**, and then click **Close**.

► **Task 2: Create several host records for web apps in the Adatum.com domain**

1. On LON-DC1, in the Server Manager console, click **Tools**, and then click **DNS**.
2. In the DNS Manager console, expand **LON-DC1**, expand **Forward Lookup Zones**, and then click **Adatum.com**.
3. Right-click **Adatum.com**, and then click **New Host (A or AAAA)**.
4. In the **New Host** window, configure the following settings:
 - Name: **www**
 - IP address: **172.16.0.200**
5. Click **Add Host**, and then click **OK**.
6. In the **New Host** window, configure the following settings:
 - Name: **ftp**
 - IP address: **172.16.0.201**
7. Click **Add Host**, click **OK**, and then click **Done**.

► **Task 3: Verify replication of new records to LON-SVR1**

1. On LON-SVR1, in the Server Manager console, click **Tools**, and then click **DNS**.
2. In the DNS Manager console, expand **LON-SVR1**, expand **Forward Lookup Zones**, and then click **Adatum.com**.
3. Ensure that both **www** and **ftp** resource records display. It might take several minutes for the records to display.



Note: If the **www** and **ftp** resource records do not display within several minutes, right-click **Adatum.com**, and then click **Refresh**.

► **Task 4: Use the ping command to locate new records from LON-CL1**

1. On LON-CL1, on the taskbar, right-click the **Windows** icon, and then click **Run**.
2. In the Run pop-up window, in the **Open** text box, type **cmd**, and then press Enter.
3. In the Command Prompt window, at a command prompt, type **ping www.adatum.com**, and then press Enter.
4. Ensure that the name resolves to **172.16.0.200**.



Note: You will not receive replies.

5. At a command prompt, type **ping ftp.adatum.com**, and then press Enter.

6. Ensure that name resolves to **172.16.0.201**. (You will not receive replies.)
7. Leave the Command Prompt window open.

Results: After completing this exercise, you should have configured DNS records.

Exercise 3: Managing the DNS Server Cache

► Task 1: Use the ping command to locate an Internet record from LON-CL1

1. On LON-CL1, in the Command Prompt window, at a command prompt, type **ping www.contoso.com**, and then press Enter.
2. Ping does not work. Ensure that the name resolves to the IP address 172.16.0.100.
3. Leave the Command Prompt window open.

► Task 2: Update an Internet record to point to the LON-DC1 IP address

1. On LON-DC1, open **DNS Manager**.
2. In the DNS Manager console, expand **LON-DC1**, expand **Forward Lookup Zones**, and then click **contoso.com**.
3. In the right pane, right-click **www**, and then click **Properties**.
4. Change the IP address to **172.16.0.10**, and then click **OK**.
5. Switch back to LON-CL1.
6. In the Command Prompt window, at a command prompt, type the following, and then press Enter:

```
ping www.contoso.com
```

Note that ping does not work, and that the old IP address (which is 172.16.0.100) is still displayed.

► Task 3: Examine the content of the DNS cache

1. Switch to LON-SVR1.
2. In the Server Manager console, click **Tools**, and then click **DNS**.
3. Click **LON-SVR1**, click the **View** menu, and then click **Advanced**.
4. Expand **LON-SVR1**, expand the **Cached Lookups** node, expand **.(root)**, expand **com**, and then click **contoso**.
5. In the right pane, examine the cached content and note that the **www** record has the IP address: **172.16.0.100**.
6. Switch to LON-CL1.
7. In the Command Prompt window, at a command prompt, type **ipconfig /displaydns**, and then press Enter.
8. Look for cached entries and notice that **www.contoso.com** is resolving to **172.16.0.100**.

► **Task 4: Clear the cache, and retry the ping command**

1. On LON-SVR1, on the taskbar, click the **Windows PowerShell** icon.
2. At the Windows PowerShell prompt, type **Clear-DNSServerCache**, and then press Enter.
3. Type **y**, and then press Enter.
4. Switch to LON-CL1.
5. In a Command Prompt window, at a command prompt, type **ping www.contoso.com**, and then press Enter.

The result still returns the old IP address.

6. In the Command Prompt window, at a command prompt, type **ipconfig /flushdns**, and then press Enter.
7. In the Command Prompt window, type **ping www.contoso.com**, and then press Enter.
Ping now should work on address **172.16.0.10**.

Results: After completing this exercise, you should have examined the DNS server cache.

► **Prepare for the next module**

After you finish the lab, revert the virtual machines to their initial state.

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20410C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20410C-LON-SVR1** and **20410C-LON-CL1**.

Module 8: Implementing IPv6

Lab: Implementing IPv6

Exercise 1: Configuring an IPv6 Network

► Task 1: Verify IPv4 routing

1. On LON-SVR2, on the taskbar, click the **Windows PowerShell** icon.
2. At the Windows PowerShell prompt, type **ping lon-dc1**, and then press Enter.
Notice that there are four replies from 172.16.0.10.
3. Type **ipconfig**, and then press Enter.
Verify that the only IPv6 address listed is a link-local address that cannot be routed.
4. Type **Get-NetIPAddress**, and then press Enter.
Notice that **Get-NetIPAddress** cmdlet returns a link-local IPv6 address.

► Task 2: Disable IPv6 on LON-DC1

1. On LON-DC1, in Server Manager, click **Local Server**.
2. In the Properties window, next to the **Ethernet** section, click **172.16.0.10, IPv6 enabled**.
3. In the Network Connections window, right-click **Ethernet**, and then click **Properties**.
4. In the **Ethernet Properties** dialog box, clear the **Internet Protocol Version 6 (TCP/IPv6)** check box, and then click **OK**.
5. Close the Network Connections window.
6. In Server Manager, verify that **Ethernet** lists only **172.16.0.10**. You may need to refresh the view.
LON-DC1 is now an IPv4-only host.

► Task 3: Disable IPv4 on LON-SVR2

1. On LON-SVR2, in Server Manager, click **Local Server**.
2. In the **Local Server Properties** dialog box, next to Ethernet, click **10.10.0.11, IPv6 enabled**.
3. In the Network Connections window, right-click **Ethernet**, and then click **Properties**.
4. In the **Ethernet Properties** dialog box, clear the **Internet Protocol Version 4 (TCP/IPv4)** check box, and then click **OK**.
5. Close the Network Connections window.
6. In Server Manager, verify that **Ethernet** now lists only **IPv6 enabled**. You may need to refresh the view.
LON-SVR2 is now an IPv6-only host.

► Task 4: Configure an IPv6 network on LON-RTR

1. On LON-RTR, on the taskbar, click the **Windows PowerShell** icon.
2. Configure a network address that will be used on the IPv6 network. At the Windows PowerShell prompt, type the following cmdlet, and then press Enter:

```
New-NetRoute -InterfaceAlias " Ethernet 2" -DestinationPrefix 2001:db8:0:1::/64
-Publish Yes
```

3. Allow clients to obtain the IPv6 network address automatically from LON-RTR. At the Windows PowerShell prompt, type the following cmdlet, and then press Enter:

```
Set-NetIPInterface -InterfaceAlias "Ethernet 2" -AddressFamily IPv6 -Advertising
Enabled
```

4. Type **ipconfig**, and then press Enter.

Notice that Ethernet 2 now has an IPv6 address on the 2001:db8:0:1::/64 network. This address is used for communication on the IPv6-only network.

► Task 5: Verify IPv6 on LON-SVR2

1. On LON-SVR2, on the taskbar, click the **Windows PowerShell** icon.
2. At the Windows PowerShell prompt, type **ipconfig**, and then press Enter.

Notice that the Ethernet now has an IPv6 address on the 2001:db8:0:1::/64 network. The network address was obtained from the router through stateless configuration.

Results: After completing the exercise, you will have configured an IPv6-only network.

Exercise 2: Configuring an ISATAP Router

► Task 1: Add an ISATAP host record to DNS

1. On LON-DC1, in Server Manager, click **Tools**, and then click **DNS**.
2. In DNS Manager, expand **LON-DC1**, expand **Forward Lookup Zones**, and then click **Adatum.com**.
3. Right-click **Adatum.com**, and then click **New Host (A or AAAA)**.
4. In the New Host window, in the **Name** box, type **ISATAP**.
5. In the **IP address** box, type **172.16.0.1**, and then click **Add Host**. ISATAP clients resolve this host name to find the ISATAP router.
6. Click **OK** to clear the success message.
7. Click **Done** to close the New Host window.
8. Close **DNS Manager**.

► Task 2: Enable the ISATAP router on LON-RTR

1. On LON-RTR, configure the IP address of the **Ethernet** adapter as the ISATAP router. At the Windows PowerShell prompt, type the following cmdlet, and then press Enter:

```
Set-NetIsatapConfiguration -Router 172.16.0.1
```

2. Type the following command, and then press Enter:

```
Get-NetIPAddress | Format-Table InterfaceAlias,InterfaceIndex,IPv6Address
```

3. Record the InterfaceIndex of the ISATAP interface that has an IPv6 address that includes **172.16.0.1**.

Interface index:	
------------------	--

4. Type the following command, and then press Enter:

```
Get-NetIPInterface -InterfaceIndex IndexYouRecorded -PolicyStore ActiveStore |  
Format-List
```

5. Verify that Forwarding is enabled for the interface and that Advertising is disabled.
6. The ISATAP interface for an ISATAP router must have forwarding enabled and advertising enabled. Type the following command, and then press Enter:

```
Set-NetIPInterface -InterfaceIndex IndexYouRecorded -Advertising Enabled
```

7. Create a new IPv6 network that will be used for the ISATAP network. Type the following command, and then press Enter:

```
New-NetRoute -InterfaceIndex IndexYouRecorded -DestinationPrefix 2001:db8:0:2::/64  
-Publish Yes
```

8. View the IP address configuration for the ISATAP interface. Type the following command, and then press Enter:

```
Get-NetIPAddress -InterfaceIndex IndexYouRecorded
```

9. Verify that an IPv6 address is listed on the 2001:db8:0:2::/64 network.

► Task 3: Remove ISATAP from the Global Query Block List

1. On LON-DC1, at the Windows PowerShell prompt, type the following command, and then press Enter.

```
dnscmd /config /globalqueryblocklist wpad
```

2. At the Windows PowerShell prompt, type **Restart-Service DNS -Verbose**, and then press Enter.
3. Type **ping isatap**, and then press Enter.

The name should resolve and you should receive four replies from 172.16.0.1.

► Task 4: Enable LON-DC1 as an ISATAP client

1. On LON-DC1, at the Windows PowerShell prompt, type the following command, and then press Enter:

```
Set-NetIsatapConfiguration -State Enabled
```

2. Type **ipconfig**, and then press Enter.

3. Verify that the Tunnel adapter for ISATAP has an IPv6 address on the 2001:db8:0:2/64 network.
Notice that this address includes the IPv4 address of LON-DC1.

► Task 5: Test connectivity

1. On LON-SVR2, at the Windows PowerShell prompt, type the following command, and then press Enter:

```
ping 2001:db8:0:2:0:5efe:172.16.0.10
```

2. In Server Manager, if necessary, click **Local Server**.
3. In the **Local Server Properties** dialog box, next to Ethernet, click **IPv6 enabled**.
4. In the Network Connections window, right-click **Ethernet**, and then click **Properties**.
5. In the **Ethernet Properties** dialog box, click **Internet Protocol Version 6 (TCP/IPv6)**, and then click **Properties**.
6. In the **Internet Protocol Version 6 (TCP/IPv6) Properties** dialog box, click **Use the following DNS server addresses**.
7. In the **Preferred DNS server** box, type **2001:db8:0:2:0:5efe:172.16.0.10**, and then click **OK**.
8. In the **Ethernet Properties** dialog box, click **Close**.
9. Close the Network Connections window.
10. At the Windows PowerShell prompt, type **ping LON-DC1**, and then press Enter.
Notice that four replies are received from LON-DC1.



Note: A ping from LON-DC1 to LON-SVR2 does not respond because the firewall configuration on LON-SVR2 blocks ping requests.

Results: After completing this exercise, you will have configured an ISATAP router on LON-RTR to allow communication between an IPv6-only network and an IPv4-only network.

► Prepare for the next module

After you finish the lab, revert the virtual machines back 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 **20410C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20410C-LON-SVR1**, **20410C-LON-RTR** and **20410C-LON-SVR2**.

Module 9: Implementing Local Storage

Lab: Implementing Local Storage

Exercise 1: Installing and Configuring a New Disk

► Task 1: Initialize a new disk

1. Sign in to **LON-SVR1** with the username **Adatum\Administrator** and the password **Pa\$\$w0rd**.
2. In Server Manager, click the **Tools** menu, and then click **Computer Management**.
3. In the Computer Management console, under the Storage node, click **Disk Management**.
4. In the Disks pane, right-click **Disk2**, and then click **Online**.
5. Right-click **Disk2**, and then click **Initialize Disk**.
6. In the **Initialize Disk** dialog box, select the **Disk 2** check box, click **GPT (GUID Partition Table)**, and then click **OK**.

► Task 2: Create and format two simple volumes on the disk

1. In the Computer Management console, in Disk Management, right-click the black marked box right of Disk 2, and then click **New Simple Volume**.
2. In the New Simple Volume Wizard, on **Welcome to the New Simple Volume Wizard** page, click **Next**.
3. On the **Specify Volume Size** page, in the **Simple volume size MB** field, type **4000**, and then click **Next**.
4. On **Assign Drive Letter or Path** page, ensure that the **Assign the following drive letter** check box is selected, and that **F** is selected in from the drop-down menu, and then click **Next**.
5. On the **Format Partition** page, from the **File system** drop-down menu, click **NTFS**, in the **Volume label** text box, type **Volume1**, and then click **Next**.
6. On **Completing the New Simple Volume Wizard** page, click **Finish**.
7. In the Disk Management window, right-click the black box right of Disk 2, and then click **New Simple Volume**.
8. In the New Simple Volume Wizard, on **Welcome to the New Simple Volume Wizard** page, click **Next**.
9. On the **Specify Volume Size** page, in the **Simple volume size in MB** field, type **5000**, and then click **Next**.
10. On the **Assign Drive Letter or Path** page, ensure that the **Assign the following drive letter** check box is selected, verify that **G** is listed as the drive letter, and then click **Next**.
11. On the **Format Partition** page, from the **File system** drop-down menu, click **ReFS**, in the **Volume label** text box, type **Volume2**, and then click **Next**.
12. On the **Completing the New Simple Volume Wizard** page, click **Finish**.

► Task 3: Verify the drive letter in a File Explorer window

1. On the taskbar, open a File Explorer window, expand **This PC**, and then click **Volume1 (F:)**.
2. In File Explorer, click **Volume2 (G:)**, right-click **Volume2 (G:)**, point to **New**, and then click **Folder**.
3. In the **New folder** field, type **Folder1**, and then press Enter.

Results: After you complete this lab, you should have initialized a new disk, created two simple volumes, and then formatted them. You should also have verified that the drive letters you assigned are available in File Explorer.

Exercise 2: Resizing Volumes

► Task 1: Shrink Volume1

1. On LON-SVR1, switch to the Computer Management console.
2. In the Computer Management console, in Disk Management, in the middle-pane, right-click **Volume1 (F:)**, and then click **Shrink Volume**.
3. In the Shrink F: window, in the **Enter the amount of space to shrink in MB** field, type **1000**, and then click **Shrink**.

► Task 2: Extend Volume2

1. On LON-SVR1, in Disk Management, in the middle-pane, right-click **Volume2 (G:)**, and then click **Extend Volume**.
2. In Extend Volume Wizard, on the **Welcome to the Extended Volume Wizard** page, click **Next**.
3. On the **Select Disks** page, in the **Select the amount of space in MB** field, type **1000**, and then click **Next**.
4. On the **Completing the Extended Volume Wizard** page, click **Finish**.
5. In a File Explorer window, click **Volume2 (G:)**, and verify that **Folder1** is available on the volume.

Results: After this lab, you should have made one volume smaller, and extended another.

Exercise 3: Configuring a Redundant Storage Space

► Task 1: Create a storage pool from five disks that are attached to the server

1. On LON-SVR1, on the taskbar, click the **Server Manager** icon.
2. In Server Manager, in the left 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 menu, 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 the following physical disks, and then click **Next**:
 - **PhysicalDisk3**
 - **PhysicalDisk4**
 - **PhysicalDisk5**
 - **PhysicalDisk6**
 - **PhysicalDisk7**
7. On the **Confirm selections** page, click **Create**.
8. On the **View results** page, wait until the task completes, and then click **Close**.

► **Task 2: Create a three-way mirrored virtual disk**

1. On LON-SVR1, in Server Manager, in the Storage Spaces pane, click **StoragePool1**.
2. In the VIRTUAL DISKS pane, click **TASKS**, and then from the **TASKS** drop-down menu, 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 storage pool** page, click **StoragePool1**, and then click **Next**.
5. On the **Specify the virtual disk name** page, in the **Name** box, type **Mirrored Disk**, and then click **Next**.
6. On the **Select the storage layout** page, in the **Layout** list, click **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 **Specify 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 task completes.
12. Ensure that the **Create a volume when this wizard closes** check box is selected, and then click **Close**.
13. In the New Volume Wizard window, on the **Before you begin** page, click **Next**.
14. On the **Select the server and disk** page, in the Disk pane, click the **Mirrored Disk** virtual disk, and then click **Next**.
15. On the **Specify the size of the volume** page, click **Next** to confirm the default selection.
16. On the **Assign to a drive letter or folder** page, in the **Drive letter** drop-down menu, ensure that **H** is selected, and then click **Next**.
17. On the **Select file system settings** page, in the **File system** drop-down menu, click **ReFS**, in the **Volume label** box, type **Mirrored Volume**, and then click **Next**.
18. On the **Confirm selections** page, click **Create**.
19. On the **Completion** page, wait until the creation completes, and then click **Close**.

► Task 3: Copy a file to the volume, and verify that it is visible in File Explorer

1. Click to the Start screen, type **command prompt**, and then press Enter.
2. In the Command Prompt window, at the command prompt, type the following command, and then press Enter:

```
Copy C:\windows\system32\write.exe H:\
```

3. Close the Command Prompt window.
4. On the taskbar, click the **File Explorer** icon.
5. In the File Explorer window, click **Mirrored Volume (H:)**.
6. Verify that write.exe displays in the file list.
7. Close File Explorer.

► Task 4: Remove a physical drive

1. On Host machine, in Hyper-V Manager®, in the Virtual Machines pane, right-click **0410C-LON-SVR1**, and then click **Settings**.
2. In Settings for 20410C-LON-SVR1, in the Hardware pane, click Hard Drive that begins with **20410C-LON-SVR1-Disk5**.
3. In the Hard Drive pane, click **Remove**, and then click **OK**. Click **Continue**.

► Task 5: Verify that the write.exe file is still accessible

1. Switch to LON-SVR1.
2. On the taskbar, click the **File Explorer** icon.
3. In the File Explorer window, click **Mirrored Volume (H:)**.
4. In the file list pane, verify that **write.exe** is still available.
5. Close File Explorer.
6. In Server Manager, in the STORAGE POOLS pane, on the menu bar, click the **Refresh "Storage Pools"** button.

Notice the warning that displays next to Mirrored Disk.

7. In the VIRTUAL DISK pane, right-click **Mirrored Disk**, and then click **Properties**.
8. In the **Mirrored Disk Properties** dialog box, in the left pane, click **Health**.

Notice that the Health Status indicates a Warning. The Operational Status should indicate Incomplete, Unknown or Degraded.

9. Click **OK** to close the **Mirrored Disk Properties** dialog box.

► Task 6: Add a new disk to the storage pool and remove a broken disk

1. Switch to **LON-SVR1**.
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 click **Add Physical Disk**.
4. In the Add Physical Disk window, click **PhysicalDisk8 (LON-SVR1)**, and then click **OK**.
5. Click **Windows Powershell** on the Task Bar.

6. Type **Get-PhysicalDisk**, and press Enter.
7. Note the FriendlyName for the disk that shows an OperationalStatus of Lost Communication.
8. Type **\$Disk = Get-PhysicalDisk -FriendlyName *diskname*** and press ENTER. Replace *diskname* with the name of the disk you noted in Step 7.
9. Type **Remove-PhysicalDisk -PhysicalDisks \$disk -StoragePoolFriendlyName StoragePool1** and press Enter.
10. Type **Y**, and then press Enter.
11. If you get a warning that the disk cannot be removed, wait five minutes and then run the last command again. It can take some time for the mirrored disk to be resynchronized after a disk has been removed and another disk added. If you still cannot remove the disk after five minutes, restart LON-SVR1, sign in as **Adatum\Administrator** using the password **Pa\$\$w0rd**, and repeat steps 5-10.
12. In Server Manager, in the STORAGE POOLS pane, on the menu bar, click the **Refresh "Storage Pools"** button to see the warnings disappear.

Results: After completing this lab, you should have created a storage pool and added five disks to it. Then you should have created a three-way mirrored, thinly provisioned virtual disk from the storage pool. You should have also copied a file to the new volume and verified that it is accessible. Next, after removing a physical drive, you should have verified that the virtual disk was still available and could be accessed. Finally, you should have added another physical disk to the storage pool.

► Prepare for the next module

After you finish the lab, revert the virtual machines back 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 **20410C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20410C-LON-SVR1**.

Module 10: Implementing File and Print Services

Lab: Implementing File and Print Services

Exercise 1: Creating and Configuring a File Share

► Task 1: Create the folder structure for the new share

1. On LON-SVR1, on the taskbar, click the **File Explorer** icon.
2. In File Explorer, in the navigation pane, expand **This PC**, and then click **Allfiles (E:)**.
3. On the menu toolbar, click **Home**, click **New folder**, type **Data**, and then press Enter.
4. Double-click the **Data** folder.
5. On the menu toolbar, click **Home**, click **New folder**, type **Development**, and then press Enter.
6. Repeat step 5 for the following new folder names:
 - **Marketing**
 - **Research**
 - **Sales**

► Task 2: Configure NTFS permissions on the folder structure

1. In File Explorer, navigate to drive E, right-click the **Data** folder, and then click **Properties**.
2. In the **Data Properties** dialog box, click **Security**, and then click **Advanced**.
3. In the **Advanced Security Settings for Data** dialog box, click **Disable Inheritance**.
4. In the **Block Inheritance** dialog box, click **Convert inherited permissions into explicit permissions on this object**.
5. Click **OK** to close the **Advanced Security Settings for Data** dialog box.
6. Click **OK** to close the **Data Properties** dialog box.
7. In File Explorer, double-click the **Data** folder.
8. Right-click the **Development** folder, and then click **Properties**.
9. In the **Development Properties** dialog box, click **Security**, and then click **Advanced**.
10. In the **Advanced Security Settings for Development** dialog box, click **Disable Inheritance**.
11. In the **Block Inheritance** dialog box, click **Convert inherited permissions into explicit permissions on this object**.
12. Remove the two permissions entries for Users (LON-SVR1\Users), and then click **OK**.
13. On the **Security** tab, click **Edit**.
14. In the **Permissions for Development** dialog box, click **Add**.
15. Type **Development**, click **Check names**, and then click **OK**.
16. In the **Permissions for Development** dialog box, under **Allow**, select **Modify** permission.
17. Click **OK** to close the **Permissions for Development** dialog box.
18. Click **OK** to close the **Development Properties** dialog box.
19. Repeat steps 8 through 18 for the **Marketing**, **Research**, and **Sales** folders, assigning Modify permissions to the **Marketing**, **Research**, and **Sales** groups for their respective folders.

► Task 3: Create the shared folder

1. In File Explorer, navigate to drive E, right-click the **Data** folder, and then click **Properties**.
2. In the **Data Properties** dialog box, click the **Sharing** tab, and then click **Advanced Sharing**.
3. In the **Advanced Sharing** dialog box, select **Share this folder**, and then click **Permissions**.
4. In the **Permissions for Data** dialog box, click **Add**.
5. Type **Authenticated Users**, click **Check names**, and then click **OK**.
6. In the **Permissions for Data** dialog box, click **Authenticated Users**, and then under **Allow**, select **Change** permission.
7. Click **OK** to close the **Permissions for Data** dialog box.
8. Click **OK** to close the **Advanced Sharing** dialog box.
9. Click **Close** to close the **Data Properties** dialog box.

► Task 4: Test access to the shared folder

1. Sign in to LON-CL1 as **Adatum\Bernard** with the password **Pa\$\$w0rd**.



Note: Bernard is a member of the Development group.

2. On the Start screen, click **Desktop**.
3. On the taskbar, click the **File Explorer** icon.
4. In File Explorer, in the address bar, type **\\LON-SVR1\Data**, and then press Enter.
5. Double-click the **Development** folder.



Note: Bernard should have access to the Development folder.

6. Attempt to access the **Marketing**, **Research**, and **Sales** folders.
NTFS permissions on these folders prevents you from doing this.



Note: Bernard can still see the other folders, even though he does not have access to their contents.

7. Sign out of LON-CL1.

► Task 5: Enable access-based enumeration

1. Switch to LON-SVR1.
2. On the taskbar, click the **Server Manager** icon.
3. In Server Manager, in the navigation pane, click **File and Storage Services**.
4. In the File and Storage Services window, in the navigation pane, click **Shares**.
5. In the **Shares** pane, right-click **Data**, and then click **Properties**.
6. In the **Data Properties** dialog box, click **Settings**, and then select **Enable access-based enumeration**.

7. Click **OK** to close the **Data Properties** dialog box.
8. Close Server Manager.

► **Task 6: Test access to the share**

1. Sign in to LON-CL1 as **Adatum\Bernard** with the password **Pa\$\$w0rd**.
2. Click the **Desktop** tile.
3. On the taskbar, click the **File Explorer** icon.
4. In File Explorer, in the address bar, type **\\LON-SVR1\Data**, and then press Enter.



Note: Bernard can now view only the Development folder, the folder for which he has been assigned permissions.

5. Double-click the **Development** folder.



Note: Bernard should have access to the Development folder.

6. Sign out of LON-CL1.

► **Task 7: Disable Offline Files for the share**

1. Switch to LON-SVR1.
2. On the taskbar, click the **File Explorer** icon.
3. In File Explorer, navigate to drive E, right-click the **Data** folder, and then click **Properties**.
4. In the **Data Properties** dialog box, click the **Sharing** tab, click **Advanced Sharing**, and then click **Caching**.
5. In the **Offline Settings** dialog box, click **No files or programs from the shared folder are available offline**, and then click **OK**.
6. Click **OK** to close the **Advanced Sharing** dialog box.
7. Click **Close** to close the **Data Properties** dialog box.

Results: After completing this exercise, you will have created a new shared folder for use by multiple departments.

Exercise 2: Configuring Shadow Copies

► **Task 1: Configure shadow copies for the file share**

1. On LON-SVR1.
2. Open File Explorer.
3. Navigate to drive E, right-click **Allfiles (E:)**, and then click **Configure Shadow Copies**.
4. In the **Shadow Copies** dialog box, click drive **E**, and then click **Enable**.
5. In the **Enable Shadow Copies** dialog box, click **Yes**.
6. In the drive **Shadow Copies** dialog box, click **Settings**.

7. In the **Settings** dialog box, click **Schedule**.
This opens the drive **E:** dialog box.
8. In drive **E:** dialog box, change Schedule Task to **Daily**, change Start time to **12:00 AM** and then click **Advanced**.
9. In the **Advanced Schedule Options** dialog box, select **Repeat task**, and then set the frequency to **every 1 hours**.
10. Select **Time**, and change the time value to **11:59 PM**.
11. Click **OK** twice.
12. Click **OK** to close the **Settings** dialog box.
13. Leave the drive **Shadow Copies** dialog box open.

► **Task 2: Create multiple shadow copies of a file**

1. On LON-SVR1, open File Explorer.
2. Navigate to **E:\Data\Development**.
3. On the menu toolbar, click **Home**, click **New item**, and then click **Text Document**.
4. Type **Report**, and then press Enter.
5. Switch back to the **Shadow Copies** dialog box; it should still be opened on the **Shadow Copies** tab.
6. Click **Create Now**.

► **Task 3: Recover a deleted file from a shadow copy**

1. On LON-SVR1, switch back to File Explorer.
2. Right-click **Report.txt**, and then click **Delete**.
3. In File Explorer, right-click the **Development** folder, and then click **Properties**.
4. In the **Development Properties** dialog box, click the **Previous Versions** tab.
5. Click the most recent folder version for **Development**, and then click **Open**.
6. Confirm that Report.txt is in the folder, right-click **Report.txt**, and then click **Copy**.
7. Close the File Explorer window that just opened.
8. In the other File Explorer window, right-click the **Development** folder, and then click **Paste**.
9. Close File Explorer.
10. Click **OK** and close all open windows.

Results: After completing this exercise, you will have enabled shadow copies on the file server.

Exercise 3: Enabling and Configuring Work Folders

► **Task 1: Install the Work Folders role service**

1. On LON-SVR1, on the taskbar, click the **Windows PowerShell®** icon.
2. At the command prompt type the following command and press Enter:

Add-WindowsFeature FS-SyncShareService

Note that the name of the feature is case-sensitive.

► **Task 2: Create a Sync Share on the File Server**

1. On LON-SVR1, in the Windows PowerShell command window type the following command and press Enter:

New-SyncShare Corp -path C:\CorpData -User "Adatum\Domain Users"

2. If required, on the taskbar, click the **Server Manager** icon to open Server Manager.
3. Click **File and Storage Services**.
4. Click **Work Folders** and ensure the Corp sync share exists.

► **Task 3: Automate settings for users via Group Policy**

1. On LON-DC1, in Server Manager, click **Tools** and click **Group Policy Management**.
2. In the **Group Policy Management Console**, go to **Forest:Adatum.com\Domains\Adatum.com**.
3. Right-click **Adatum.com** and click **Create a GPO in this domain, and Link it here**.
4. In the **New GPO** dialog box, in **Name**, type **Work Folders**, and then click **OK**.
5. Right-click the **Work Folders** GPO and then click **Edit**.
6. In the Group Policy Management Editor window, go to **User Configuration\Policies\Administrative Templates\Windows Components\Work Folders**.
7. In the details pane, double-click **Specify Work Folders settings**.
8. Click **Enabled** and, in **Work Folders URL**, type **http://lon-svr1.Adatum.com**.
9. Select **Force automatic setup** and click **OK**.
10. Close all open windows.

► **Task 4: Test synchronization**

1. Sign in to LON-CL1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On the Start screen, click **Desktop**.
3. On the taskbar, click the **File Explorer** icon.
4. Navigate to **C:\Labfiles\Mod10** and double-click **WorkFolders.bat**.
This adds a registry entry to allow unsecured connections to the work folders.
5. In the lower-left corner of the screen, click the **Start** button.
6. Sign out of LON-CL1.
7. Sign in to LON-CL1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
8. Click the **Desktop** tile and click **File Explorer**.
9. Double-click the **Work Folders** folder.
10. In the Work Folders folder, right-click an empty space, point to **New**, and then click **Text Document**.
11. Name the new text document **TestFile2**, and then press Enter.
12. Switch to LON-SVR1 and click **File Explorer**.

13. Navigate to **C:\CorpData\Administrator**. Ensure the new text file named TestFile2 exists.
14. Close all open Windows.

Results: After completing this exercise, you will have installed the Work Folders role service, created a sync share, and created a Group Policy Object to deliver the settings to the users automatically. You will have also tested the settings.

Exercise 4: Creating and Configuring a Printer Pool

► Task 1: Install the Print and Document Services server role

1. On LON-SVR1, on the taskbar, click the **Server Manager** icon.
2. In Server Manager, on the menu toolbar, click **Manage**.
3. Click **Add Roles and Features**, click **Next**.
4. Click **Role-based or feature-based Installation**, click **Next**.
5. On the **Select destination server** page, click the server on which you want to install the Print and Document Services, and then click **Next**.
The default server is the local server.
6. On the **Select Server Roles** page, select **Print and Document Services**.
7. In the Add Roles and Features Wizard, click **Add Features**.
8. On the **Select server roles** page, click **Next**.
9. On the **Select Features** page, click **Next**.
10. On the **Print and Document Services** page, review the Notes for the administrator, and then click **Next**.
11. On the **Select Role Services** page, click **Next** until the **Confirm Installation Selections** page displays.
12. Click **Install** to install the required role services.
13. Click **Close**.

► Task 2: Install a printer

1. On LON-SVR1, in the Server Manager, click **Tools**, and then click **Print Management**.
2. Expand **Printer Servers**, expand **LON-SVR1**, right-click **Printers**, and then click **Add Printer**.
The Network Printer Installation Wizard starts.
3. On the **Network Printer Installation Wizard** page, click **Add a TCP/IP or Web Services Printer by IP address or hostname**, and then click **Next**.
4. Change the Type of Device to **TCP/IP Device**.
5. In **Host name or IP address**, type **172.16.0.200**, clear **Auto detect the printer driver to use**, and then click **Next**.
6. Under Device Type, click **Generic Network Card**, and then click **Next**.
7. Click **Install a new driver**, and then click **Next**.
8. Click **Microsoft** as the Manufacturer, under Printers, click **Microsoft XPS Class Driver**, and then click **Next**.

9. Change the Printer Name to **Branch Office Printer**, and then click **Next**.
10. Click **Next** two times to accept the default printer name and share name, and to install the printer.
11. Click **Finish** to close the Network Printer Installation Wizard.
12. In the Print Management console, right-click the **Branch Office Printer**, and then click **Enable Branch Office Direct Printing**.
13. In the Print Management console, right-click the **Branch Office Printer**, and then select **Properties**.
14. Click the **Sharing** tab, select **List in the directory**, and then click **OK**.

► Task 3: Configure printer pooling

1. In the Print Management console, under **LON-SVR1**, right-click **Ports**, and then click **Add Port**.
2. In the **Printer Ports** dialog box, click **Standard TCP/IP Port**, and then click **New Port**.
3. In the Add Standard TCP/IP Printer Port Wizard, click **Next**.
4. In **Printer Name or IP Address**, type **172.16.0.201**, and then click **Next**.
5. In the **Additional port information required** dialog box, click **Next**.
6. Click **Finish** to close the Add Standard TCP/IP Printer Port Wizard.
7. Click **Close** to close the **Printer Ports** dialog box.
8. In the Print Management console, click **Printers**, right-click **Branch Office Printer**, and then click **Properties**.
9. In the **Branch Office Printer Properties** dialog page, click the **Ports** tab, select **Enable printer pooling**, and then click the **172.16.0.201** port to select it as the second port.
10. Click **OK** to close the **Branch Office Printer Properties** dialog box.
11. Close the Print Management Console.

► Task 4: Install a printer on a client computer

1. On LON-CL1, in the lower-left corner of the screen, right-click the **Start** button, and click **Control Panel**.
2. In Control Panel, under **Hardware and Sound**, click **Add a device**.
3. In the **Add a device** dialog box, click **Branch Office Printer on LON-SVR1**, and then click **Next**. The device installs automatically.

Results: After completing this exercise, you will have installed the Print and Document Services server role and installed a printer with printer pooling.

► Prepare for the next module

After you finish the lab, revert the virtual machines back 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 **20410C-LON-SVR1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20410C-LON-CL1** and **20410C-LON-DC1**.

Module 11: Implementing Group Policy

Lab: Implementing Group Policy

Exercise 1: Configuring a Central Store

► **Task 1: View the location of administrative templates in a Group Policy Object (GPO)**

1. Sign in to LON-DC1 as **Administrator** with a password of **Pa\$\$w0rd**.
2. In Server Manager, click **Tools**, and then click **Group Policy Management**.
3. In the Group Policy Management Console (GPMC), expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**, and then expand the **Group Policy Objects** folder.
4. Right-click the **Default Domain Policy**, and then click **Edit**. This opens the Group Policy Management Editor window.
5. In the Group Policy Management Editor window, expand the **Default Domain Policy**, under **User Configuration**, expand **Policies**, and then click **Administrative Templates**.
6. Point to the **Administrative Templates** folder, and note that the location is **Administrative Templates: Policy definitions (.admx files) retrieved from the local computer**.
7. Close the Group Policy Management Editor window.

► **Task 2: Create a Central Store**

1. On the taskbar, click the **File Explorer** icon.
2. In the File Explorer window, expand **Local Disk (C:)**, expand **Windows**, expand **SYSVOL**, expand **sysvol**, expand **Adatum.com**, and then double-click **Policies**.
3. In the details pane, right-click a blank area, click **New**, and then click **Folder**.
4. Name the folder **PolicyDefinitions**.

► **Task 3: Copy administrative templates to the Central Store**

1. In File Explorer, go to **C:\Windows**, and open the **PolicyDefinitions** folder.
2. Select the entire contents of the **PolicyDefinitions** folder.



Note: Hint: To select all the content, click in the details pane, and then press CTRL+A.

3. Right-click the selection, and then click **Copy**.
4. Expand **Local Disk (C:)**, expand **Windows**, expand **SYSVOL**, expand **sysvol**, expand **Adatum.com**, expand **Policies**, and then open the **PolicyDefinitions** folder.
5. Right-click in the empty folder area, and then click **Paste**.

► Task 4: Verify the administrative template location in GPMC

1. In the GPMC, right-click the **Default Domain Policy**, and then click **Edit**.
2. In the Group Policy Management Editor window, expand **Policies**, point to the **Administrative Templates** folder and read the local information text, which reads: "Administrative Templates: Policy definitions (ADMX files) retrieved from the Central Store."
3. Close the **Group Policy Management Editor** window.

Results: After completing this exercise, you should have configured a Central Store.

Exercise 2: Creating GPOs

► Task 1: Create a Windows Internet Explorer Restriction default starter GPO

1. In the GPMC, right-click the **Starter GPOs** folder, and then click **New**.
2. In the **New Starter GPO** dialog box, in the **Name** field, type **Internet Explorer Restrictions**, in the **Comment** field, type **This GPO disables the General page in Internet Options**, and then click **OK**.

► Task 2: Configure the Internet Explorer Restriction starter GPO

1. In the GPMC, under the Starter GPOs folder, right-click the **Internet Explorer Restrictions** GPO, and then click **Edit**.
2. In the Group Policy Management Editor window, expand **User Configuration, Administrative Templates**, and then click **All Settings**.
3. Right-click **All Settings**, and then click **Filter Options**.
4. In the **Filter Options** dialog box, select the **Enable Keyword Filters** check box.
5. In the **Filter for word(s):** field, type **General page**.
6. Beside **Within**, deselect the **Help Text** and the **Comment** checkboxes.
7. Beside the **Filter for word(s)** field, click the drop-down list box, click **Exact**, and then click **OK**.
8. Double-click the **Disable the General page** setting, click **Enabled**, and then click **OK**.
9. Close the **Group Policy Starter GPO Editor** window.

► Task 3: Create an Internet Explorer Restrictions GPO from the Internet Explorer Restrictions starter GPO

1. In the GPMC, right-click the **Adatum.com** domain, and then click **Create a GPO in this domain, and Link it here**.
2. In the **New GPO** dialog box, in the **Name** field, type **IE Restrictions**.
3. Under Source Starter GPO, click the drop-down box, select **Internet Explorer Restrictions**, and then click **OK**.

► Task 4: Test the GPO for Domain Users

1. Sign in to **LON-CL1** as **Adatum\Brad** with a password of **Pa\$\$w0rd**.
2. Point the mouse at the lower-right edge of the screen and click the **Search** charm when it appears.
3. In the **Everywhere** search box, type **Control Panel**.
4. In the search results, click **Control Panel**.

5. In Control Panel, click **Network and Internet**.
6. In the **Network and Internet** dialog box, click **Change your homepage**.
7. Read the message box that appears informing you that this feature has been disabled, and then click **OK**.
8. In the Control Panel, click **Internet Options**. Notice that in the **Internet Properties** dialog box the **General** tab does not display.
9. Close all open windows and sign out from **LON-CL1**.

► **Task 5: Use security filtering to exempt the IT Department from the Internet Explorer Restrictions policy**

1. Switch to **LON-DC1**.
2. In the GPMC, expand the **Group Policy Objects** folder, and then in the left pane, click the **IE Restrictions** policy.
3. In the details pane, click the **Delegation** tab.
4. On the **Delegation** tab, click the **Advanced** button.
5. In the **IE Restrictions Security Settings** dialog box, click **Add**.
6. In the Select Users, Computers, Service Accounts, or Groups window, in the **Enter the object names to select (examples)** box, type **IT**, and then click **OK**.
7. In the **IE Restrictions Security Settings** dialog box, click the **IT (Adatum\IT)** group, next to the Apply group policy permission, select the **Deny** check box, and then click **OK**.
8. Click **Yes** to acknowledge the **Windows Security** dialog box.

► **Task 6: Test the GPO app for IT department users**

1. Sign in to **LON-CL1** as **Brad**, with the password **Pa\$\$w0rd**.
2. Switch to **LON-CL1**.
3. Point the mouse at the lower-right edge of the screen and click the **Search** charm when it appears.
4. In the **Everywhere** search box, type **Control Panel**.
5. In the search results window, click **Control Panel**.
6. In Control Panel, click **Network and Internet**.
7. In the **Network and Internet** dialog box, click **Change your homepage**. The **Internet Properties** dialog box opens to the **General** tab, and all settings are available.
8. Close all open windows, and sign out from **LON-CL1**.

► **Task 7: Test Application of the GPO for other domain users**

1. Sign in to **LON-CL1** as **Boris** with a password of **Pa\$\$w0rd**.
2. Point the mouse at the lower-right edge of the screen and click the **Search** charm when it appears.
3. In the **Everywhere** search box, type **Control Panel**.
4. In the search results window, click **Control Panel**.
5. In Control Panel, click **Network and Internet**.
6. In the **Network and Internet** dialog box, click **Change your homepage**. A message box appears informing you that this feature has been disabled.

7. Click **OK** to acknowledge the message.
8. Click **Internet Options**. In the **Internet Properties** dialog box, notice that the **General** tab does not display.
9. Close all open windows, and sign out from **LON-CL1**.

Results: After completing this lab, you should have created a GPO.

► Prepare for the next module

After you finish the lab, revert the virtual machines back 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 **20410C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for 20410C-LON-CL1.

Module 12: Securing Windows Servers Using Group Policy Objects

Lab A: Increasing Security for Server Resources

Exercise 1: Using Group Policy to Secure Member Servers

► **Task 1: Create a Member Servers organizational unit (OU) and move servers into it**

1. On LON-DC1, in Server Manager, click **Tools**, and then click **Active Directory Users and Computers**.
2. In Active Directory Users and Computers, in the navigation pane, right-click **Adatum.com**, click **New**, and then click **Organizational Unit**.
3. In the New Object - Organizational Unit window, in the **Name** box, type **Member Servers OU**, and then click **OK**.
4. In Active Directory Users and Computers, in the navigation pane, click **Computers** container.
5. Press and hold the CTRL key. In the details pane, click both **LON-SVR1** and **LON-SVR2**, right-click the selection, and then click **Move**.
6. In the Move window, click **Member Servers OU**, and then click **OK**.

► **Task 2: Create a Server Administrators group**

1. On LON-DC1, in Active Directory Users and Computers, in the navigation pane, right-click the **Member Servers OU**, click **New**, and then click **Group**.
2. In the New Object – Group window, in **Group Name**, type **Server Administrators**, and then click **OK**.

► **Task 3: Create a Member Server Security Settings Group Policy Object (GPO) and link it to the Member Servers OU**

1. On LON-DC1, in the Server Manager window, click **Tools**, and then click **Group Policy Management**.
2. In the Group Policy Management Console, expand **Forests: Adatum.com**, expand **Domains**, expand **Adatum.com**, right-click **Group Policy Objects**, and then click **New**.
3. In the New GPO window, in **Name**, type **Member Server Security Settings**, and then click **OK**.
4. In the Group Policy Management Console, right-click **Member Servers OU**, and then click **Link an Existing GPO**.
5. In the Select GPO window, in the Group Policy Objects window, click **Member Server Security Settings**, and then click **OK**.

► **Task 4: Configure group membership for local administrators to include Server Administrators and Domain Admins**

1. On LON-DC1, open the Group Policy Management Console.
2. If necessary, click the Group Policy Objects container, right-click **Default Domain Policy**, and then click **Edit**.
3. In the Group Policy Management Editor window, go to **Computer Configuration\Policies\Windows Settings\Security Settings\Restricted Groups**.

4. Right-click **Restricted Groups**, and then click **Add Group**.
5. In the **Add Group** dialog box, in **Group name**, type **Administrators**, and then click **OK**.
6. In the **Administrators Properties** dialog box, next to **Members of this group**, click **Add**.
7. In the **Add Member** dialog box type **Adatum\Server Administrators**, and then click **OK**.
8. Next to **Members of this group**, click **Add**.
9. In the **Add Member** dialog box type **Adatum\Domain Admins**, and then click **OK** twice.
10. Close the Group Policy Management Editor window.

► **Task 5: Verify that Computer Administrators has been added to the local Administrators group**

1. Switch to LON-SVR1.
2. On the taskbar, click the **Windows PowerShell®** icon.
3. At the Windows PowerShell prompt, type the following command, and then press Enter:

```
Gpupdate /force
```

4. In the Server Manager window, click **Tools**, and then click **Computer Management**.
5. In the Computer Management console, expand **Local Users and Groups**, click **Groups**, and then in the right-hand pane, double-click **Administrators**.
6. Confirm that the **Administrators** group contains both **ADATUM\Domain Admins** and **ADATUM\Server Administrators** as members. Click **Cancel**.
7. Close the Computer Management console.

► **Task 6: Modify the Member Server Security Settings GPO to remove Users from Allow Log On Locally**

1. On LON-DC1, in the Group Policy Management Console, click **Group Policy Objects**.
2. In the right-hand pane, right-click **Member Server Security Settings**, and then click **Edit**.
3. In the Group Policy Management Editor window, go to **Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\User Rights Assignment**.
4. In the right-hand pane, right-click **Allow log on locally**, and then click **Properties**.
5. In the **Allow log on locally Properties** dialog box, select the **Define these policy settings** check box, and then click **Add User or Group**.
6. In the Add User or Group window, type **Domain Admins**, and then click **OK**.
7. Click **Add User or Group**.
8. In the Add User or Group window, type **Administrators**, and then click **OK** twice.

► **Task 7: Modify the Member Server Security Settings GPO to enable User Account Control: Admin Approval Mode for the Built-in Administrator account**

1. On LON-DC1, in the Group Policy Management Editor window, go to **Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\Security Options**.
2. In the right-hand pane, right-click **User Account Control: Admin Approval Mode for the Built-in Administrator account**, and then click **Properties**.

3. In the **User Account Control: Admin Approval Mode for the Built-in Administrator account Properties** dialog box, select the **Define this policy settings** check box, ensure that **Enabled** is selected, and then click **OK**.
4. Close the Group Policy Management Editor window.

► **Task 8: Verify that a non-administrative user cannot sign in to a member server**

1. Switch to LON-SVR1.
2. On the taskbar, click the **Windows PowerShell** icon.
3. At the Windows PowerShell prompt, type the following command, and then press Enter:

```
Gpupdate /force
```

4. Sign out of LON-SVR1.
5. Try to sign in to LON-SVR1 as **Adatum\Adam** with the password **Pa\$\$w0rd**.
Verify that you cannot sign in to LON-SVR1, and that a logon error message is displayed.
6. To prepare for the next exercise, sign out of LON-SVR1, and sign back in to LON-SVR1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.

Results: After completing this exercise, you should have used Group Policy to secure Member servers.

Exercise 2: Auditing File System Access

► **Task 1: Modify the Member Server Security Settings GPO to enable object access auditing**

1. Switch to LON-DC1.
2. In the Group Policy Management Console, go to **Forest: Adatum.com\Domains\Adatum.com**.
3. Click **Group Policy Objects**.
4. In the right-hand pane, right-click **Member Server Security Settings**, and then click **Edit**.
5. In the Group Policy Management Editor window, go to **Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies**.
6. Click **Audit Policy**.
7. In the right-hand pane, right-click **Audit object access**, and then click **Properties**.
8. In the **Audit object access Properties** dialog box, select the **Define these policy settings** check box, select both the **Success** and **Failure** check boxes, and then click **OK**.
9. Sign out from LON-DC1.

► **Task 2: Create and share a folder**

1. Switch to LON-SVR1.
2. On LON-SVR1, on the taskbar, click the **File Explorer** icon.
3. In File Explorer, in the navigation pane, double-click **Local Disk (C)**, and then click **Home**.
4. Click **New folder**, type **Marketing**, and then press Enter.

5. In the Computer window, right-click the **Marketing** folder, click **Share with**, and then click **Specific people**.
6. In the File Sharing window, type **Adam**, and then click **Add**.
7. Change the Permission Level to **Read/Write**, click **Share**, and then click **Done**.

► **Task 3: Enable auditing on the Marketing folder for Domain Users**

1. On LON-SVR1, in the Local Disk (C:) window, right-click the **Marketing** folder, and then click **Properties**.
2. In the Marketing Properties window, click the **Security** tab, and then click **Advanced**.
3. In the Advanced Security Settings for Marketing window, click the **Auditing** tab, click **Continue**, and then click **Add**.
4. In the Auditing Entry for Marketing window, click **Select a principal**.
5. In the Select User, Computer, Service Account or Group window, in **Enter the object name to select**, type **Domain Users**, and then click **OK**.
6. In the Auditing Entry for Marketing window, from the **Type** drop-down menu, select **All**.
7. In the Auditing Entry for Marketing window, under the **Permission** list, select the **Write** check box, and then click **OK** three times.
8. On the taskbar, click the **Windows PowerShell** icon.
9. At the Windows PowerShell prompt, type the following command, and then press Enter:

```
gpupdate /force
```

10. Close the Windows PowerShell window.

► **Task 4: Create a new file in the file share from LON-CL1**

1. Sign in to LON-CL1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Point to the lower-right corner of the screen, and then click the **Search** charm when it appears.
3. In the **Search** box type **cmd**, and then press Enter.
4. Open the Command Prompt window, and at the command prompt, type the following command, and then press Enter:

```
gpupdate /force
```

5. Close the Command Prompt window.
6. Sign out from LON-CL1, and then sign in again as **Adatum\Adam** with the password **Pa\$\$w0rd**.
7. Point to the lower-right corner of the screen, and then click the **Search** charm when it appears.
8. In the **Search** box type **\\LON-SVR1\Marketing**, and then press Enter.
9. In the Marketing window, click **Home**, click **New item**, click **Text Document**, in **File name**, type **Employees**, and then press Enter.
10. Sign out from LON-CL1.

► **Task 5: View the results in the security log on the domain controller**

1. Switch to LON-SVR1.
2. In the Server Manager window, click **Tools**, and then click **Event Viewer**.
3. In the Event Viewer window, expand **Windows Logs**, and then click **Security**.

4. Verify that the following event and information is displayed:
 - Source: **Microsoft Windows Security Auditing**
 - Event ID: **4663**
 - Task category: **File System**
 - An attempt was made to access an object

Results: After completing this exercise, you should have enabled file system access auditing.

Exercise 3: Auditing Domain Logons

► Task 1: Modify the Default Domain Policy GPO

1. Sign in to LON-DC1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On LON-DC1, on the taskbar, click the **Server Manager** icon.
3. In the Server Manager window, click **Tools**, and then click **Group Policy Management**.
4. On LON-DC1, in the Group Policy Management Console, go to **Forest: Adatum.com \Domains\Adatum.com**.
5. Click **Group Policy Objects**.
6. In the right-hand pane, right-click **Default Domain Policy**, and then click **Edit**.
7. In the Group Policy Management Editor window, go to **Computer Configuration\Policies \Windows Settings\Security Settings\Local Policies**.
8. Click **Audit Policy**.
9. In the right-hand pane, right-click **Audit account logon events**, and then click **Properties**.
10. In the **Audit account logon events Properties** dialog box, select the **Define these policy settings** check box, select both the **Success** and **Failure** check boxes, and then click **OK**.
11. Point to the lower-right corner of the screen, and then click the **Search** charm when it appears.
12. In the **Search** box type **cmd**, and then press Enter.
13. At the command prompt, type the following command, and then press Enter:

```
gpupdate /force
```

► Task 2: Run GPOUpdate

1. Sign in to LON-CL1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Point to the lower-right corner of the screen, and then click the **Search** charm when it appears.
3. In the **Search** box type **cmd**, and then press Enter.
4. At the command prompt, type the following command, and then press Enter:

```
gpupdate /force
```

5. Close the Command Prompt window, and sign out from LON-CL1.

► **Task 3: Sign in to LON-CL1 with an incorrect password**

- Sign in to LON-CL1 as **Adatum\Adam** with the password **password**.



Note: This password is intentionally incorrect to generate a security log entry that shows that an unsuccessful sign in attempt has been made.

► **Task 4: Review event logs on LON-DC1**

1. On LON-DC1, in Server Manager, click **Tools**, and then click **Event Viewer**.
2. In the Event Viewer window, expand **Windows Logs**, and then click **Security**.
3. Review the event logs for following message: "**Event ID 4771 Kerberos pre-authentication failed. Account Information: Security ID: ADATUM\Adam.**"

► **Task 5: Sign in to LON-CL1 with the correct password**

1. Sign in to LON-CL1 as **Adatum\Adam** with the password **Pa\$\$w0rd**.



Note: This password is correct, and you should be able to sign in successfully as **Adam**.

2. Sign out of LON-CL1.

► **Task 6: Review event logs on LON-DC1**

1. Switch to LON-DC1.
2. In the Server Manager window, click **Tools**, and then click **Event Viewer**.
3. In the Event Viewer window, expand **Windows Logs**, and then click **Security**.
4. Review the event logs for the following message: "**Event ID 4624 An account was successfully logged on. New Logon: Security ID: ADATUM\Adam.**"

Results: After completing this exercise, you should have enabled domain logon auditing.

► **Prepare for the next lab**

- To prepare for the next lab, leave the virtual machines running.

Lab B: Configuring AppLocker and Windows Firewall

Exercise 1: Configuring AppLocker Policies

► Task 1: Create an OU for client computers

1. Switch to LON-DC1.
2. In Server Manager, click **Tools**, and then click **Active Directory Users and Computers**.
3. In Active Directory Users and Computers, in the navigation pane, right-click **Adatum.com**, click **New**, and then click **Organizational Unit**.
4. In the New Object - Organizational Unit window, type **Client Computers OU**, and then click **OK**.

► Task 2: Move LON-CL1 to the Client Computers OU

1. On LON-DC1, in Active Directory Users and Computers, in the navigation pane, click **Computers** container.
2. In the details pane, right-click **LON-CL1**, and then click **Move**.
3. In the Move window, click **Client Computers OU**, and then click **OK**.

► Task 3: Create a Software Control GPO and link it to the Client Computers OU

1. On LON-DC1, in Server Manager, click **Tools**, and then click **Group Policy Management**.
2. In the Group Policy Management Console, go to **Forests: Adatum.com/Domains/Adatum.com**.
3. Right-click **Group Policy Objects**, and then click **New**.
4. In New GPO window, in the **Name** text box type **Software Control GPO**, and then click **OK**.
5. In the right-hand pane, right-click **Software Control GPO**, and then click **Edit**.
6. In the Group Policy Management Editor window, go to **Computer Configuration/Policies /Windows Settings/Security Settings/Application Control Policies/AppLocker**.
7. Under AppLocker, right-click **Executable Rules**, and then click **Create Default Rules**.
8. Repeat the previous step for Windows Installer Rules, Script Rules, and Packaged app Rules.
9. In the navigation pane, click **AppLocker**, and then in the right-hand pane, click **Configure rule enforcement**.
10. In the **AppLocker Properties** dialog box, under Executable rules, select the **Configured** check box, and then from the drop-down menu, select **Audit only**.
11. Repeat the previous step for Windows Installer Rules, Script Rules, and Packaged app Rules, and then click **OK**.
12. In the Group Policy Management Editor window, go to **Computer Configuration/Policies /Windows Settings/Security Settings**.
13. Click **System Services**, and then double-click **Application Identity**.
14. In the **Application Identity Properties** dialog box, click **Define this policy setting**.

15. Under Select service startup mode, click **Automatic**, and then click **OK**.
16. Close the Group Policy Management Editor window.
17. In the Group Policy Management Console, right-click **Client Computers OU**, and then click **Link an Existing GPO**.
18. In the Select GPO window, in the **Group Policy Objects** list, click **Software Control GPO**, and then click **OK**.

► Task 4: Run GPUpdate

1. Switch to LON-CL1.
2. Point to the lower-right corner of the screen, and then click the **Search** charm when it appears.
3. In the **Search** box type **cmd**, and then press Enter.
4. In the Command Prompt window, type following command, and then press Enter:

```
gpupdate /force
```

5. Close the Command Prompt window.
6. Point to the lower-right corner of the screen, and then click the **Settings** charm when it appears.
7. Click **Power**, and then click **Restart**.

► Task 5: Run app1.bat in the C:\CustomApp folder

1. Sign in to LON-CL1 as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Point to the lower-right corner of the screen, and then click the **Search** charm when it appears.
3. In the **Search** box type **cmd**, and then press Enter.
4. At the command prompt, type following command, and then press Enter:

```
gpresult /R
```

Review the result of the command and ensure that **Software Control GPO** is displayed under Computer Settings, Applied Group Policy Objects.

5. If **Software Control GPO** is not displayed, restart LON-CL1 and repeat steps 1 through 4.
6. Point to the lower-right corner of the screen, and then click the **Search** charm when it appears.
7. In the **Search** box type **cmd**, and then press Enter.
8. At the command prompt, type the following command, and then press Enter:

```
C:\CustomApp\app1.bat
```

► Task 6: View AppLocker events in an event log

1. On LON-CL1, point to the lower-right corner of the screen, and then click the **Search** charm when it appears.
2. In the **Search** box type **eventvwr.msc**, and then press Enter.
3. In the Event Viewer window, expand **Application and Services Logs**, expand **Microsoft**, expand **Windows**, and then expand **AppLocker**.

- Click **MSI and Scripts** and review event log 8005 that contains the following text:
%OSDRIVE%\CUSTOMAPP\APP1.BAT was allowed to run.



Note: If no events are displayed, ensure that the Application Identity service has started, and try again.

► **Task 7: Create a rule that allows software to run from a specific location**

- On LON-DC1, in Server Manager, click **Tools**, and then click **Group Policy Management**.
- In the Group Policy Management Console, expand the **Group Policy Objects** node, right-click **Software Control GPO**, and click **Edit**.
- In the Group Policy Management Editor window, go to **Computer Configuration/Policies /Windows Settings/Security Settings/Application Control Policies/AppLocker**.
- Right-click **Script rules** and click **Create New Rule**.
- On the **Before You Begin** page, click **Next**.
- On the **Permissions** page, click **Allow**, and then click **Next**.
- On the **Conditions** page, click **Path**, and then click **Next**.
- On the **Path** page, in **Path**, type the path **%OSDRIVE%\CustomApp\app1.bat**, and then click **Next**.
- On the **Exception** page, click **Next**.
- On the **Name and Description** page, in **Name**, type **Custom App Rule**, and then click **Create**.

► **Task 8: Modify the Software Control GPO to enforce rules**

- In the Software control GPO window, in the navigation pane, click **AppLocker**, and then in the right-hand pane, click **Configure rule enforcement**.
- In **AppLocker Properties** dialog box, under Executable rules, select the **Configured** check box, and then from drop-down menu, click **Enforce rules**.
- Repeat the previous step for Windows Installer Rules, Script Rules, and Packaged app Rules, and then click **OK**.
- Close the Group Policy Management Editor window.

► **Task 9: Verify that an application can still be run**

- Switch to LON-CL1.
- Point to the lower-right corner of the screen, and then click the **Search** charm when it appears.
- In the **Search** box type **cmd**, and then press Enter.
- In the Command Prompt window, type the following command, and then press Enter:

```
gpupdate /force
```

- Close the Command Prompt window.
- Point to the lower-right corner of the screen, and then click the **Settings** charm when it appears.
- Click **Power**, and then click **Restart**.
- Sign in to LON-CL1 as **Adatum\Tony** with the password **Pa\$\$w0rd**.

9. Point to the lower-right corner of the screen, and then click the **Search** charm when it appears.
10. In the **Search** box type **cmd**, and then press Enter.
11. In the Command Prompt window, type following command, and then press Enter:

```
C:\customapp\app1.bat
```

► **Task 10: Verify that an app cannot be run**

1. On LON-CL1, on the taskbar, click the **File Explorer** icon.
2. In File Explorer, in the navigation pane, click **Computer**.
3. In the Computer window, double-click **Local Disk (C:)**, double-click the **CustomApp** folder, right-click **app1.bat**, and then click **Copy**.
4. In the CustomApp window, on the navigation pane, right-click the **Documents** folder, and then click **Paste**.
5. In the Command Prompt window, type **C:\Users\Tony\Documents\app1.bat**, and then press Enter.
6. Verify that apps cannot be run from the Documents folder, and that the following message is displayed: "This program is blocked by Group Policy. For more information, contact your system administrator."
7. Close all open windows, and sign out from LON-CL1.

Results: After completing this exercise, you should have configured AppLocker policies for all users whose computer accounts are located in the Client Computers OU. The policies you configured should allow these users to run apps that are located in the folders C:\Windows and C:\Program Files, and run the custom-developed app app1.bat in the C:\CustomApp folder.

Exercise 2: Configuring Windows Firewall

► **Task 1: Create a group named Application Servers**

1. Switch to LON-DC1.
2. In the Server Manager window, click **Tools**, and then click **Active Directory Users and Computers**.
3. In Active Directory Users and Computers, in the navigation pane, right-click the **Member Servers OU**, click **New**, and then click **Group**.
4. In the New Object – Group window, in **Group Name**, type **Application Servers**, and then click **OK**.

► **Task 2: Add LON-SVR1 as a group member**

1. In Active Directory Users and Computers, in the navigation pane, click the **Member Servers OU**, in the details pane right-click **Application Servers group**, and then click **Properties**.
2. In the **Application Server Properties** dialog box, click the **Members** tab, and then click **Add**.
3. In **Select Users, Computers, Service Accounts or Groups**, click **Object Types**, click **Computers**, and then click **OK**.
4. In the **Enter the object names to select** box type **LON-SVR1**, and then click **OK**.
5. In the **Application Server Properties** dialog box, click **OK**.

► **Task 3: Create a new Application Servers GPO**

1. On LON-DC1, in Server Manager, click **Tools**, and then click **Group Policy Management**.
2. In the Group Policy Management Console, expand **Forests: Adatum.com**, expand **Domains**, expand **Adatum.com**, right-click **Group Policy Objects**, and then click **New**.
3. In the New GPO window, in **Name**, type **Application Servers GPO**, and then click **OK**.
4. In the Group Policy Management Console, right-click **Application Servers GPO**, and then click **Edit**.
5. In the Group Policy Management Editor window, go to **Computer Configuration/Policies /Windows Settings/Security Settings/Windows Firewall with Advanced Security**.
6. Click **Windows Firewall with Advanced Security - LDAP://CN={GUID}**.
7. In the Group Policy Management Editor window, click **Inbound Rules**.
8. Right-click **Inbound Rules**, and then click **New Rule**.
9. In the New Inbound Rule Wizard, on the **Rule Type** page, click **Custom**, and then click **Next**.
10. On the **Program** page, click **Next**.
11. On the **Protocol and Ports** page, in the **Protocol type** list, click **TCP**.
12. In the **Local port** list, click **Specific Ports**, in the text box type **8080**, and then click **Next**.
13. On the **Scope** page, click **Next**.
14. On the **Action** page, click **Allow the connection**, and then click **Next**.
15. On the **Profile** page, clear both the **Private** and **Public** check boxes, and then click **Next**.
16. On the **Name** page, in the **Name** box type **Application Server Department Firewall Rule**, and then click **Finish**.
17. Close the Group Policy Management Editor window.

► **Task 4: Link the Application Servers GPO to the Member Servers OU**

1. On LON-DC1, in the Group Policy Management Console, right-click **Member Servers OU**, and then click **Link an Existing GPO**.
2. In the Select GPO window, in the **Group Policy objects** list, click **Application Servers GPO**, and then click **OK**.

► **Task 5: Use security filtering to limit the Application Server GPO to members of Application Server group**

1. On LON-DC1, in the Group Policy Management Console, click **Member Servers OU**.
2. Expand the **Member Servers OU**, and then click the **Application Servers GPO** link.
3. In the Group Policy Management Console message box, click **OK**.
4. In the right-hand pane, under **Security Filtering**, click **Authenticated Users**, and then click **Remove**.
5. In the **Confirmation** dialog box, click **OK**.
6. In the details pane, under Security Filtering, click **Add**.
7. In the **Select User, Computer, or Group** dialog box, type **Application Servers**, and then click **OK**.

► Task 6: Run GPUUpdate on LON-SVR1

1. Switch to LON-SVR1 and sign in as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Point to the lower-right corner of the screen, and then click the **Search** charm when it appears.
3. In the **Search** box type **cmd**, and then press Enter.
4. In the Command Prompt window, type following command, and then press Enter:

```
gpupdate /force
```

5. Close the Command Prompt window.
6. Restart LON-SVR1, and then sign back in as **Adatum\Administrator** with the password of **Pa\$\$w0rd**.

► Task 7: View the firewall rules on LON-SVR1

1. Switch to LON-SVR1.
2. In Server Manager, click **Tools**, and then click **Windows Firewall with Advanced Security**.
3. In the Windows Firewall with Advanced Security window, click **Inbound rules**.
4. In the right-hand pane, verify that the **Application Server Department Firewall Rule** that you created earlier using Group Policy is configured.
5. Verify that you cannot edit the **Application Server Department Firewall Rule**, because it is configured through Group Policy.

Results: After completing this exercise, you should have used Group Policy to configure Windows Firewall with Advanced Security to create rules for app servers.

► Prepare for the next module

When you finish the lab, revert the virtual machines to their initial state by performing the following steps:

1. On the host computer, start Hyper-V® Manager.
2. In the **Virtual Machines** list, right-click **20410C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for 20410C-LON-SVR1 and 20410C-LON-CL1.

Module 13: Implementing Server Virtualization with Hyper-V

Lab: Implementing Server Virtualization with Hyper-V

Exercise 1: Installing the Hyper-V® Role onto a Server

► Task 1: Install the Hyper-V role onto a server

1. On LON-HOST1, in Server Manager, click **Local Server**.
2. In the Properties pane, click the **IPv4 address assigned by DHCP, IPv6 enabled** link.
3. In the **Network Connections** dialog box, right-click the network object, and then click **Properties**.
4. In the **Properties** dialog box, click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
5. In the **Properties** dialog box, on the **General** tab, click **Use the following IP address**, and then configure the following:
 - IP Address: **172.16.0.31**
 - Subnet mask: **255.255.0.0**
 - Default gateway: **172.16.0.1**
6. On the **General** tab, click **Use the following DNS server addresses**, and then configure the following:
 - Preferred DNS server: **172.16.0.10**
7. Click **OK** to close the **Properties** dialog box.
8. In the **Properties** dialog box of the network object click **Close**.
9. Close the **Network Connections** dialog box.
10. In the Server Manager console, from the **Manage** menu, click **Add Roles and Features**.
11. In the Add Roles and Features Wizard, on the **Before you begin** page, click **Next**.
12. On the **Select installation type** page, click **Role-based or feature-based installation**, and then click **Next**.
13. On the **Select destination server** page, ensure that LON-HOST1 is selected, and then click **Next**.
14. On the **Select server roles** page, select **Hyper-V**.
15. In the Add Roles and Features Wizard, click **Add Features**.
16. On the **Select server roles** page, click **Next**.
17. On the **Select features** page, click **Next**.
18. On the **Hyper-V** page, click **Next**.
19. On the **Virtual Switches** page, verify that no selections have been made, and then click **Next**.
20. On the **Virtual Machine Migration** page, click **Next**.
21. On the **Default Stores** page, review the location of the **Default Stores**, and then click **Next**.
22. On the **Confirm installation selections** page, click **Restart the destination server automatically if required**.

23. In the Add Roles and Features Wizard, review the message regarding automatic restarts, and then click **Yes**.

24. On the **Confirm Installation Selections** page, click **Install**.

After a few minutes, the server restarts automatically. Ensure that you restart the machine from the boot menu as **20410C-LON-HOST1**. The computer will restart several times.

► **Task 2: Complete Hyper-V role installation and verify settings**

1. Sign in to **LON-HOST1** using the account **Administrator** with the password **Pa\$\$word**.
2. When the installation of the **Hyper-V** tools completes, click **Close** to close the Add Roles and Features Wizard.
3. In the Server Manager console, click the **Tools** menu, and then click **Hyper-V Manager**.
4. In the Hyper-V Manager console, click **LON-HOST1**.
5. In the Hyper-V Manager console, in the Actions pane, with **LON-HOST1** selected, click **Hyper-V Settings**.
6. In the **Hyper-V Settings for LON-HOST1** dialog box, click the **Keyboard** item. Verify that the Keyboard is set to the **Use on the virtual machine** option.
7. In the **Hyper-V Settings for LON-HOST1** dialog box, click the **Virtual Hard Disks** item.
8. Verify that the location of the default folder to store Virtual Hard Disk files is **C:\Users\Public\Documents\Hyper-V\Virtual Hard Disks**, and then click **OK**.

Results: After completing this exercise, you should have installed the Hyper-V role onto a physical server.

Exercise 2: Configuring Virtual Networking

► **Task 1: Configure the external network**

1. In the Hyper-V Manager console, click **LON-HOST1**.
2. From the **Actions** menu, click **Virtual Switch Manager**.
3. In the **Virtual Switch Manager for LON-HOST1** dialog box, click **New virtual network switch**. Ensure that **External** is selected, and then click **Create Virtual Switch**.
4. In the Virtual Switch Properties area, enter the following information, and then click **OK**:
 - Name: **Switch for External Adapter**
 - External Network: Mapped to the host computer's physical network adapter. (This varies depending on the host computer.)
5. In the **Apply Networking Changes** dialog box, review the warning, and then click **Yes**.

► **Task 2: Create a private network**

1. In Server Manager, on the **Tools** menu, open **Hyper-V Manager**, and then click **LON-HOST1**.
2. From the **Actions** menu, click **Virtual Switch Manager**.
3. Under Virtual Switches, click **New virtual network switch**.
4. Under Create virtual switch, click **Private**, and then click **Create Virtual Switch**.

5. In the Virtual Switch Manager dialog box, in the Virtual Switch Properties section, configure the following settings, and then click **OK**:
 - Name: **Private Network**
 - Connection type: **Private network**

► **Task 3: Create an internal network**

1. In Server Manager, on the **Tools** menu, open **Hyper-V Manager**, and then click **LON-HOST1**.
2. From the **Actions** menu, click **Virtual Switch Manager**.
3. Under Virtual Switches, click **New virtual network switch**.
4. Under Create virtual switch, click **Internal** and then click **Create Virtual Switch**.
5. In the **Virtual Switch Manager** dialog box, in the **Virtual Switch Properties** section, configure the following settings, and then click **OK**:
 - Name: **Internal Network**
 - Connection type: **Internal network**

► **Task 4: Configure the Media Access Control (MAC) address range**

1. In Server Manager, on the **Tools** menu, open **Hyper-V Manager**, and then click **LON-HOST1**.
2. On the **Actions** menu, click **Virtual Switch Manager**.
3. Under Global Network Settings, click **MAC Address Range**.
4. On MAC Address Range settings, configure the following values, and then click **OK**:
 - Minimum: **00-15-5D-0F-AB-A0**
 - Maximum: **00-15-5D-0F-AB-EF**
5. Close the Hyper-V Manager console.

Results: After completing this exercise, you should have configured virtual switch options on a physically-deployed Windows Server 2012 server running the Hyper-V role.

Exercise 3: Creating and Configuring a Virtual Machine

► **Task 1: Create differencing virtual hard disks**

1. On the taskbar, click the **File Explorer** icon.
2. Expand **This PC**, expand drive **E**, expand **Program Files**, expand **Microsoft Learning**, and then expand **Base**.



Note: The drive letter may depend upon the number of drives on the physical host machine.

3. In the Base folder, verify that the **Base14A-WS12R2.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 rename the folders to the following names:
 - **LON-GUEST1**
 - **LON-GUEST2**
5. Close File Explorer.
6. In the Server Manager console, click **Tools**, and then click **Hyper-V Manager**.
7. In the Hyper-V Manager console, in the Actions pane, click **New**, and then click **Hard Disk**.
8. In the New Virtual Hard Disk Wizard, on the **Before You Begin** page, click **Next**.
9. On the **Choose Disk Format** page, click **VHD**, and then click **Next**.
10. On the **Choose Disk Type** page, click **Differencing**, and then click **Next**.
11. On the **Specify Name and Location** page, specify the following details, and then click **Next**:
 - Name: **LON-GUEST1.vhd**
 - Location: **E:\Program Files\Microsoft Learning\Base\LON-GUEST1**
12. On the **Configure Disk** page, type the location: **E:\Program Files\Microsoft Learning\Base****Base14A-WS12R2.vhd**, and then click **Finish**.
13. On the desktop, on the taskbar, click the **Windows PowerShell**® icon.
14. At the Windows PowerShell prompt, type the following command to create a new differencing virtual hard 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\ Base14A-WS12R2.vhd"
```
15. Close Windows PowerShell.
16. In the Hyper-V Manager console, in the Actions pane, click **Inspect Disk**.
17. In the **Open** dialog box, browse to **E:\Program Files\Microsoft Learning\Base\LON-GUEST2**, click **LON-GUEST2.vhd**, and then click **Open**.
18. 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****Base14A-WS12R2.vhd** as a parent, and then click **Close**.

► Task 2: Create virtual machines

1. In Server Manager, on the **Tools** menu, open **Hyper-V Manager**, and then click **LON-HOST1**.
2. In the Hyper-V Manager console, in the Actions pane, click **New**, and then click **Virtual Machine**.
3. In the New Virtual Machine Wizard, on the **Before You Begin** page, click **Next**.
4. On the **Specify Name and Location** page, click **Store the virtual machine in a different location**, enter the following values, and then click **Next**:
 - Name: **LON-GUEST1**
 - Location: **E:\Program Files\Microsoft Learning\Base\LON-GUEST1**
5. On the **Specify Generation** page, select **Generation 1** and then click **Next**.
6. On the **Assign Memory** page, enter a value of **1024 MB**, select the **Use Dynamic Memory for this virtual machine** option, and then click **Next**.

7. On the **Configure Networking** page, for the connection, click **Private Network**, and then click **Next**.
8. On the **Connect Virtual Hard Disk** page, click **Use an existing virtual hard disk**. Click **Browse**, browse to **E:\Program Files\Microsoft Learning\Base\LON-GUEST1\lon-guest1.vhd**, click **Open**, and then click **Finish**.
9. On the desktop, on the taskbar, click the **Windows PowerShell** icon.
10. At the Windows PowerShell prompt, type the following command to create a new virtual machine named **LON-GUEST2**, and then press Enter:

```
New-VM -Name LON-GUEST2 -MemoryStartupBytes 1024MB -VHDPATH "E:\Program
Files\Microsoft Learning\Base\LON-GUEST2\LON-GUEST2.vhd" -SwitchName "Private
Network"
```

11. Close Windows PowerShell.
12. In the Hyper-V Manager console, click **LON-GUEST2**.
13. In the Actions pane, under **LON-GUEST2**, click **Settings**.
14. In the **Settings for LON-GUEST2 on LON-HOST1** dialog box, click **Automatic Start Action**, and set the Automatic Start Action to **Nothing**.
15. In the **Settings for LON-GUEST2 on LON-HOST1** dialog box, click **Automatic Stop Action**, and set the Automatic Stop Action to **Shut down the guest operating system**.
16. Click **OK** to close the **Settings for LON-GUEST2 on LON-HOST1** dialog box.

► Task 3: Enable resource metering

1. On the taskbar, click the **Windows PowerShell** icon.
2. At the Windows PowerShell prompt, enter the following commands to enable resource metering on the virtual machines, pressing Enter at the end of each line:

```
Enable-VMResourceMetering LON-GUEST1
Enable-VMResourceMetering LON-GUEST2
```

Results: After completing this exercise, you should have deployed two separate virtual machines using a sysprepped virtual hard disk file as a parent disk for two differencing virtual hard disks.

Exercise 4: Using Virtual Machine Checkpoints

► Task 1: Deploy Windows Server 2012 in a virtual machine

1. In the Hyper-V Manager console, click **LON-GUEST1**.
2. In the Actions pane, click **Start**.
3. Double-click **LON-GUEST1** to open the Virtual Machine Connection Window.
4. In the LON-GUEST1 on LON-HOST1 - Virtual Machine Connection window, perform the following steps:
 - On the **Settings** page, click **Next** to accept the Region and Language settings.
 - On the **Settings** page, click **I Accept**.
 - On the **Settings** page, type the password **Pa\$\$w0rd** twice, and then click **Finish**.

5. In the LON-GUEST1 on LON-HOST1 - Virtual Machine Connection window, from the **Action** menu, click **CTRL+Alt+Delete**.
6. Sign in to the virtual machine using the account **Administrator** and the password **Pa\$\$w0rd**.
7. On the virtual machine, in the Server Manager console, click **Local Server**, and then click the randomly assigned name next to the computer name.
8. In the **System Properties** dialog box, on the **Computer Name** tab, click **Change**.
9. In the **Computer Name** field, type **LON-GUEST1**, and then click **OK**.
10. In the **Computer Name/Domain Changes** dialog box, click **OK**.
11. Click **Close** to close the **System Properties** dialog box.
12. In the **Microsoft Windows** dialog box, click **Restart Now**.

► **Task 2: Create a virtual machine checkpoint**

1. Sign in to the LON-GUEST1 virtual machine using the **Administrator** account and the password **Pa\$\$w0rd**.
2. In the Server Manager console, click the **Local Server** node, and verify that the name of the computer is set to **LON-GUEST1**.
3. In the Virtual Machine Connection window, from the **Action** menu, click **Checkpoint**.
4. In the **Checkpoint Name** dialog box, type the name **Before Change**, and then click **Yes**.

► **Task 3: Modify the virtual machine**

1. In the Server Manager console, click **Local Server**, and then next to **Computer name**, click **LON-GUEST1**.
2. In the **System Properties** dialog box, on the **Computer Name** tab, click **Change**.
3. In the **Computer Name** field, type **LON-Computer1**, and then click **OK**.
4. In the **Computer Name/Domain Changes** dialog box, click **OK**.
5. Close the **System Properties** dialog box.
6. In the **Microsoft Windows** dialog box, click **Restart Now**.
7. Sign back in to the **LON-GUEST1** virtual machine using the **Administrator** account and the password **Pa\$\$w0rd**.
8. In the Server Manager console, click **Local Server**, and verify that the server name is set to **LON-Computer1**.

► **Task 4: Revert to the existing virtual machine checkpoint**

1. In the Virtual Machine Connection window, from the **Action** menu, click **Revert**.
2. In the **Revert Virtual Machine** dialog box, click **Revert**.
3. In the Server Manager console, in the Local Server node, in the **Virtual Machines** list, verify that the **Computer Name** is now set to **LON-GUEST1**.

► Task 5: View resource metering data

1. On LON-HOST1, on the taskbar, click the **Windows PowerShell** icon.
2. To retrieve resource metering information, at the Windows PowerShell prompt, enter the following command, and then press Enter:

```
Measure-VM LON-GUEST1
```

Note the average CPU, average random access memory (RAM), and total disk usage figures.

3. Close the Windows PowerShell window.

Results: After completing this exercise, you should have used virtual machine checkpoints to recover from a virtual machine misconfiguration.

► Revert the virtual machines

After you finish the lab, restart the computer in Windows Server 2012 by performing the following steps:

1. On the taskbar, click the **Windows PowerShell** icon.
2. At the Windows PowerShell command prompt, type the following command, and then press Enter:

```
Shutdown /r /t 5
```

3. From the Windows Boot Manager, select **Windows Server 2012**.

