Microsoft

OFFICIAL MICROSOFT LEARNING PRODUCT

20412C

Configuring Advanced Windows Server® 2012 Services

Product Number: 20412C

Part Number (if applicable): X19-09980

Released: January 28, 2014

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 silver or gold-level 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 Programs and Services**. The Licensed Content may contain third party programs or services. These license terms will apply to your use of those third party programs or services, unless other terms accompany those programs and services.

2.5 **Additional Terms**. Some Licensed Content may contain components with additional terms, conditions, and licenses regarding its use. Any non-conflicting terms in those conditions and licenses also apply to 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 software, Microsoft product, or service that includes the feedback. You will not give feedback that is subject to a license that requires Microsoft to license its software, technologies, or products to third parties because we include your feedback in them. These rights survive this agreement.

c. **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 dues 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 September 2012

Welcome!

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

- **Microsoft Certified Trainers and Instructors**—Your instructor is a technical and instructional expert who meets ongoing certification requirements. And, if instructors are delivering training at one of our Certified Partners for Learning Solutions, they are also evaluated throughout the year by students and by Microsoft.

- **Certification Exam Benefits**—After training, consider taking a Microsoft Certification exam. Microsoft Certifications validate your skills on Microsoft technologies and can help differentiate you when finding a job or boosting your career. In fact, independent research by IDC concluded that 75% of managers believe certifications are important to team performance[1]. 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 individuals for their contribution toward developing this title. Their effort at various stages in the development has ensured that you have a good classroom experience.

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 Exchange Server and Active Directory 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 Microsoft Certified Trainer (MCT) for 14 years.

Byron Wright – Subject Matter Expert

Byron Wright is a partner in a consulting firm, where he performs network consulting, computer-systems Implementation, and technical training. Byron is also a sessional instructor for the Asper School of Business at the University of Manitoba, where he teaches management information systems and networking. Byron has authored and coauthored a number of books on Windows servers, Windows clients, and Exchange Server, including the Windows Server 2008 Active Directory Resource Kit. To recognize Byron's commitment to sharing knowledge with the technical community, he has been given the Microsoft MVP Award for Exchange Server.

Damir Dizdarevic– Subject Matter Expert

Damir Dizdarevic is an MCT, Microsoft Certified Solutions Expert (MCSE), Microsoft Certified Technology Specialist (MCTS), and a Microsoft Certified Information Technology Professional (MCITP). He is a manager and trainer of the Learning Center at Logosoft d.o.o., in Sarajevo, Bosnia, and Herzegovina. He also works as a consultant on IT infrastructure and messaging projects. Damir has more than 17 years of experience on Microsoft platforms, and he specializes in Windows Server, Exchange Server, security, and virtualization. He has worked as a subject matter expert and technical reviewer on many Microsoft Official Courses (MOC) courses, and has published more than 400 articles in various IT magazines, such as Windows IT Pro and INFO Magazine. He's also a frequent and highly rated speaker on most of Microsoft conferences in Eastern Europe. Additionally, Damir is a Microsoft Most Valuable Professional (MVP) for Windows Server, seven years in a row. His technical blog is available at http://dizdarevic.ba/ddamirblog.

Orin Thomas – Subject Matter Expert

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

David M. Franklyn – Subject Matter Expert

David M. Franklyn, MCT, MCSE, Microsoft Certified IT Professional (MCITP), Microsoft Most Valuable Professional (MVP) Windows Expert--II 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. Dave began working with in 1976, when he started out in the mainframe world and moved early into the networking arena. Before joining

Auburn University, Dave spent 22 years in the U.S. Air Force as an electronic communications and computer systems specialist, retiring in 1998. Dave is president of the Montgomery Windows IT Professional Group, and is a guest speaker at many events involving Microsoft products.

Gary Dunlop – Subject Matter Expert

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

David Susemiehl – Subject Matter Expert

David Susemiehl 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 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.

Ulf B. Simon-Weinder – Technical Reviewer

Ulf B. Simon-Weidner got his first jobs in digital electronics and microprocessor programming, then moved into programming and building network infrastructures, the area in which he has worked for more than 20 years. . He is also an independent author, consultant, speaker, and trainer. Ulf has received the yearly award as Microsoft Most Valuable Professional (MVP) for Windows Server – Directory Services 10 times, and has been a Microsoft Certified Trainer since 2001. Throughout his professional career, he has had numerous consulting engagements with major European or global corporations. He also published many books and articles about Active Directory, Windows Server Infrastructures, Client and Security. Ulf is a frequent visiting speaker for conferences such as Microsoft TechEd North America and Europe, the Directory Experts Conference and The Experts Conference.

Contents

About This Course

This course is intended for information technology (IT) professionals who have hands-on experience implementing, managing, and maintaining a Windows Server 2012 or Windows Server 2012 R2 environment who wish to acquire the skills and knowledge necessary to perform advanced services management and provisioning within that Windows Server 2012 environment.

Course Description

Get hands-on instruction and practice configuring advanced Windows Server 2012, including Windows Server 2012 R2, services in this five-day Microsoft Official Course. This course is the third part in a series of three courses 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. Although there is some cross-over of skills and tasks across these courses, this course focuses on advanced configuration of services necessary to deploy, manage, and maintain a Windows Server 2012 infrastructure, such as advanced networking services, Active Directory Domain Services (AD DS), Active Directory Rights Management Services (AD RMS), Active Directory Federation Services (AD FS), Network Load Balancing, failover clustering, business continuity, and disaster-recovery services. This course also covers access and information provisioning, and protection technologies such as Dynamic Access Control (DAC), and Web Application Proxy integration with ADFS and Workplace Join.

This course maps directly to and is the preferred choice for hands-on preparation for Microsoft Certified Solutions Associate (MCSA): Exam 412: Configuring Advanced Windows Server 2012 Services, which is the third of three exams required for MCSA: Windows Server 2012 certification.

Note: Labs in this course are based on the General Availability release of Windows Server 2012 R2 and Windows 8.1.

Module 1 starts the course with topics on advanced network configuration. Students will already be familiar with Domain Name System (DNS) and Dynamic Host Configuration Protocol (DHCP) services, and this course is designed for more advanced configurations that they may not have encountered. IP Address Management (IPAM) is a new Windows Server 2012 feature that will help students streamline the management of IP addressing in the organization.

Modules 2 and 3 provide a block of topics that are focused on file services. Module 2 expands on previous knowledge that students have acquired on how to configure file services in a Windows Server environment by introducing some advanced configuration options. Module 3 describes the new Windows Server 2012 feature that provides even more advanced options for managing and auditing access to file server resources in Windows Server 2012.

Modules 4 through 8 discuss the more advanced topics in implementing AD DS and other Active Directory role services. Modules 4 and 5 describe the scenario where an organization has a highly complicated environment that cannot be easily managed with a single AD DS domain and site. Therefore, these modules describe how to implement multi-domain and multi-site AD DS environments.

Modules 6 through 8 take AD DS implementation in a different direction. While modules 4 and 5 focused on providing AD DS services to users inside the organization, modules 6 to 8 switch the focus to providing some AD DS services outside of the organization. This includes authentication and authorization to users or services that might be in the same forest, but that might also be in a different AD DS forest, or might not even have any AD DS accounts.

Module 6 describes how to implement a public key infrastructure (PKI) environment that will meet internal certificate services requirements and external requirements. Module 7 describes how to implement an Active Directory Rights Management Services (AD RMS) deployment to enable internal access restrictions to be extended outside the organization's boundaries. Module 8 describes how to implement Active Directory Federation Services (AD FS) environments to extend authentication services to users who might not have any accounts in the internal AD DS forest.

Modules 9 and 10 provide details on two different options for making applications and services highly available in a Windows Server 2012 environment. Module 9 describes Network Load Balancing (NLB), which is used primarily for web-based applications. Module 10 describes failover clustering, which can be used to make many other applications and services highly available. Module 11 expands on the failover clustering content from Module 10, by describing how to integrate Hyper-V™ virtual machines with failover clustering.

Module 12 provides instruction on how to plan for and recover from various data and server loss scenarios in Windows Server 2012. Because of the options for integrating high availability with disaster recovery, this module will build on the high-availability content that was presented in the previous modules, but will also include scenarios and procedures for ensuring data and service availability in the event of failure in a highly available environment.

Audience

This course is intended for candidates who would typically be experienced Windows Server Administrators who have real-world experience working in a Windows Server 2008 or Windows Server 2012 enterprise environment. The audience also includes IT professionals who want to take the course 70-412, Configuring Advanced Windows Server 2012 Services. Lastly, the audience includes IT professionals who wish to take the Microsoft Certified Solutions Expert (MCSE) exams in DataCenter, Desktop Infrastructure, Messaging, Collaboration and Communications. This course may help them 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 you meet the following prerequisites:

- Experience working with Windows Server 2008 or Windows Server 2012 servers day to day in an enterprise environment.
- Knowledge equivalent to the content covered in courses 20410C: Installing and Configuring Windows Server 2012; and 20411C: Administering Windows Server 2012.

Course Objectives

After completing this course, the students will be able to:

- Configure advanced features for DHCP and DNS, and configure IP address management.
- Configure file services to meet advanced business requirements.
- Configure Dynamic Access Control (DAC) to manage and audit access to shared files.
- Plan and implement an AD DS deployment that includes multiple domains and forests.
- Plan and implement an AD DS deployment that includes multiple locations and data centers.
- Implement and configure an Active Directory Certificate Services (AD CS) deployment.
- Implement and configure an Active Directory Rights Management Services (AD RMS) deployment.

- Implement and configure an Active Directory Federation Services (AD FS) deployment.
- Provide high availability and load balancing for web-based applications by implementing Network Load Balancing (NLB).
- Provide high availability for network services and applications by implementing failover clustering.
- Deploy and manage Windows Server 2012 Hyper-V virtual machines in a failover cluster.
- Implement a backup and disaster-recovery solution based on business and technical requirements.

Course Outline

The course outline is as follows:

- Module 1: "Implementing Advanced Network Services"
- Module 2: "Implementing Advanced File Services"
- Module 3: "Implementing Dynamic Access Control"
- Module 4: "Implementing Distributed Active Directory Domain Services Deployments"
- Module 5: "Implementing Active Directory Domain Services Sites and Replication"
- Module 6: "Implementing AD CS"
- Module 7: "Implementing Active Directory Rights Management Services"
- Module 8: "Implementing and Administering AD FS"
- Module 9: "Implementing Network Load Balancing"
- Module 10: "Implementing Failover Clustering"
- Module 11: "Implementing Failover Clustering with Hyper-V"
- Module 12: "Implementing Business Continuity and Disaster Recovery"

Exam/Course Mapping

This course, *20412C: Configuring Advanced Windows Server® 2012 Services*, has a direct mapping of its content to the objective domain for the Microsoft exam *70-412: Configuring Advanced Windows Server 2012 Services*.

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-70-412.aspx,%20under%20Skills%20Measured.

Exam Objective Domain: 70-412: Configuring Advanced Windows Server 2012 Services		Course Content		
1. Configure and Manage High Availability (16%)		**Module**	**Lesson**	**Lab**
1.1 Configure Network Load Balancing (NLB).	This objective may include but is not limited to: Installing NLB nodes; configuring NLB prerequisites; configuring affinity; configuring port rules; configuring cluster operation mode; upgrading an NLB cluster	Mod 9	Lesson 1/2/3	Mod 9 Ex 1/2/3
1.2 Configure failover clustering.	This objective may include but is not limited to: Configuring Quorum; configuring cluster networking; restoring single node or cluster configuration; configuring cluster storage; implement Cluster Aware Updating; upgrade a cluster ; configure and optimize clustered shared volumes; configure clusters without network names; configure storage spaces	Mod 10	Lesson 1/2/3/4/5	Mod 10 Ex 1/2/3/4
1.3 Manage failover clustering roles.	This objective may include but is not limited to: Configuring role-specific settings including continuously available shares; configure VM monitoring ; configuring failover and preference settings; configure guest clustering	Mod 10	Lesson 1/3	Mod 10 Lab Ex 2
		Mod 11	Lesson 1/2	Mod 11 Lab Ex 1/2
1.4 Manage virtual machine (VM) movement.	This objective may include but is not limited to: Perform live migration; perform quick migration; performing storage migration; import, export, and copy VMs; configure Virtual Machine network health protection; configure drain on shutdown	Mod 11	Lesson 2/3	Mod 11 Lab Ex 2/3

2. Configure File and Storage Solutions (18%)				
2.1 Configure advanced file services.	This objective may include but is not limited to: Configuring Network File System (NFS) data store; configuring BranchCache; configuring File Classification Infrastructure (FCI) using the File Server Resource Manager (FSRM); configuring file access auditing	Mod 2	Lesson 2/3	Mod 2 Lab A Ex 2 Lab B Ex 1/2/3/4
2.2 Implement Dynamic Access Control (DAC).	This objective may include but is not limited to: Configuring user and device claim types; implementing policy changes and staging; performing access-denied remediation; configuring file classification; create and configure Central Access rules and policies; create and configure resource properties and lists	Mod 3	Lesson 1/2/3	Mod 3 LAB Ex 1/2/3
2.3 Configure and optimize storage.	This objective may include but is not limited to: Configuring iSCSI Target and Initiator; configuring Internet Storage Name Server (iSNS); implementing thin provisioning and trim; managing server free space using Features on Demand; configure tiered storage	Mod 2	Lesson 1/3	Mod 2 Lab Ex 1
3 Implement Business Continuity and Disaster Recovery (14%)				
3.1 Configure and manage backups.	This objective may include but is not limited to: Configuring Windows Server backups; configuring Windows Azure backups; configuring role-specific backups; managing VSS settings using VSSAdmin;	Mod 12	Lesson 1/2	Mod 12 Lab Ex 1/2

3.2 Recover servers.	This objective may include but is not limited to: Restore from backups; perform a Bare Metal Restore (BMR); recover servers using Windows Recovery Environment (Win RE) and safe mode; configuring the Boot Configuration Data (BCD) store	Mod 12	Lesson 1/2/3	Mod 12 Ex 1/2
3.3 Configure site-level fault tolerance.	This objective may include but is not limited to: Configuring Hyper-V Replica including Hyper-V Replica Broker and VMs; configuring multi-site clustering including network settings, Quorum, and failover settings; configure Hyper-V Replica extended replication; configure Global Update Manager; recover a multi-site failover cluster	Mod 11	Lessons 1/3	Mod 11 Lab Ex 1
		Mod 10	Lesson 1	Mod 10 Lab Ex 1
4. Configure Network Services (17%)				
4.1 Implement an advanced Dynamic Host Configuration Protocol (DHCP) solution.	This objective may include but is not limited to: Create and configure superscopes and multicast scopes; implementing DHCPv6; configuring high availability for DHCP including DHCP failover and split scopes; configuring DHCP Name Protection; configure DNS registration	Mod 1	Lesson 1	Mod 1 Lab Ex 1
4.2 Implement an advanced DNS solution.	This objective may include but is not limited to: Configuring security for DNS including DNSSEC, DNS Socket Pool, and cache locking; configuring DNS logging; configuring delegated administration; configuring recursion; configuring netmask ordering; configuring a GlobalNames zone; analyze zone level statistics	Mod 1	Lesson 2	Mod 1 Lab Ex 2

4.3 Deploy and manage IPAM.	This objective may include but is not limited to: Provision IPAM manually or by using Group Policy ; configuring server discovery; creating and managing IP blocks and ranges; monitoring utilization of IP address space; migrate to IPAM; delegate IPAM administration; manage IPAM collections; configure IPAM database storage	Mod 1	Lesson 3	Mod 1 Lab Ex 3
5. Configure the Active Directory Infrastructure (15%)				
5. 1 Configure a forest or a domain	This objective may include but is not limited to: Implement multi-domain and multi-forest Active Directory environments including interoperability with previous versions of Active Directory; upgrade existing domains and forests including environment preparation and functional levels; configuring multiple user principal name (UPN) suffixes	Mod 4	Lesson 1/2	Mod 4 Lab Ex 1
5.2 Configure trusts.	This objective may include but is not limited to: Configuring external, forest, shortcut, and realm trusts; configuring trust authentication; configuring SID filtering; configuring name suffix routing	Mod 4	Lesson 3	Mod 4 Lab Ex 2
5.3 Configure sites.	This objective may include but is not limited to: Configure sites and subnets; create and configure site links; manage site coverage; manage registration SRV records; move domain controllers between sites	Mod 5	Lesson 2/3	Mod 5 Lab Ex 1/2

5.4 Manage Active Directory and System Volume (SYSVOL) replication.	This objective may include but is not limited to: Configuring replication to read-only domain controllers (RODCs); configuring Password Replication Policy (PRP) for RODCs; monitoring and managing replication; upgrading SYSVOL replication to Distributed File System Replication (DFSR)	Mod 5	Lesson 1/3	Mod 5 Lab Ex 3/4
6. Configure Access and Information Protection Solutions (19%)				
6.1 Implement Active Directory Federation Services (AD FS).	This objective may include but is not limited to: Install AD FS; Implement claims-based authentication including Relying Party Trusts; configure authentication policies; configure Workplace Join; configure multi-factor authentication	Mod 8	Lesson 1/2/3/4/5	Mod 8 Lab A Ex 1/2/3 Lab B Ex 1/2
6.2 Install and configure Active Directory Certificate Services (AD CS).	This objective may include but is not limited to: Install an Enterprise Certificate Authority (CA); Configure CRL distribution points; install and configure Online Responder; implement administrative role separation; configuring CA backup and recovery	Mod 6	Lesson 1/2/3	Mod 6 Lab A Ex 1/2 Lab B
6.3 Manage certificates.	This objective may include but is not limited to: Manage certificate templates; implement and manage certificate deployment, validation, and revocation; manage certificate renewal; managing certificate enrollment and renewal to computers and users using Group Policies; configure and manage key archival and recovery	Mod 6	Lesson 4/5/6	Mod 6 Lab B Ex 1/2/3/4

6.4 Install and configure Active Directory Rights Management Services (AD RMS).	This objective may include but is not limited to: Installing a licensing or certificate AD RMS server; managing AD RMS Service Connection Point (SCP); managing RMS templates; configuring Exclusion Policies; backup and restore AD RMS	Mod 7	Lesson 1/2/3/4	Mod 7 Lab Ex 1/2/3/4

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-70-412.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=20412C, 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 the critical technical information in a crisp, tightly focused format, which is essential 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 on-the-job reference material to boost knowledge and skills retention.
- **Lab Answer Keys**: provide step-by-step lab solution guidance.

Course Companion Content on the http://www.microsoft.com/learning/en/us/companion-moc.aspx Site*:* searchable, easy-to-browse digital content with integrated premium online resources that 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 current premium content on TechNet, MSDN®, or Microsoft Press®.

Note: For this version of the Courseware, Companion Content is not available. However, the Companion Content will be published when the next (D) version of this course is released; and students who have taken this course will be able to download the Companion Content at that time from the **http://www.microsoft.com/learning/en/us/companion-moc.aspx** site. Please check with your instructor when the D version is scheduled for release to find out when you can access the Companion Content.

- ***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 an email to support@mscourseware.com. To inquire about the Microsoft Certification Program, send an email 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 may need to 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 that is used in this course:

Virtual machine	Role
20412C-LON-DC1/-B	Windows Server 2012 Domain controller in the Adatum.com domain
20412C-LON-CA1	Windows Server 2012 Standalone server
20412C-LON-CL1	Windows 8 client computer Member of the Adatum.com domain
20412C-LON-CL2	Windows 8 client computer Member of the Adatum.com domain
20412C-LON-CORE	Windows Server 2012 Member server in the Adatum.com domain
20412C-LON-SVR1/-B	Windows Server 2012 Member server in the Adatum.com domain
20412C-LON-SVR2	Windows Server 2012 Member server in the Adatum.com domain
20412C-LON-SVR3	Windows Server 2012 Member server in the Adatum.com domain
20412C-LON-SVR4	Windows Server 2012 Member server in the Adatum.com domain

Virtual machine	Role
20412C-TREY-CL1	Windows 8 client computer Member of the Treyresearch.net domain
20412C-TREY-DC1	Windows Server 2012 Domain controller in the Treyresearch.net domain
20412C-LON-HOST1	Windows Server 2012 Member server in the Adatum.com domain
20412C-LON-HOST2	Windows Server 2012 Member server in the Adatum.com domain
20412C-TOR-DC1	Windows Server 2012 Member server in the Adatum.com domain

Software Configuration

The following software is installed in the course

- Windows Server 2012 R2
- Windows 8.1
- Microsoft Office 2013
- Windows Identity Foundation SDK 4.0

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 is taught. This includes:

- Intel Virtualization Technology (Intel VT) or AMD Virtualization (AMD-V) processor
- Dual 120 gigabyte (GB) hard disks 7200 RM Serial ATA (SATA) or better*
- 16 GB RAM

- DVD drive
- Network adapter
- Super VGA (SVGA) 17-inch monitor
- Microsoft Mouse or compatible pointing device
- Sound card with amplified speakers

*Striped

In addition, the instructor computer must be connected to a projection display device that supports SVGA 1024 x 768 pixels, 16-bit colors.

Module 1

Implementing Advanced Network Services

Contents:

Module Overview

In Windows Server® 2012, network services such as Domain Name System (DNS) provide critical support for name resolution of network and Internet resources. Within DNS, DNS Security Extensions (DNSSEC) is an advanced feature that provides a means of securing DNS responses to client queries so that malicious users cannot tamper with them. With Dynamic Host Configuration Protocol (DHCP), you can manage and distribute IP addresses to client computers. DHCP is essential for managing IP-based networks. DHCP failover is an advanced feature that can prevent clients from losing access to the network in case of a DHCP server failure. IP Address Management (IPAM) provides a unified means of controlling IP addressing.

This module introduces DNS and DHCP improvements, IP address management, and provides details about how to implement these features.

Objectives

After completing this module, you will be able to:

- Configure advanced DHCP features.
- Configure advanced DNS settings.
- Implement IPAM.

Lesson 1
Configuring Advanced DHCP Features

DHCP plays an important role in the Windows Server 2012 operating system 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 and Network Access Protection (NAP). To support a Windows Server-based network infrastructure, it is important that you understand the DHCP server role. Windows Server 2012 improves the functionality of DHCP by providing failover capabilities.

Lesson Objectives

After completing this lesson you will be able to:

- Describe DHCP components.
- Explain how to configure DHCP interaction with DNS.
- Explain how to configure advanced DHCP scope designs.
- Explain how DHCP works with IPv6.
- Describe DHCP name protection.
- Describe DHCP failover.
- Explain how to configure DHCP failover.

DHCP Components Overview

DHCP is a server role that you can install on Windows Server 2012. With the DHCP server role, you can ensure that all clients have appropriate IP addresses and network configuration information, which can help eliminate human error during configuration. A *DHCP client* is any device that is taking a DHCP address, and that can request and retrieve network settings from a DHCP server service. DHCP clients may be computers, mobile devices, printers, or switches. DHCP may also provide IP address information to network boot clients.

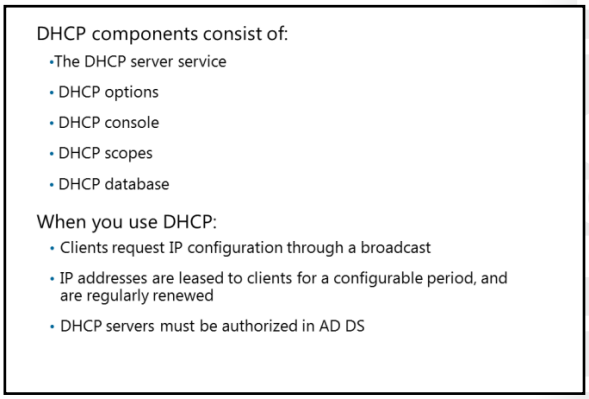

When key network configuration information changes in the network, (such as the default gateway address), you can update the configuration using the DHCP server role without having to change the information directly on each computer. DHCP is also a key service for mobile users who change networks often. You can install the DHCP Server role on a stand-alone server, a domain member server, or a domain controller.

DHCP consists of the components that are listed in the following table.

Component	Description
DHCP Server service	After installing the DHCP Server role, the DHCP server is implemented as a service. This service can distribute IP addresses and other network configuration information to clients who request it.
DHCP scopes	The DHCP administrator configures the range of IP addresses and related information that is allotted to the server for distribution to requesting clients. Each scope can only be associated with a single IP subnet. A scope must consist of: • A name and description • A range of addresses that can be distributed • A subnet mask A scope can also define: • IP addresses that should be excluded from distribution • The duration of the IP address lease • DHCP options You can configure a single DHCP server with multiple scopes, but the server must be either connected directly to each subnet that it serves, or have a supporting and configured DHCP relay agent in place. Scopes also provide the primary way for the server to manage and distribute any related configuration parameters (DHCP options) to clients on the network.
DHCP options	When you assign the IP address to the client, you can also simultaneously assign many other network configuration parameters. The most common DHCP options include: • Default Gateway IP address • DNS server IP address • DNS domain suffix • Windows Internet Name Service (WINS) server IP address You can apply the options at different levels. They can be applied as follows: • Globally to all scopes • Specifically to particular scopes • To specific clients based on a class ID value • To clients that have specific IP address reservations configured  **Note:** IPv6 scopes are slightly different, and will be discussed later in this lesson.
DHCP database	The DHCP database contains configuration data about the DHCP server, and stores information about the IP addresses that have been distributed. By default, the DHCP database files are stored in the %systemroot%\System32\Dhcp folder.

Component	Description
DHCP console	The DHCP console is the main administrative tool for managing all aspects of the DHCP server. This management console is installed automatically on any server that has the DHCP role installed. However, you can also install it on a remote server or Windows 8 client by using the Remote Server Administration Tools (RSAT) and by connecting to the DHCP server for remote management.

How Clients Acquire IP Addresses

When you configure a Windows client operating system to use the DHCP service, upon startup the client will use an ARP broadcast in its subnet to request IP configuration from any DHCP server that may receive the request. Because DHCP uses broadcasts to initiate communications, DHCP servers are limited to communication within their IP subnets. This means that either there must be a DHCP server on each IP subnet, or a router configured to forward BOOTP traffic DHCP relay agent configured on the remote subnet. The DHCP relay service, or BOOTP forwarding, can relay DHCP broadcast packets as directed messages into other IP subnets across a router. The relay agent acquires an IP address configuration on behalf of the requesting client on the remote subnet, and then forwards that configuration to the client.

DHCP Leases

DHCP allocates IP addresses on a dynamic basis. This is known as a *lease*. You can configure the duration of the lease. The default lease time for wired clients is eight days.

When the DHCP lease has reached 50 percent of the lease time, the client attempts to renew the lease. This automatic process occurs in the background. Computers might have the same IP address for a long time if they operate continually on a network without being shut down. Client computers also attempt renewal during the startup process.

DHCP Server Authorization

If the server is a domain member, you must authorize the Windows Server 2012 DHCP server role in Active Directory Domain Services (AD DS) before it can begin leasing IP addresses. You must be an Enterprise Administrator to authorize the DHCP server. Stand-alone Microsoft servers verify whether there is a DHCP server on the network, and do not start the DHCP service if this is the case.

Windows PowerShell

Windows PowerShell® cmdlets are used to provide command line support for managing DHCP. In addition to providing command-line support, PowerShell cmdlets are used if you want to script your DHCP management. A subset of the nearly 100 Windows Server 2012 PowerShell cmdlets for managing DHCP are included in the following table:

cmdlet	Additional information
Add-DhcpServerInDC	You use cmdlet to add the specified computer running the DHCP server service as an authorized DHCP server in AD DS.
Add-DhcpServerv4Class	You use this cmdlet to add an IPv4 vendor or user class to the DHCP server service.
Add-DhcpServerv4ExclusionRange	You use this cmdlet to add an IP address exclusion range to an IPv4 scope.
Add-DhcpServerv4Failover	You use this cmdlet to add a new IPv4 failover relationship on the DHCP server service.

cmdlet	Additional information
Add-DhcpServerv4FailoverScope	You use this cmdlet to add one or more scopes to an existing failover relationship.
Add-DhcpServerv4Filter	You use this cmdlet to add a media access control (MAC) address filter of the DHCP server service, the filter can be used on an allow list or Deny list.
Add-DhcpServerv4Lease	You use this cmdlet to add a new IPv4 address lease in the DHCP server service for testing purposes.
Add-DhcpServerv4OptionDefinition	You use this cmdlet to add a new DHCPv4 option definition to the DHCP server service.
Add-DhcpServerv4Policy	You use this cmdlet to add a new IPv4 policy to a DCHP server or a DHCP scope.
Add-DhcpServerv4PolicyIPRange	You use this cmdlet to add an IP range to an existing scope policy.
Add-DhcpServerv4Reservation	You use this cmdlet to reserve the specified IPv4 address in the specified DHCP scope for a specified client.
Add-DhcpServerv4Scope	You use this cmdlet to add an IPv4 scope on the DHCP server service.

For a complete list of the available cmdlets see, DHCP Server Cmdlets in Windows PowerShell:

http://go.microsoft.com/fwlink/?LinkID=386639

Windows Server® 2012 R2 added or improved the DHCP cmdlets for additional functionality and to support the features added in Windows Server 2012 R2. The following table lists some of the cmdlets that have been added or improved:

Cmdlet	New or Improved	Additional information
Add-DhcpServerSecurityGroup	New	You use this cmdlet to add security groups to a DHCP server.
Add-DhcpServerv4MulticastExclusionRange	New	You use this cmdlet to add an IP address exclusion range to a multicast scope.
Add-DhcpServerv4MulticastScope	New	You use this cmdlet to add a multicast scope on the DHCP server.
Add-DhcpServerv4Policy	Improved	You use this cmdlet to add a new policy to either a server or a scope. This cmdlet has been improved so that it can now be used to specify lease duration and also add fully qualified domain name (FQDN)-based policies.

Cmdlet	New or Improved	Additional information
Get-DhcpServerDnsCredential	New	You use this cmdlet to get the credentials for an account that the DHCP Server service uses to register or deregister client records on a DNS server.
Get-DhcpServerv4DnsSetting	Improved	You can now use this cmdlet to display DNS settings of DHCP policies.
Get-DhcpServerv4MulticastExclusionRange	New	You use this cmdlet to retrieve the exclusion range for a specified multicast scope.
Get-DhcpServerv4MulticastLease	New	You use this cmdlet to retrieve multicast leases for a specified scope name.
Get-DhcpServerv4MulticastScope	New	You use this cmdlet to get information on multicast scope objects.
Get-DhcpServerv4MulticastScopeStatistics	New	You use this cmdlet to get information on multicast scope statistics.

For information about the DHCP cmdlets that were added or improved Windows Server 2012 see, What's New in DHCP in Windows Server 2012 R2:

http://go.microsoft.com/fwlink/?LinkID=386638

Configuring DHCP Interaction with DNS

During dynamic IP address allocation, the DHCP server creates resource records automatically for DHCP clients in the DNS database. However, those records may not be deleted automatically when the client DHCP lease expires. You can configure DHCP options to allow the DHCP server to own and fully control the creation and deletion of those DNS resource records.

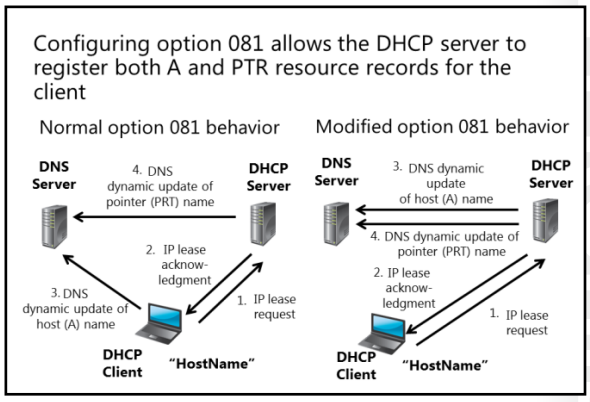

Configuring Dynamic DNS Updates

You can configure the DHCP service to control the way that resource records are updated in the DNS database. The default setting for the Enable DNS dynamic updates, according to the settings below option (DHCP option 081), permits the client to provide its fully qualified domain name (FQDN) and instructions to the DHCP server about how it would like the server to process DNS dynamic updates on its behalf. You configure this option on the DNS tab of the Properties dialog box for the protocol node (IPv4 or IPv6), or per scope in the DHCP console. You can also

configure DHCP to perform updates on behalf of its clients to any DNS servers that support dynamic updates.

By default, the DHCP server behaves in the following manner:

- The DHCP server dynamically updates DNS address host (A) resource records and pointer (PTR) resource records only if requested by the DHCP clients. By default, the client requests that the DHCP server register the DNS pointer (PTR) resource record, while the client registers its own DNS host (A) resource record.
- The DHCP server discards the host (A) and pointer (PTR) resource records when the client's lease is deleted.

You can change the Enable DNS dynamic updates according to the settings below option to Always dynamically update DNS records so that it instructs the DHCP server to always dynamically update DNS host (A) and pointer (PTR) resource records no matter what the client requests. In this way, the DHCP server becomes the owner of the resource record because the DHCP server performed the registration of the resource records. Once the DHCP server becomes the owner of the client computer's host (A) and pointer (PTR) resource records, only that DHCP server can update the DNS resource records for the client computer based on the duration and renewal of the DHCP lease.

The following Windows PowerShell cmdlets can be used to manage DHCP Scopes:

cmdlet	Description
Add-DhcpServerv4Scope	You use this cmdlet to add a new IPv4 policy to a DCHP server or a DHCP scope.
Get-DhcpServerv4Scope	You use this cmdlet to view the IPv4 scope configuration of the specified scopes.
Remove-DhcpServerv4Scope	You use this cmdlet to delete the specified IPv4 scopes from the DHCP server service.
Set-DhcpServerv4Scope	You use this cmdlet to configure the properties of an existing IPv4 scope on the DHCP server service.

DHCP Policies

Windows Server 2012 introduced DHCP policies to allow more control over the settings assigned to a client. A DHCP policy consists of the Conditions being evaluated and the Settings that will be applied if the conditions are met. Conditions can be configured as either Equals or Not Equals to the value you specify. The following table shows the conditions that can be configured in a DHCP policy.

Condition	Additional Information
MAC Address	Uses the MAC address of the DHCP requesting client
Vendor Class	Vendor classes are configured on the protocol node (IPv4 or IPv6)
User Class	User classes are configured on the protocol node (IPv4 or IPv6)
Client Identifier	Hex-based client identifier
Fully Qualified Domain Name	Introduced with Windows 2012 R2
Relay Agent Information	A relay agent can insert information, such as the client's network ID, into a DHCP request. Also known as option 082.

DHCP Policy can contain multiple conditions joined together using either an And or an Or operator. Once the conditions are specified, you can configure an address range and the DHCP options to apply to the clients. If a client does not match any conditions specified in any policies they will receive the options specified for the appropriate scope and protocol.

You can use the following Windows PowerShell cmdlets to manage DHCP policies:

cmdlet	Description
Add-DhcpServerv4Policy	You use this cmdlet to add a new IPv4 policy either to a server or to a scope. Windows Server 2012 R2 added options for lease duration and for adding FQDN-based policies.
Set-DhcpServerv4DnsSetting	You use this cmdlet to configure how the DHCP server service updates the DNS server. Windows Server 2012 R2 added the ability to set the DNS settings of policies.
Set-DhcpServerv4Policy	You use this cmdlet to configure the properties of an existing policy either on a server or in a scope. Windows Server 2012 R2 added the ability to configure policy lease duration and also modify FQDN-based policies.

DNS PTR Registration

By default the DHCP server dynamically updates both the client computer's host (A) and pointer (PTR) resource records. If you have not configured a DNS reverse lookup zone for the IP address range being distributed, you can disable DNS registration. Prior to Windows Server 2012 R2, you had to disable both host (A) and pointer (PTR) resource record registration. Windows Server 2012 R2 introduces the ability to disable just the PTR record registration. You configure the Disable dynamic updates for DNS PTR records option on the DNS tab of the Properties dialog box for the protocol node (IPv4 or IPv6), or per scope in the DHCP console.

Configuring Advanced DHCP Scope Designs

You can configure advanced DHCP scope designs called *superscopes*. A superscope is a collection of individual scopes that are grouped together for administrative purposes. This allows client computers to receive an IP address from multiple logical subnets even when the clients are located on the same physical subnet. You can only create a superscope if you have two or more IP scopes already created in DHCP. You can use the New Superscope Wizard to select the scopes that you wish to combine to create a superscope.

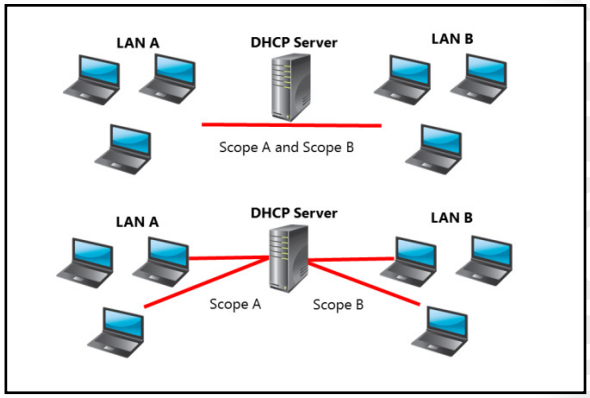

Benefits of Superscopes

A superscope is useful in several situations. For example, if a scope runs out of addresses and you cannot add more addresses from the subnet, you can instead add a new subnet to the DHCP server. This scope will lease addresses to clients in the same physical network, but the clients will be in a separate network logically. This is known as *multinetting*. Once you add a new subnet, you must configure routers to recognize the new subnet so that you ensure local communications in the physical network.

A superscope is also useful when you need to move clients gradually into a new IP numbering scheme. When you have both numbering schemes coexist for the original lease's duration, you can move clients into the new subnet transparently. When you have renewed all client leases in the new subnet, you can retire the old subnet.

Multicast Scopes

A multicast scope is a collection of multicast addresses from the class D IP address range of 224.0.0.0 to 239.255.255.255 (224.0.0.0/3). These addresses are used when applications need to communicate with numerous clients efficiently and simultaneously. This is accomplished with multiple hosts that listen to traffic for the same IP address.

A multicast scope is commonly known as a Multicast Address Dynamic Client Allocation Protocol (MADCAP) scope. Applications that request addresses from these scopes need to support the MADCAP application programming interface (API). Windows Deployment Services is an example of an application that supports multicast transmissions.

Multicast scopes allow applications to reserve a multicast IP address for data and content delivery.

DHCP Integration with IPv6

IPv6 can configure itself without DHCP. IPv6-enabled clients have a self-assigned link-local IPv6 address. A link-local address is intended only for communications within the local network. It is equivalent to the 169.254.0.0 self-assigned addresses used by IPv4. IPv6-enabled network interfaces can, and often do, have more than one IPv6 address. For example, addresses might include a self-assigned link-local address and a DHCP-assigned global address. By using DHCP for IPv6 (DHCPv6), an IPv6 host can obtain subnet prefixes, global addresses, and other IPv6 configuration settings.

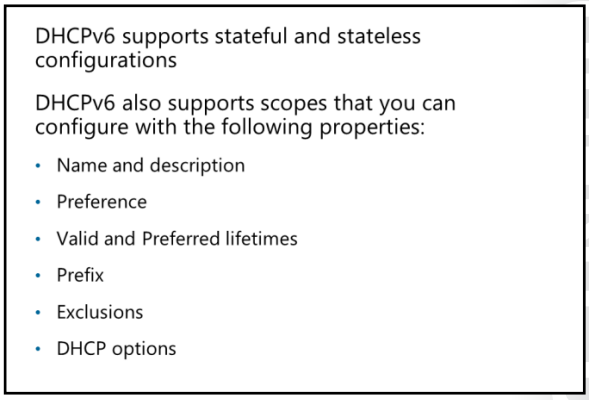

Note: You should obtain a block of IPv6 addresses from a Regional Internet Registry. There are five regional Internet registries in the world. They are:

- African Network Information Centre (AfriNIC) for Africa
- Asia-Pacific Network Information Centre (APNIC) for Asia, Australia, New Zealand, and neighboring countries
- American Registry for Internet Numbers (ARIN) for Canada, many Caribbean and North Atlantic islands, and the United States
- Latin America and Caribbean Network Information Centre (LACNIC) for Latin America and parts of the Caribbean region
- Réseaux IP Européens Network Coordination Centre (RIPE NCC) for Europe, Russia, the Middle East, and Central Asia

Stateful and Stateless Configuration

Whenever you add the DHCP server role to a Windows Server 2012 computer, you also automatically install a DHCPv6 server. Windows Server 2012 supports both DHCPv6 stateful and stateless configurations, described as follows:

- 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 and client agree on an IPv6 automatically, and the DHCPv6 server only assigns other IPv6 configuration settings. The IPv6 address is built by using the network portion from the router, and the host portion of the address, which is generated by the client.

DHCPv6 Scopes for IPv6

DHCPv6 scopes for IPv6 must be created separately from IPv4 scopes. IPv6 scopes have an enhanced lease mechanism and several different options. When configuring a DHCPv6 scope, you must define the properties listed in the following table.

Property	Use
Name and description	This property identifies the scope.
Prefix	The IPv6 address prefix is analogous to the IPv4 address range. It defines the network portion of the IP address.
Preference	This property informs DHCPv6 clients which server to use if you have multiple DHCPv6 servers.
Exclusions	This property defines single addresses or blocks of addresses that fall within the IPv6 prefix but will not be offered for lease.
Valid and Preferred lifetimes	This property defines how long leased addresses are valid.
DHCP options	As with IPv4, there are many available options.

Configuring an IPv6 Scope

You can use the New Scope Wizard to create IPv6 scopes:

1. In the DHCP console, right-click the **IPv6 node**, and then click **New Scope**.
2. Configure a scope prefix and preference.
3. Define the starting and ending IP addresses, and any exclusions.
4. Configure the Preferred and Valid lifetime properties.
5. Activate the scope to enable it.

What Is DHCP Name Protection?

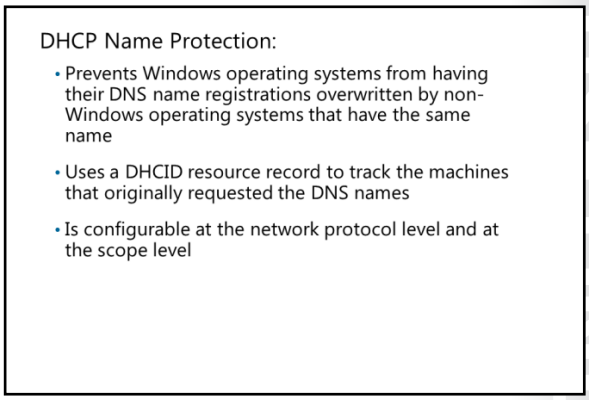

You should protect the names that DHCP registers in DNS on behalf of other computers or systems from being overwritten by other systems that use the same names. In addition, you should also protect the names from being overwritten by systems that use static addresses that conflict with DHCP-assigned addresses when they use unsecure DNS, and when DHCP is not configured for conflict detections. For example, a UNIX-based system named Client1 could potentially overwrite the DNS address that was assigned and registered by DHCP on behalf of a Windows-based system also named Client1. A new feature in Windows Server 2012—DHCP Name Protection—addresses this issue.

Name squatting is the term used to describe the conflict that occurs when one client registers a name with DNS but that name is already used by another client. This problem causes the original machine to become inaccessible, and it typically occurs with systems that have the same names as Windows operating systems. DHCP Name Protection addresses this by using a resource record known as a Dynamic Host Configuration Identifier (DHCID) to track which machines originally requested which names. The DHCP server provides the DHCID record, which is stored in DNS. When the DHCP server receives a request by a machine with an existing name for an IP address, the DHCP server can refer to the DHCID in DNS to verify that the machine that is requesting the name is the original machine that used the name. If it is not the same machine, then the DNS resource record is not updated.

You can implement name protection for both IPv4 and IPv6. In addition, you can configure DHCP Name Protection at both the server level and the scope level. Implementation at the server level will only apply for newly created scopes.

To enable DHCP Name Protection for an IPv4 or IPv6 node:

1. Open the DHCP console.
2. Right-click the **IPv4 or IPv6** node, and then open the **Property** page.
3. Click DNS, click **Advanced**, and then select the **Enable Name Protection** check box.

To enable DHCP Name Protection for a scope:

1. Open the DHCP Microsoft Management Console (MMC).
2. Expand the IPv4 or IPv6 node, right-click the **scope**, and the open the **Property** page.
3. Click DNS, click Advanced, and then select the **Enable Name Protection** check box.

What Is DHCP Failover?

DHCP manages the distribution of IP addresses in TCP/IP networks of all sizes. When this service fails, clients lose connectivity to the network and all of its resources. A new feature in Windows Server 2012, DHCP failover, addresses this issue.

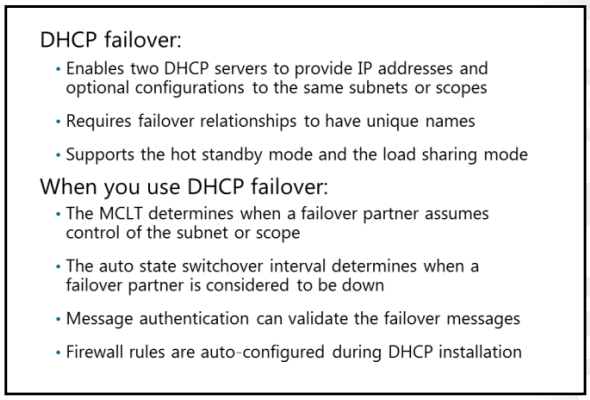

DHCP Failover

DHCP clients renew their leases on their IP addresses at regular, configurable intervals. When the DHCP service fails and the leases time out, the clients no longer have IP addresses. In the past, DHCP failover was not possible because DHCP servers were independent and unaware of each other. Therefore, configuring two separate DHCP servers to distribute the same pool of addresses could lead to duplicate addresses. Additionally, providing redundant DHCP services required that you configure clustering and perform a significant amount of manual configuration and monitoring.

The new DHCP failover feature enables two DHCP servers to provide IP addresses and optional configurations to the same subnets or scopes. Therefore, you can now configure two DHCP servers to replicate lease information. If one of the servers fails, the other server services the clients for the entire subnet.

Note: In Windows Server 2012, you can only configure two DHCP servers for failover, and only for IPv4 scopes and subnets.

Configuring DHCP Failover

To configure DHCP failover, you need to establish a failover relationship between the two DHCP servers' services. You must also give this relationship a unique name. The failover partners exchange this name during configuration. This enables a single DHCP server to have multiple failover relationships with other DHCP servers if all servers have unique names. To configure failover, use the Configuration Failover Wizard that you can launch by right-clicking the IP node or the scope node.

Note: DHCP failover is time sensitive. You must synchronize time between the partners in the relationship. If the time difference is greater than one minute, the failover process will halt with a critical error.

You can configure failover in one of the two following modes.

Mode	Characteristics
Hot standby	In this mode, one server is the primary server and the other is the secondary server. The primary server actively assigns IP configurations for the scope or subnet. The secondary DHCP server only assumes this role if the primary server becomes unavailable. A DHCP server can simultaneously act as the primary for one scope or subnet, and be the secondary for another. Administrators must configure a percentage of the scope addresses to be assigned to the standby server. These addresses are supplied during the Maximum Client Lead Time (MCLT) interval if the primary server is down. The default MCLT value is five percent of the scope. The secondary server takes control of the entire IP range after the MCLT interval has passed. Hot Standby mode is best suited to deployments in which a disaster recovery site is located at a different location. That way, the DHCP server will not service clients unless there is a main server outage.
Load sharing	This is the default mode. In this mode both servers supply IP configuration to clients simultaneously. The server that responds to IP configuration requests depends on how the administrator configures the load distribution ratio. The default ratio is 50:50.

MCLT

The administrator configures the MCLT parameter to determine the amount of time a DHCP server should wait when a partner is unavailable, before assuming control of the address range. This value cannot be zero, and the default is one hour.

Auto State Switchover Interval

A communication-interrupted state occurs when a server loses contact with its partner. Because the server has no way of knowing what is causing the communication loss, it remains in this state until the administrator manually changes it to a partner down state. The administrator can also enable automatic transition to partner down state by configuring the auto state switchover interval. The default value for this interval is 10 minutes.

Message Authentication

Windows Server 2012 enables you to authenticate the failover message traffic between the replication partners. The administrator can establish a shared secret—much like a password—in the Configuration Failover Wizard for DHCP failover. This validates that the failover message comes from the failover partner.

Firewall Considerations

DHCP uses TCP port 647 to listen for failover traffic. The DHCP installation creates the following inbound and outbound firewall rules:

- Microsoft-Windows-DHCP-Failover-TCP-In
- Microsoft-Windows-DHCP-Failover-TCP-Out

Demonstration: Configuring DHCP Failover

In this demonstration, you will see how to configure a DHCP failover relationship.

Demonstration Steps

Configure a DHCP failover relationship

1. Sign in on LON-SVR1 as **Adatum\Administrator** with the password **Pa$$w0rd**. Note that the server is authorized, but that no scopes are configured.
2. Switch to LON-DC1. In Server Manager, click **Tools**, and then on the drop-down list, click **DHCP**.
3. In the DHCP console, launch the **Configure Failover Wizard**.
4. Configure failover replication with the following settings:
   - Partner server: **172.16.0.21**
   - Relationship Name: **Adatum**
   - Maximum Client Lead Time: **15 minutes**
   - Mode: **Load balance**
   - Load Balance Percentage: **50%**
   - State Switchover Interval: **60 minutes**
   - Message authentication shared secret: **Pa$$w0rd**
5. Complete the Configure Failover Wizard.
6. Switch back to LON-SVR1, refresh the IPv4 node and note that the Adatum scope is **configured** and is **active**.

Lesson 2
Configuring Advanced DNS Settings

In TCP/IP networks of any size, certain services are essential. DNS is one of the most critical network services for any network, because many other applications and services—including AD DS—rely on DNS to resolve resource names to IP addresses. Without DNS, user authentications fail, and network-based resources and applications may become inaccessible. For these reasons, you need to manage and protect DNS.

This lesson discusses management techniques and options for optimizing DNS resolution. Windows Server 2012 implements DNSSEC to protect DNS responses. Windows Server 2012 also includes support for global name zones to provide single-label name resolution.

Lesson Objectives

After completing this lesson, you will be able to:

- Manage DNS services.
- Optimize DNS name resolution.
- Describe global name zones.
- Describe options for implementing DNS security.
- Explain how DNSSEC works.
- Describe the new DNSSEC features for Windows Server 2012.
- Explain how to configure DNSSEC.
- Describe global name zones.

Managing DNS Services

Like other important network services, you must manage DNS. DNS management consists of the following tasks:

- Delegating DNS administration
- Configuring logging for DNS
- Aging and scavenging
- Backing up the DNS database

To manage DNS services:
- Delegate DNS administration through membership in the DNS Admins group
- View DNS logs in Event Viewer
- Enable DNS debug logging in the DNS server properties
- Enable aging and scavenging to remove stale records

Backup methods for the DNS database depend on how the database is deployed:
- Back up Active Directory-integrated zones through system state backups, by using dnscmd, or by using Windows PowerShell
- Non-integrated primary zone are single files that you can copy or back up

Delegating Administration of DNS

By default, the Domain Admins group has full permissions to manage all aspects of the DNS server in its home domain, and the Enterprise Admins group has full permissions to manage all aspects of all DNS servers in any domain in the forest. If you need to delegate the administration of a DNS server to a different user or group, then you can add that user or global group to the DNS Admins group for a given domain in the forest. Members of the DNS Admins group can view and modify all DNS data, settings, and configurations of DNS servers in their home domain.

The DNS Admins group is a Domain Local security group, and by default has no members in it.

Configuring DNS Logging

By default, DNS maintains a DNS server log, which you can view in the Event Viewer. This event log is located in the Applications and Services Logs folder in Event Viewer. It records common events such as:

- Starting and stopping the DNS service.
- Background loading and zone signing events.
- Changes to DNS configuration settings.
- Various warnings and error events.

For more verbose logging, you can enable debug logging. Debug logging options are disabled by default, but can be selectively enabled. Debug logging options include the following:

- Direction of packets
- Contents of packets
- Transport protocol
- Type of request
- Filtering based on IP address
- Specifying the name and location of the log file, which is located in the %windir%\System32\DNS directory
- Log file maximum size limit

Debug logging can be resource intensive. It can affect overall server performance and consume disk space. Therefore, you should only enable it temporarily when you require more detailed information about server performance. To enable debug logging on the DNS server, do the following:

1. Open the DNS console.
2. Right-click the applicable DNS server, and then click **Properties**.
3. In the **Properties** dialog box, click the **Debug Logging** tab.
4. Select **Log packets for debugging**, and then select the events for which you want the DNS server to record debug logging.

 **Note:** Logging can generate a large number of files, and if it is left on too long it can fill a drive. We highly recommend that you only turn logging on while you are actively troubleshooting; at all other times, logging should be turned off.

Aging and Scavenging

DNS dynamic updates add resource records to the zone automatically, but in some cases, those records are not deleted automatically when they are no longer required. For example, if a computer registers its own host (A) resource record and is improperly disconnected from the network, the host (A) resource record might not be deleted. These records, known as *stale records*, take up space in the DNS database and may result in an incorrect query response being returned. Windows Server 2012 can search for those stale records and, based on the aging of the record, scavenge them from the DNS database.

Aging and scavenging is disabled by default. You can enable automatic scavenging and the interval at which it will take place in the Advanced properties of the DNS server. Each individual zone is then configured to indicate whether or not the stale records should be scavenged and the aging settings that determine when records become stale. The aging settings are found in the zones properties General tab.

Aging is determined by using parameters known as the No-refresh interval and the Refresh interval. The *No-refresh interval* is the period of time that the record is not eligible to be refreshed. By default, this is seven days. The *Refresh interval* is the date and time that the record is eligible to be refreshed by the client. The default is seven days. In the normal course of events, a client host record cannot be refreshed in the database for seven days after it is first registered or refreshed. However, it then must be refreshed within the next seven days after the No-refresh interval, or the record becomes eligible to be scavenged out of the database. A client will attempt to refresh its DNS record at startup, and every 24 hours while the system is running.

 **Note:** Records that are added dynamically to the database are time stamped. Static records that you enter manually have a time stamp value of zero (0); therefore, they will not be affected by aging, and will not be scavenged out of the database.

Backing Up the DNS Database

How you back up the DNS database depends on how DNS was implemented in your organization. If your DNS zone was implemented as an Active Directory-integrated zone, then your DNS zone is included in the Active Directory database ntds.dit file. If the DNS zone is a primary zone and is not stored in AD DS, then the file is stored as a .dns file in the %SystemRoot%\System32\Dns folder.

Backing Up Active Directory-Integrated Zones

Active Directory-integrated zones are stored in AD DS and are backed up as part of a System State or a full server backup. Additionally, you can back up just the Active Directory–integrated zone by using the dnscmd command-line tool.

To back up an Active Directory–integrated zone, perform the following steps:

1. Launch an elevated command prompt.
2. Run the following command:

   ```
   dnscmd /ZoneExport <zone name> <zone file name>
   ```

   <zone name> is the name of your DNS zone, and *<zone file name>* is the file that you want to create to hold the backup information.

The dnscmd tool exports the zone data to the file name that you designate in the command, to the %windir%\System32\DNS directory.

You can also use Windows PowerShell® to perform the same task. In Windows PowerShell, you use the **Export-DnsServerZone** cmdlet. For example, if you want to export a zone named contoso.com, type the following command:

```
Export-DnsServerZone –Name contoso.com –Filename contoso
```

 **Note:** If DNSSEC is configured, the security information will not be exported with these commands.

Backing Up Primary Zones

To back up a primary zone that is not stored in AD DS, simply copy or back up the individual zone file, *zonename.dns*, which is located in the %windir%\System32\DNS directory. For example, if your DNS primary zone is named Adatum.com, then the DNS zone file will be named Adatum.com.dns.

Optimizing DNS Name Resolution

In a typical DNS query event, a client computer attempts to resolve a FQDN to an IP address. For example, if a user tries to go to the FQDN www.microsoft.com, the client computer will perform a recursive query to the DNS server that it is configured to discover the IP address associated with that FQDN. The local DNS server must then respond with an authoritative response. If the local DNS server has a copy of the DNS zone for which it was queried, it will respond with an authoritative answer to the client computer. If the local DNS server does not have that information, it will perform recursion.

Option	Description
Forwarding	Forwards DNS requests that cannot be resolved locally to other specific DNS servers
Conditional forwarding	Forwards queries for specific DNS suffixes to specific DNS servers
Stub zones	A regularly replicated copy of certain resource records that identify authoritative DNS servers for specific DNS domains
Netmask ordering	Responds with addresses of hosts that are close in proximity based in IP address information of the client to DNS queries

Recursion refers to the process of having the local DNS server itself make a recursive query to another DNS server until it finds the authoritative answer, and then returns that answer to the client that made the original request. By default, this server will be one of the servers on the Internet that is listed as a root hint. When the local DNS server receives a response, it will return that information to the original requesting client computer.

There are a number of options available for optimizing DNS name resolution, which include features such as:

- Forwarding
- Conditional forwarding
- Stub zones
- Netmask ordering

Forwarding

A *forwarder* is a network DNS server that you configure to forward DNS queries for host names that it cannot resolve to other DNS servers for resolution. In a typical environment, the internal DNS server forwards queries for external DNS host names to DNS servers on the Internet. For example, if the local network DNS server cannot resolve a query for www.microsoft.com, then the local DNS server can forward the query to the Internet service provider's (ISP's) DNS server for resolution.

Conditional Forwarding

You also can use conditional forwarders to forward queries according to specific domain names. A conditional forwarder is a setting that you configure on a DNS server that enables forwarding DNS queries based on 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.adatum.com to the IP address of a specific DNS server, or to the IP addresses of multiple DNS servers. This can be useful when you have multiple DNS namespaces in a forest, or a partner's DNS namespace across firewalls. For example, suppose Contoso.com and Adatum.com are merged. Rather than requiring each domain to host a complete replica of the other domain's DNS database, you could create conditional forwarders so that they point to each other's specific DNS servers for resolution of internal DNS names.

Stub Zones

A *stub zone* is a copy of a zone that contains only those resource records necessary to identify that zone's DNS servers. A stub zone resolves names between separate DNS namespaces, which might be necessary when you want a DNS server that is hosting a parent zone to remain aware of all the DNS servers for one of its child zones. A stub zone that is hosted on a parent domain DNS server will receive a list of all new

DNS servers for the child zone, when it requests an update from the stub zone's master server. By using this method, the DNS server that is hosting the parent zone maintains a current list of the DNS servers for the child zone as they are added and removed.

A stub zone consists of the following:

- The delegated zone's start of authority (SOA) resource record, name server (NS) resource records, and host (A) resource records
- The IP address of one or more master servers that you can use to update the stub zone

Stub zones have the following characteristics:

- You create stub zones using the New Zone Wizard.
- You can store stub zones in AD DS.
- You can replicate stub zones either in the domain only, or throughout the entire forest or any other replication scope configured by Active Directory application partitions.
- Stub zone master servers are one or more DNS servers that are responsible for the initial copy of the zone information, and are usually the DNS server that is hosting the primary zone for the delegated domain name.

Conditional Forwarding vs. Stub Zones

Conditional forwarder and stub zones perform similar functions. The distinguishing difference between conditional forwarders and stub zones are that conditional forwarders work better across firewalls, while stub zones are more dynamic when DNS-servers are added and removed. If you have firewalls, you usually configure two DNS servers that can be accessed by a partner behind the firewall; therefore, you need to configure conditional forwarding. For internal DNS servers, where you usually do not have firewalls or permit DNS traffic to all DNS servers behind the firewall, you can use stub zones that automatically learn about new DNS servers.

Netmask Ordering

There are various reasons for having multiple IP addresses associated with a single name, for example, load balancing a web page. Netmask ordering returns addresses for type A address records (A record) DNS queries that prioritize resources on the client computer's local subnet to the client. In other words, addresses of hosts that are on the same subnet as the requesting client will have a higher priority in the DNS response to the client computer.

Localization is based on IP addresses. For example, if there are multiple A records that are associated with the same DNS name, and each of the A records are located on a different IP subnet, netmask ordering returns an A record that is on the same IP subnet as the client computer that made the request.

What Is the GlobalNames Zone?

The GlobalNames zone was introduced with Windows 2008, and support for this zone continues in Windows Server 2012. The GlobalNames zone contains 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. GlobalNames zones provide single-label name resolution for large enterprise networks that do not deploy WINS and that have multiple DNS domain environments. GlobalNames zones are created manually and do not support dynamic record registration.

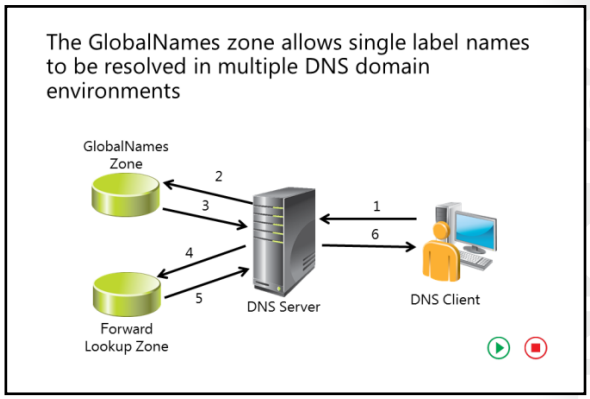

When clients try to resolve short names, they append their DNS domain name automatically. Depending on the configuration, they also try to find the name in upper-level domain name, or work through their name suffix list. Therefore, short names are primarily resolved in the same domain.

You use a GlobalNames zone to provide a short name to multiple DNS suffixes. For example, if an organization supports two DNS domains, such as adatum.com and contoso.com, and they have a server called intranet in contoso.com, only contoso domain users would be able to query it using the short name. Users of the adatum domain would not be able to use the short name to access the server.

Global names are based on creating alias (CNAME) resource records in a special forward lookup zone that uses single names to point to FQDNs. For example, GlobalNames zones would enable clients in both the adatum.com domain and the contoso.com domain to use a single label name, such as intranet, to locate a server whose FQDN is intranet.contoso.com without having to use the FQDN.

Creating a GlobalNames Zone

To create a GlobalNames zone, do the following:

1. Use the **dnscmd** tool to enable GlobalNames zones support.
2. Create a new forward lookup zone named GlobalNames (not case sensitive). Do not allow dynamic updates for this zone.
3. Manually create CNAME records that point to records that already exist in the other zones that are hosted on your DNS servers.

For example, you could create a CNAME record in the GlobalNames zone named Data that points to Data.contoso.com. This enables clients from any DNS domain in the organization to find this server by the single label name of Data.

You can also use the Windows PowerShell cmdlets **Get-DnsServerGlobalNameZone** and **Set-DnsServerGlobalNameZone** to configure GlobalNames zones.

Options for Implementing DNS Security

Because DNS is a critical network service, you must protect it as much as possible. A number of options are available for protecting the DNS server, including:

- DNS cache locking
- DNS socket pool
- DNSSEC

Option	Description
DNS cache locking	Prevents entries in the cache from being overwritten until a percentage of the TTL has expired
DNS socket pool	Randomizes the source port for issuing DNS queries Enabled by default in Windows Server 2012
DNSSEC	Enables cryptographically signing DNS records so that client computers can validate responses

DNS Cache Locking

Cache locking is a Windows Server 2012 security feature that allows you to control when information in the DNS cache can be overwritten. When a recursive DNS server responds to a query, it caches the results so that it can respond quickly if it receives another query requesting the same information. The period of time the DNS server keeps information in its cache is determined by the Time to Live (TTL) value for a resource record.

Information in the cache can be overwritten before the TTL expires if updated information about that resource record is received. If a malicious user successfully overwrites information in the cache, then the malicious user might be able to redirect your network traffic to a malicious site. When you enable cache locking, the DNS server prohibits cached records from being overwritten for the duration of the TTL value.

You configure cache locking as a percentage value. For example, if the cache locking value is set to 50, then the DNS server will not overwrite a cached entry for half of the duration of the TTL. By default, the cache locking percentage value is 100. This means that cached entries will not be overwritten for the entire duration of the TTL.

You can configure cache locking with the **dnscmd** tool as follows:

1. Launch an elevated command prompt.
2. Run the following command:

```
dnscmd /Config /CacheLockingPercent <percent>
```

3. Restart the DNS service to apply the changes.

   Alternatively, you can use the Windows PowerShell **Set-DnsServerCache –LockingPercent** cmdlet to set this value. For example:

```
Set-DnsServerCache –LockingPercent <value>
```

DNS Socket Pool

The DNS socket pool enables a DNS server to use source port randomization when issuing DNS queries. When the DNS service starts, the server chooses a source port from a pool of sockets that are available for issuing queries. Instead of using a predicable source port, the DNS server uses a random port number that it selects from the DNS socket pool. The DNS socket pool makes cache-tampering attacks more difficult because a malicious user must correctly guess both the source port of a DNS query and a random transaction ID to successfully run the attack. The DNS socket pool is enabled by default in Windows Server 2012.

The default size of the DNS socket pool is 2,500. When you configure the DNS socket pool, you can choose a size value from 0 to 10,000. The larger the value, the greater the protection you will have against

DNS spoofing attacks. If the DNS server is running Windows Server 2012, you can also configure a DNS socket pool exclusion list.

You can configure the DNS socket pool size by using the **dnscmd** tool as follows:

1. Launch an elevated command prompt.
2. Run the following command:

```
dnscmd /Config /SocketPoolSize <value>
```

3. Restart the DNS service to apply the changes.

DNSSEC

DNSSEC enables a DNS zone and all records in the zone to be signed cryptographically so that client computers can validate the DNS response. DNS is often subject to various attacks, such as spoofing and cache-tampering. DNSSEC helps protect against these threats and provides a more secure DNS infrastructure.

How DNSSEC Works

Intercepting and tampering with an organization's DNS query response is a common attack method. If malicious users can alter responses from DNS servers, or send spoofed responses to point client computers to their own servers, they can gain access to sensitive information. Any service that relies on DNS for the initial connection—such as e-commerce web servers and email servers—are vulnerable. DNSSEC protects clients that are making DNS queries from accepting false DNS responses.

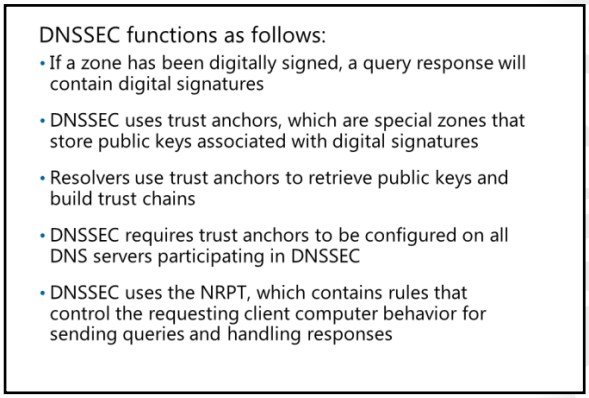

When a DNS server that is hosting a digitally signed zone receives a query, it returns the digital signatures along with the requested records. A resolver or another server can obtain the public key of the public/private key pair from a trust anchor, and then validate that the responses are authentic and have not been tampered with. To do this, the resolver or server must be configured with a trust anchor for the signed zone or for a parent of the signed zone.

Trust Anchors

A *trust anchor* is an authoritative entity that is represented by a public key. The TrustAnchors zone stores preconfigured public keys that are associated with a specific zone. In DNS, the trust anchor is the DNSKEY or DS resource record. Client computers use these records to build trust chains. You must configure a trust anchor from the zone on every domain DNS server to validate responses from that signed zone. If the DNS server is a domain controller, then Active Directory-integrated zones can distribute the trust anchors.

Name Resolution Policy Table

The Name Resolution Policy Table (NRPT) contains rules that control the DNS client behavior for sending DNS queries and processing the responses from those queries. For example, a DNSSEC rule prompts the client computer to check for validation of the response for a particular DNS domain suffix. As a best practice, you should use Group Policy as the preferred method for configuring the NRPT. If there is no NRPT present, the client computer accepts responses without validating them.

Deploying DNSSEC

To deploy DNSSEC:

1. Install Windows Server 2012, and assign the DNS role to the server. Typically, a domain controller also acts as the DNS server. However, this is not a requirement.
2. Sign the DNS zone by using the DNSSEC Configuration Wizard, which is located in the DNS console.
3. Configure trust anchor distribution points.
4. Configure the NRPT on the client computers.

Assigning the DNS Server Role

To assign the DNS server role, in the Server Manager Dashboard, use the Add Roles and Features Wizard. You can also add this role can when you add the AD DS role. Then configure the primary zones on the DNS server. After a zone is signed, any new DNS servers in Windows Server 2012 automatically receive the DNSSEC parameters.

Signing the Zone

The following signing options are available:

- **Configure the zone signing parameters**. This option guides you through the steps and enables you to set all values for the key signing key (KSK) and the zone signing key (ZSK).
- **Sign the zone with parameters of an existing zone**. This option enables you to keep the same values and options that are set in another signed zone.
- **Use recommended settings**. This option signs the zone by using the default values.

 **Note:** Zones can also be unsigned by using the DNSSEC management user interface to remove zone signatures.

Configuring Trust Anchor Distribution Points

If the zone is Active Directory-integrated, and if all domain controllers are running Windows Server 2012, you can select to distribute the trust anchors to all the servers in the forest. Make this selection with caution because the wizard turns on DNSSEC validation. If you enable DNS trust anchors without thorough testing, you could cause DNS outages. If trust anchors are required on computers that are not domain joined—for example, a DNS server in the perimeter network (also known as screened subnet)—then you should enable automated key rollover.

 **Note:** A key rollover is the act of replacing one key pair with another at the end of a key's effective period.

Configuring NRPT on Client Computers

The DNS client computer only performs DNSSEC validation on domain names where the NRPT has configured the DNS client computer to do so. A client computer that is running Windows 7 is DNSSEC-aware, but it does not perform validation. Instead, it relies on the security-aware DNS server to perform validation on its behalf.

New DNSSEC Features for Windows Server 2012

Windows Server 2012 has simplified DNSSEC implementation. Although DNSSEC was supported in Windows Server 2008 R2, most of the configuration and administration tasks were performed manually, and zones were signed when they were offline.

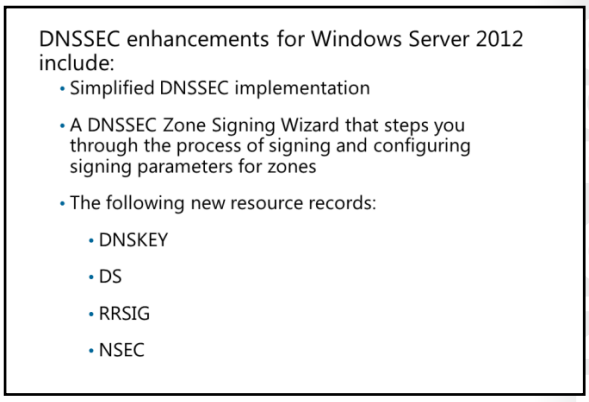

DNSSEC Zone Signing Wizard

Windows Server 2012 includes a DNSSEC Zone Signing Wizard to simplify the configuration and signing process, and to enable online signing. The wizard allows you to choose the zone signing parameters as indicated in the previous topic. If you choose to configure the zone-signing settings rather than using parameters from an existing zone or using default values, you can use the wizard to configure settings such as:

- Key signing key (KSK) options
- Zone signing key (ZSK) options
- Trust anchor distribution options
- Signing and polling parameters

New Resource Records

DNS response validation is achieved by associating a private/public key pair (as generated by the administrator) with a DNS zone, and then defining additional DNS resource records to sign and publish keys. Resource records distribute the public key, while the private key remains on the server. When the client requests validation, DNSSEC adds data to the response that enables the client to authenticate the response.

The following table describes the new resource records in Windows Server 2012.

Resource record	Purpose
DNSKEY	This record publishes the public key for the zone. It checks the authority of a response against the private key held by the DNS server. These keys require periodic replacement through key rollovers. Windows Server 2012 supports automated key rollovers. Every zone has multiple DNSKEYs that are then broken down to the ZSK and KSK.
Delegation Signer (DS)	This record is a delegation record that contains the hash of the public key of a child zone. This record is signed by the parent zone's private key. If a child zone of a signed parent is also signed, the DS records from the child must be manually added to the parent so that a chain of trust can be created.
Resource Record Signature (RRSIG)	This record holds a signature for a set of DNS records. It is used to check the authority of a response.
Next Secure (NSEC)	When the DNS response has no data to provide to the client, this record authenticates that the host does not exist.
NSEC3	This record is a hashed version of the NSEC record thatprevents alphabet attacks by enumerating the zone.

Other New Enhancements

Other enhancements for Windows Server 2012 include:

- Support for DNS dynamic updates in DNSSEC signed zones.
- Automated trust anchor distribution through AD DS.
- Windows PowerShell-based command-line interface for management and scripting.

Managing DNSSEC with Windows PowerShell cmdlets

Windows Server 2012 R2 added several Windows PowerShell cmdlets to manage DNSSEC, including:

cmdlet	Description
Add-DnsServerResourceRecordDnsKey	You use this cmdlet to add a type DNSKEY resource record to a DNS zone.
Add-DnsServerResourceRecordDS	You use this cmdlet to add a type DS resource record to a DNS zone.
Add-DnsServerTrustAnchor	You use this cmdlet to add a trust anchor to a DNS server. Windows Server 2012 R2 now includes the **–Root** option. This option allows you to retrieve trust anchors from the specified URL.
Add-DnsServerSigningKey	You use this cmdlet to add a KSK or ZSK to a signed zone.
Export-DnsServerDnsSecPublicKey	You use this cmdlet to export DS and DNSKEY information for a DNSSEC-signed zone.
Get-DnsServerDnsSecZoneSetting	You use this cmdlet to get the DNSSEC settings for a zone.
Get-DnsServerSetting	You use this cmdlet to retrieve DNS server settings. Windows Server 2012 R2 adds the **RootTrustAnchorsURL** to the output.
Set-DnsServerDnsSecZoneSetting	You use this cmdlet to make changes to the settings for a DNSSEC zone.
Step-DnsServerSigningKeyRollover	You use this cmdlet to force a KSK rollover when the DS record has been manually updated in the parent.

Demonstration: Configuring DNSSEC

In this demonstration, you will see how to use the Zone Signing Wizard in the DNS console to configure DNSSEC.

Demonstration Steps

Configure DNSSEC

1. Sign in on LON-DC1 as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. Start the DNS console.
3. Use the DNSSEC Zone Signing Wizard to sign the Adatum.com zone.

4. Choose to customize zone signing parameters.
5. Ensure that DNS server **LON-DC1 is the Key Master**.
6. Add the **Key Signing Key** by accepting default values for the new key.
7. Add the **Zone Signing Key** by accepting the default values for the new key.
8. Choose to use **NSCE3** with default values.
9. Do not choose to enable the distribution of trust anchors for this zone.
10. Accept the default values for **signing** and **polling**.
11. Verify that the DNSKEY resource records were created in the **Trust Points zone**.
12. Use the Group Policy Management Console (GPMC) to configure NRPT. Create a rule that enables DNSSEC for the Adatum.com suffix, and that requires DNS client computers to verify that the name and address data is validated.

Lesson 3
Implementing IPAM

With the development of IPv6 and with more devices requiring IP addresses, networks have become complex and difficult to manage. Maintaining an updated list of static IP addresses that have been issued has often been a manual task, which can lead to errors. To help organizations manage IP addresses, Windows Server 2012 provides the IP Address Management (IPAM) tool.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe IPAM.
- Describe IPAM architecture.
- Describe the requirements for IPAM implementations.
- Explain how to manage IP addressing using IPAM.
- Explain how to install and configure IPAM.
- Explain how to manage and monitor IPAM.
- Describe considerations for implementing IPAM.

What Is IPAM?

IP address management is a difficult task in large networks, because tracking IP address usage is largely a manual operation. Windows Server 2012 introduces IPAM, which is a framework for discovering, auditing, monitoring utilization, and managing the IP address space in a network. IPAM enables the administration and monitoring of DHCP and DNS, and provides a comprehensive view of where IP addresses are used. IPAM collects information from domain controllers and Network Policy Servers (NPSs), and then stores that information in the Windows Internal Database.

IPAM facilitates IP management in organizations with complex networks by enabling administration and monitoring of DHCP and DNS

IP administration area	Description
Planning	Reduces the time and expense of the planning process when changes occur in the network
Managing	Provides a single point of management and assists in optimizing utilization and capacity planning for DHCP and DNS
Tracking	Enables tracking and forecasting of IP address utilization
Auditing	Assists with compliance requirements and provides reporting for forensics and change management

IPAM assists in the areas of IP administration, as shown in the following table.

IP administration area	IPAM capabilities
Planning	Provides a tool set that can reduce the time and expense of the planning process when changes occur in the network.
Managing	Provides a single point of management, and assists in optimizing utilization and capacity planning for DHCP and DNS.
Tracking	Enables tracking and forecasting of IP address utilization.
Auditing	Assists with compliance requirements, such as Health Insurance Portability and Accountability Act (HIPAA) and Sarbanes-Oxley Act of 2002, and provides reporting for forensics and change management.

Characteristics of IPAM

Characteristics of IPAM include:

- A single IPAM server can support up to 150 DHCP servers and 500 DNS servers.
- A single IPAM server can support up to 6,000 DHCP scopes and 150 DNS zones.
- IPAM stores three years of forensics data (IP address leases, host MAC addresses, user logon and logoff information) for 100,000 users in a Windows Internal Database when using Windows Server 2012, Windows Server 2012 R2 added the option to select a Windows Internal Database or SQL Server. There is no database purge policy provided, and the administrator must purge the data manually as needed.
- IPAM on Windows Server 2012 supports only Windows Internal Database. An external database is only supported when IPAM is implemented on Windows Server 2012 R2.
- IP address utilization trends are provided only for IPv4.
- IP address reclamation support is provided only for IPv4.
- IPAM does not check for IP address consistency with routers and switches.

Benefits of IPAM

IPAM benefits include:

- IPv4 and IPv6 address space planning and allocation
- IP address space utilization statistics and trend monitoring
- Static IP inventory management, lifetime management, and DHCP and DNS record creation and deletion
- Service and zone monitoring of DNS services
- IP address lease and logon event tracking
- Role-based access control (RBAC)
- Remote administration support through RSAT
- Reporting in the IPAM management console

 **Note:** IPAM has limited support for management and configuration of non-Microsoft network elements.

Windows Server 2012 R2 Enhancements to IPAM

Windows Server 2012 R2 improves and adds functionality for IPAM. The IPAM framework is expanded with the following:

- RBAC. RBAC for IPAM allows you to customize roles, access scopes, and access policies for IPAM administrators.
- Virtual address space management. You can use IPAM provide to manage IP addresses in a Microsoft-based network. You can manage both physical and virtual addresses. Integration between IPAM and System Center 2012 R2 Virtual Machine Managers (VMMs) allows end-to-end address space management. You can view virtual address space in the new VIRTUALIZED ADDRESS SPACE node of the IPAM console.
- Enhanced DHCP server management. DHCP management is improved in Windows Server 2012 R2 to include new DHCP scope and DHCP server operations. Additionally, views were added for DHCP failover, DHCP policies, DHCP superscopes, DHCP filters, and DHCP reservations.

- External database support. You can configure IPAM to use a Windows Internal Database (WID). Support for using Microsoft SQL Server was added in Windows Server 2012 R2.
- Upgrade and migration support. You can upgrade the IPAM database from Windows Server 2012 to Windows Server 2012 R2.
- Enhanced Windows PowerShell support. IPAM includes more than 50 different Windows PowerShell commands.

For a complete list of the available commands, review IPAM Server cmdlets in Windows PowerShell.

http://go.microsoft.com/fwlink/?LinkID=386637

IPAM Architecture

IPAM architecture consists of four main modules, as listed in the following table.

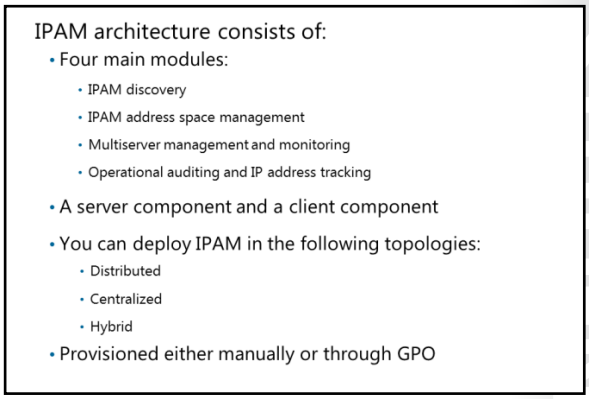

Module	Description
IPAM discovery	You use AD DS to discover servers that are running Windows Server 2008 and newer Windows Server operating systems, and that have DNS, DHCP, or AD DS installed. You can define the scope of discovery to a subset of domains in the forest. You can also add servers manually.
IP address space management	You can use this module to view, monitor, and manage the IP address space. You can dynamically issue or statically assign addresses. You can also track address utilization and detect overlapping DHCP scopes.
Multi-server management and monitoring	You can manage and monitor multiple DHCP servers. This enables tasks to execute across multiple servers. For example, you can configure and edit DHCP properties and scopes, and track the status of DHCP and scope utilization. You can also monitor multiple DNS servers, and monitor the health and status of DNS zones across authoritative DNS servers.
Operational auditing and IP address tracking	You can use the auditing tools to track potential configuration problems. You can also collect, manage, and view details of configuration changes from managed DHCP servers. You can also collect address lease tracking from DHCP lease logs, and collect logon event information from NPS and domain controllers.

The IPAM server can only manage one Active Directory forest. As such, you can deploy IPAM in one of three topologies:

- Distributed. You deploy an IPAM server to every site in the forest.
- Centralized. You deploy only one IPAM server in the forest.
- Hybrid. You deploy a central IPAM server together with a dedicated IPAM server in each site.

 **Note:** IPAM servers do not communicate with one another or share database information. If you deploy multiple IPAM servers, you must customize each server's discovery scope.

IPAM has two main components:

- IPAM server. The IPAM server performs the data collection from the managed servers. It also manages the Windows Internal Database and provides RBAC.
- IPAM client. The IPAM client provides the client computer user interface. It also interacts with the IPAM server, and invokes Windows PowerShell to perform DHCP configuration tasks, DNS monitoring, and remote management.

Provisioning for IPAM

After you install an IPAM server, servers that are managed by IPAM need to be provisioned to allow remote management. You can either manage it manually or through a GPO. If you decide to manually provision the managed servers, you will need to create all the required network shares, security groups, and firewall rules on each managed server.

If you decide to manually provision for IPAM, you must first create a group in AD DS named IPAMUG. This group contains the IPAM servers in the domain. The following table summarizes the required configuration settings that would need to be manually configured.

Configuration Setting	Domain Controller Servers and NPS	DHCP Servers	DNS Servers
<domain>\IPAMUG group	Added as a member of the **BUILTIN\Event Log Readers** group	Added as a member of the **BUILTIN\Event Log Readers** group and the **BUILTIN\DHCP Users** group	Added as a member of the **BUILTIN\Event Log Readers** group
Windows Firewall with Advanced Security	Inbound firewall rules to allow **Remote Event Log Management**	Inbound firewall rules to allow **DHCP Server Management**, **Remote Service Management**, **File and Printer Sharing** and **Remote Event Log Management**	Inbound firewall rules to allow **DNS Service**, **Remote Service Management**, and **Remote Event Log Management**
Network Share		Share the **%SYSTEMROOT%\System32\DHCP** folder as DHCPAudit. Grant IPAMUG read permissions	

Configuration Setting	Domain Controller Servers and NPS	DHCP Servers	DNS Servers
Event Log Monitoring on DNS servers			Modify the **HKLM\SYSTEM\CurrentControlSet\Services\EventLog\DNS Server** registry key
Additional settings			Add <domain>\IPAMUG group as DNS Administrator

If you choose to use GPO provisioning, you will run the **Invoke-IpamGpoProvisioning** Windows PowerShell command. Running this command will create three GPOs to configure the settings described in the table above.

- IPAM_DC_NPS. This GPO is applied to all managed AD DS servers and NPS servers.
- IPAM_DHCP. This GPO is applied to all managed DHCP servers. This GPO includes scripts to configure the network share for DHCP monitoring.
- IPAM_DNS. This GPO is applied to all managed DNS servers. This GPO includes scripts to configure the event log for DNS monitoring and to configure the IPAMUG group as a DNS administrator.

Scenarios for Using IPAM

The general scenario for using IPAM on Windows Server 2012 R2 is supporting network automation in a virtualized datacenter, such as a cloud environment provided by a third-party company or enterprise. There are many company scenarios where IPAM would be a viable solution based on the general scenarios supported by Windows Server 2012 R2 IPAM. Discuss how you envision using the following IPAM features in your environment with the class.

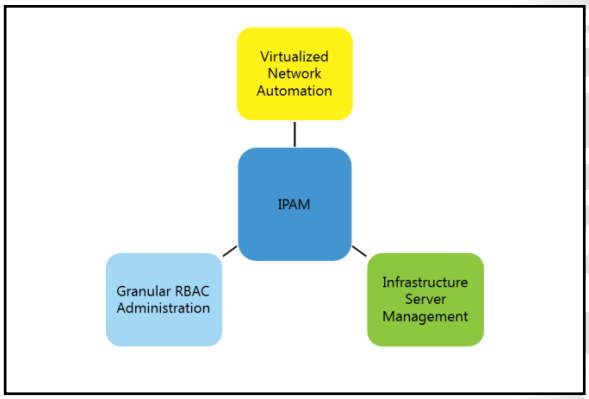

Virtualized Network Automation

IPAM provides unified administration of physical and virtual IP address spaces. When IPAM is integrated with VMM you can manage the IP addresses for your hybrid cloud solution from a single console.

Granular RBAC Administration

You can use RBAC to ensure that administrators can only manage the specified areas in larger environments that may require multiple administrators to manage the IP address spaces.

Infrastructure Administration

You can use IPAM to configure and manage the DNS and DHCP servers in your environment.

Requirements for IPAM Implementation

To ensure a successful IPAM implementation, you must meet the following prerequisites:

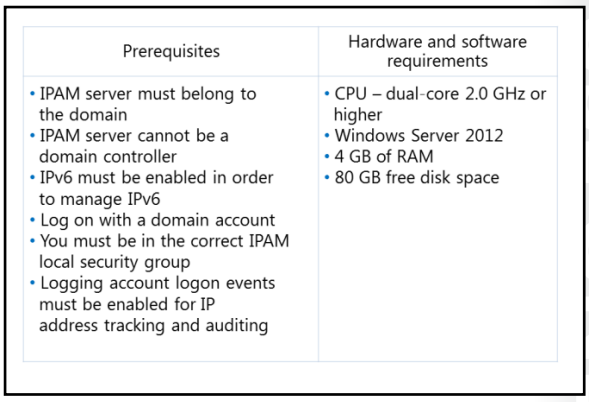

Prerequisites	Hardware and software requirements
• IPAM server must belong to the domain • IPAM server cannot be a domain controller • IPv6 must be enabled in order to manage IPv6 • Log on with a domain account • You must be in the correct IPAM local security group • Logging account logon events must be enabled for IP address tracking and auditing	• CPU – dual-core 2.0 GHz or higher • Windows Server 2012 • 4 GB of RAM • 80 GB free disk space

- The IPAM server must be a domain member, but cannot be a domain controller.
- The IPAM server should be a single-purpose server. Do not install other network roles such as DHCP or DNS on the same server.
- To manage the IPv6 address space, you must have IPv6 enabled on the IPAM server.
- Sign in on the IPAM server with a domain account, and not with a local account.
- You must be a member of the correct IPAM local security group on the IPAM server.
- You must enable logging of account logon events on domain controller and NPS servers for IPAM's IP address tracking and auditing feature.

IPAM Hardware and Software Requirements

The IPAM hardware and software requirements are as follows:

- Dual-core processor of 2.0 gigahertz (GHz) or higher
- Windows Server 2012 operating system
- 4 or more gigabytes (GB) of random access memory (RAM)
- 80 GB of free hard disk space

In addition to the previously mentioned requirements, Windows Server 2008 and Windows Server 2008 R2 require the following:

- Service Pack 2 (SP2) must be installed on Windows Server 2008.
- Microsoft® .NET Framework 4.0 full installation must be installed.
- Windows Management Framework 3.0 must be installed (KB2506146).
- For Windows Server 2008 SP2, Windows Management Framework Core (KB968930) is also required.
- Windows Remote Management must be enabled.
- Verify that service principal names (SPNs) are written.

Demonstration: Implementing IPAM

In this demonstration, you will see how to install and configure IPAM management.

Demonstration Steps

Install IPAM

1. Sign in on LON-SVR2 as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. In the Server Manager, add the IPAM feature and all required supporting features.

Configure IPAM

1. In the IPAM Overview pane, provision the IPAM server using Group Policy.
2. Enter **IPAM** as the Group Policy Object (GPO) name prefix, and provision IPAM. Provisioning will take a few moments to complete.
3. In the IPAM Overview pane, configure server discovery for the Adatum domain.
4. In the IPAM Overview pane, start the server discovery process. Discovery may take 5 to10 minutes to run. The yellow bar indicates when discovery is complete.
5. In the IPAM Overview pane, add the servers to be managed.
6. Verify that IPAM access is currently blocked.
7. Use Windows PowerShell to grant the IPAM server permission to manage LON-DC1 by using the following command:

```
Invoke-IpamGpoProvisioning –Domain Adatum.com –GpoPrefixName IPAM –IpamServerFqdn
LON-SVR2.adatum.com –DelegatedGpoUser Administrator
```

8. Set the manageability status to **Managed**.
9. Switch to LON-DC1.
10. Force the update of Group Policy.
11. Switch back to LON-SVR2 and refresh the IPv4 view. Discovery may take 5 to 10 minutes to run.
12. In the IPAM Overview pane, retrieve data from the managed server.

Virtual Address Space Management in IPAM

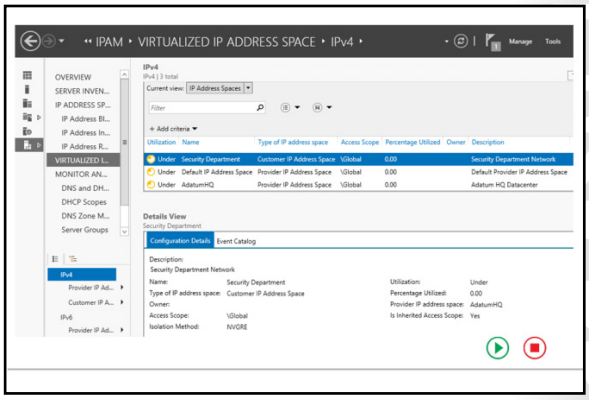

Beginning with Windows Server 2012 R2, IPAM offers a centralized management console for both physical and virtual address spaces. When IPAM is integrated with virtual machine monitor (VMM), you can use automation for your Microsoft cloud-based network. You can use IPAM to manage multiple instances of VMM to provide a single console for detecting conflicts, duplicates, and overlaps of the IP Address spaces in your data center.

Virtualization support is provided when you use the two types of virtual address spaces in IPAM, which are called the provider and the customer virtual address spaces. The provider address space typically contains the addresses associated with the datacenter, the customer address spaces typically hold the virtual addresses used by the customers. The only address space created during installation is the Default IP Address Space; the Default IP Address Space is a provider address space located in the VIRTUALIZED IP ADDRESS SPACE pane.

To create a new Address space you use the **Add-IpamAddressSpace** Windows PowerShell cmdlet. When you create a virtual address space, you must specify a friendly name for the address space, regardless of whether it is a provider or a customer address space. Additionally, you can add an optional description. When you create a customer address space, you must also specify the provider address space in which the customer address space resides, and the isolation method used for the customer network.

To create a new provider address space for the AdatumHQ datacenter based virtual systems, use the following Windows PowerShell cmdlet.

```
Add-IpamAddressSpace –Name “AdatumHQ” –ProviderAddressSpace –Description “Adatum HQ Datacenter”
```

When you create a customer address space, you must configure additional settings. A customer address space must reside in a provider address space. Additionally, you must specify how the customer network will interact with other networks when you specify the network isolation method as either IPRewrite or NVGRE. IPRewrite is a static isolation method in which each customer IP address gets rewritten when you use a physical address from the provider network. Network Virtualization using Generic Routing Encapsulation (NVGRE) is an isolation method that encapsulates the customer traffic and sends all the customer traffic using a single IP address from the provider network.

To create a new customer address space for the Security department, using the AdatumHQ provider address space and NVGRE isolation, use the following Windows PowerShell cmdlet.

```
Add-IpamAddressSpace -Name "Security Department" -CustomerAddressSpace -
AssociatedProviderAddressSpace "AdatumHQ" -IsolationMethod NVGRE –Description “Security Department Network”
```

You can create additional optional settings as part of the Windows PowerShell command or manually add them after creation. These optional settings include custom fields such as AD site or VMM IP Pool Name.

IPAM RBAC

Windows Server 2012 R2 includes RBAC for IPAM. RBAC allows you to customize how administrative permissions are defined in IPAM. For example, some people are assigned the role of administrator and are able to manage all aspects of IPAM, while other administrators may only be allowed to manage certain network objects. By default, all objects inherit the scope of their parent object. To change the Access Scope of an object, right-click the object and click on **Set Access Scope**.

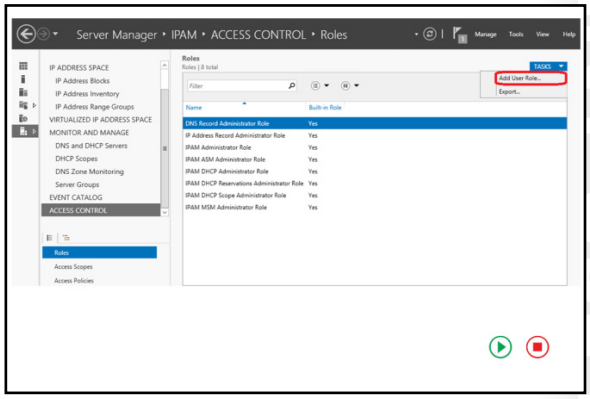

RBAC security is divided into the following three aspects, roles, access scopes, and access policies:

- Roles. A role is a collection of IPAM operations. The roles define the actions an administrator is allowed to perform. Roles are associated with Windows groups and/or users through the use of access policies. There are eight built-in RBAC roles for IPAM. New roles are created and added in the IPAM console, in the ACCESS CONTROL pane.
- The built-in roles for IPAM are:

Name	Description
DNS record administrator	Manages DNS resource records
IP address record administrator	Manages IP addresses but not IP address spaces, ranges, blocks, or subnets.

Name	Description
IPAM administrator	Manages all settings and objects in IPAM
IPAM ASM administrator	Completely manages IP addresses
IPAM DHCP administrator	Completely manages DHCP servers
IPAM DHCP reservations administrator	Manages DHCP reservations
IPAM DHCP scope administrator	Manages DHCP scopes
IPAM MSM administrator	Completely manages DHCP and DNS servers

- Access scopes. Access scopes define the objects to which an administrator has access. By default, the Global access scope is created when IPAM is installed, and all administrator-created access scopes are sub-scopes of the Global access scope. Users or groups assigned to the Global access scope can manage all the network objects in IPAM. Access scopes have up to 15 major operations that can be assigned, such as DHCP server operations. These are further defined by multiple related operations, such as Create DHCP scope, that can be assigned individually. This allows for a large customization range for administrative permissions in IPAM. New access scopes are created and added in the IPAM console, in the ACCESS CONTROL pane.
- Access Policies. An access policy combines a role with an access scope to assign RBAC permissions within IPAM. New access policies are created and added in the IPAM console, in the ACCESS CONTROL pane.

Lesson 4
Managing IP Address Spaces with IPAM

There are multiple phases to managing IP addresses with IPAM. IPAM can automatically manage IP addresses issued through DHCP servers, or you can manually create IP address ranges for management. In this lesson, you will learn how to manage all aspects of IPAM, from configuring automatic management to manually adding and updating address information. Finally, you will learn how to monitor your IP address usage.

Lesson Objectives

After completing this lesson, you will be able to:

- Use IPAM to manage IP addressing.
- Add address spaces to IPAM.
- Import and Update address spaces.
- Maintain an IPAM inventory.
- Monitor IPAM.

Using IPAM to Manage IP Addressing

IP address space management allows administrators to manage, track, audit, and report on an organization's IPv4 and IPv6 address spaces. The IPAM IP address space console provides administrators with IP address utilization statistics and historical trend data so that they can make informed planning decisions for dynamic, static, and virtual address spaces. IPAM periodic tasks discover the address space and utilization data automatically, as configured on the DHCP servers that are managed in IPAM. You can also import IP address information from comma separated values (.csv) files.

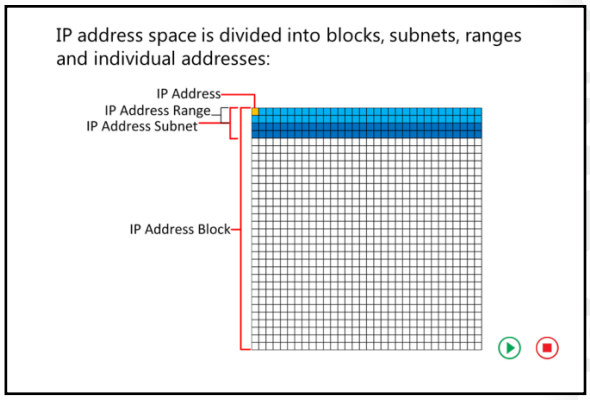

IPAM also enables administrators to detect overlapping IP address ranges that are defined on different DHCP servers, find free IP addresses within a range, create DHCP reservations, and create DNS records.

IPAM provides a number of ways to filter the view of the IP address space. You can customize how you view and manage the IP address space by using any of the following views:

- IP address blocks, which contain:
  - IP address subnets
  - IP address ranges
  - IP addresses
- IP address inventory
- IP address range groups

IP Address Blocks

IP address blocks are the highest-level entities within an IP address space organization. Conceptually, an IP block is either one of the private IP address spaces or a public IP address space as assigned to an organization by various Regional Internet Registries. Network administrators use IP address blocks to create and allocate IP address ranges to DHCP. They can add, import, edit, and delete IP address blocks. IPAM automatically maps IP address subnets to the appropriate IP address block based on the boundaries of the range. IPAM utilization statistics and trends are summarized at the block level.

IP Address Subnets

IP address subnets are the next hierarchical level of address space entities after IP address blocks. IPAM summarizes utilization statistics and trends at the IP address subnet level for the IP address ranges contained within the IP address subnet. Additionally, subnets can be created as either physical or virtual; if subnets are virtual, they can be assigned to either a provider or a customer virtual network.

IP Address Ranges

IP address ranges are the next hierarchical level of IP address space entities after IP address subnets. Conceptually, an IP address range is an IP subnet, or part of an IP subnet marked by a start and end IP address. It typically corresponds to a DHCP scope, or to a static IPv4 or IPv6 address range or address pool that is used to assign addresses to hosts. An IP address range is uniquely identifiable by the value of the mandatory Managed by Service and Service Instance options, which help IPAM manage and maintain overlapping or duplicate IP address ranges from the same console. You can add or import IP address ranges from within the IPAM console. Whenever an IP address range is created, it is automatically associated with an IP address subnet. If a subnet does not exist, one can be automatically created when the IP address range is created.

IP Addresses

IP addresses are the addresses that make up the IP address range. IPAM enables end-to-end life cycle management of IPv4 and IPv6 addresses, including record synchronization with DHCP and DNS servers. IPAM automatically maps an address to the appropriate range based on the start and end address of the range. An IP address is uniquely identifiable by the value of mandatory Managed By Service and Service Instance options that help IPAM manage and maintain duplicate IP addresses from the same console. You can add or import IP addresses from within the IPAM console.

IP Address Inventory

In the IP address inventory view, you can view a list of all IP addresses in the enterprise along with their device names and type. IP address inventory is a logical group defined by the Device Type option within the IP addresses view. These groups allow you to customize the way your address space displays for managing and tracking IP usage. You can add or import IP addresses from within the IPAM console. For example, you could add the IP addresses for printers or routers, assign IP address the appropriate device type of printer or router, and then view your IP inventory filtered by the device type that you assigned.

IP Address Range Groups

IPAM enables you to organize IP address ranges into logical groups. For example, you might organize IP address ranges geographically or by business division. Logical groups are defined by selecting the grouping criteria from built-in or user-defined custom fields.

Monitoring and Managing

IPAM enables automated, periodic service monitoring of DHCP and DNS servers across a forest. Monitoring and managing is organized into the views listed in the following table.

View	Description
DNS and DHCP servers	By default, managed DHCP and DNS servers are arranged by their network interface in /16 subnets for IPv4 and /48 subnets for IPv6. You can select the view to see just DHCP scope properties, just DNS server properties, or both.
DHCP scopes	The DHCP scope view enables scope utilization monitoring. Utilization statistics are collected periodically and automatically from a managed DHCP server. You can track important scope properties such as Name**,** ID**,** Prefix Length, and Status.
DNS zone monitoring	Zone monitoring is enabled for forward and reverse lookup zones. Zone status is based on events collected by IPAM. The status of each zone is summarized.
Server groups	You can organize your managed DHCP and DNS servers into logical groups. For example, you might organize servers by business unit or geography. Groups are defined by selecting the grouping criteria from built-in fields or user-defined fields.

Adding Address Spaces to IPAM

An address space is a container that consists of a set of connected IP blocks, IP subnets, IP ranges or IP addresses. The IP ADDRESS SPACE pane contains all the IP objects discovered or created. Non-virtualized network objects are always in the IP ADDRESS SPACE pane.

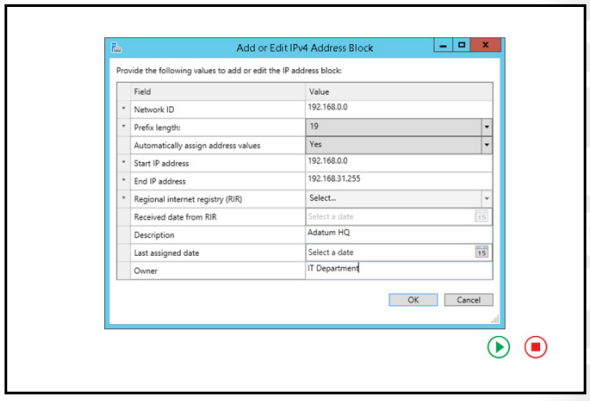

When you manually add IP addresses to IPAM, you can add either IPv4 or IPv6 addresses. When you use the IPAM console to add IP addresses, default values are automatically filled in for required fields, except for the IP addresses. You can add or import any of the following:

- IP Address Block. When you add an IP Address Block, supplying the Network ID and Prefix length allows the start IP address and End IP address to be calculated automatically for you. Additionally, if you enter a non-private IP address range, you must specify the Regional Internet Registry where the addresses are registered and the date range for the registration. Optionally, you can add a brief description and an owner.

  The following Windows PowerShell cmdlet **Add-IpamBlock** can also be used to add an IP Address block:

  ```
  Add-IpamBlock –NetworkID <network prefix, in Classless InterDomain Routing (CIDR) notation> –Rir <string>
  ```

  The RIR value is optional for private addresses. If you are specifying the RIR then the value must be one of: **AFRNIC**, **APNIC**, **ARIN**, **LACNIC**, or **RIPE**.

- IP Address Subnet. When you add an IP Address subnet, you must provide a friendly name for the subnet. Additionally, you must specify the Network ID and Prefix length.

  There are several optional settings when adding an IP Address subnet. You can specify one or more VLANs to be associated with the subnet, whether or not the subnet is virtualized, or custom fields such as AD site or VMM IP Pool Name. As with the other IP address types, you can add a brief description and an owner.
  The Windows PowerShell cmdlet **Add-IpamSubnet** can also be used to add an IP address subnet. When using Add-IpamSubnet you must also specify it the network type is NonVirtualized, Provider, or Customer IP Subnet. You must specify the address space to which the Customer IP Subnet will be added.

  ```
  Add-IpamSubnet –NetworkID <network prefix, in Classless InterDomain Routing (CIDR) notation> -
  Rir <string>
  ```

- IP Address Range. You can use an IP Address range to further divide an IP Subnet. When you create an IP address range you must specify the Network ID and either the Prefix length or Subnet mask. Additionally, if an IP address does not already exist that contains the addresses in the IP address range you are creating, you can select to have one automatically created. The other required fields, Managed by Service, Service Instance, and Assignment Type will use default values unless otherwise specified. As with the other IP address types, there is a large variety of custom fields available to describe the IP address range.

  You can also use the Windows PowerShell cmdlet **Add-IpamRange** to add an IP Address range. When you use **Add-IpamRange**, you must also specify if the network type is NonVirtualized, Provider, or Customer IP range. You must specify the address space to which the Customer IP Subnet will be added.

  ```
  Add-IpamRange –NetworkID <network prefix, in Classless InterDomain Routing (CIDR) notation> -
  CreateSubnetIfNotFound
  ```

- IP Address. IPAM provides end-to-end management of IP addresses, including synchronization with DHCP and DNS. You can use the IP address to associate the address with DHCP reservations; however, when you use Windows PowerShell to create the IP address, IPAM does not automatically create the reservation. You can discover duplicate addresses by the Managed by Service and Service Instance properties of an IP address. IPAM automatically maps an address to the range containing the address. When creating IP an address, the only required information that you have to provide is the IP address itself. The other required fields, Managed by Service, Service Instance, Device Type, Address State, and Assignment Type will use default values unless otherwise specified. As with the other IP address types, there is a large variety of custom fields available to describe the IP address.

  You can use the Windows PowerShell cmdlet **Add-IpamAddress** to add an IP Address. When you use **Add-IpamAddress**, you must also specify the IP address.

  ```
  Add-IpamAddress –IpAddress <x.x.x.x>
  ```

Importing and Updating Address Spaces

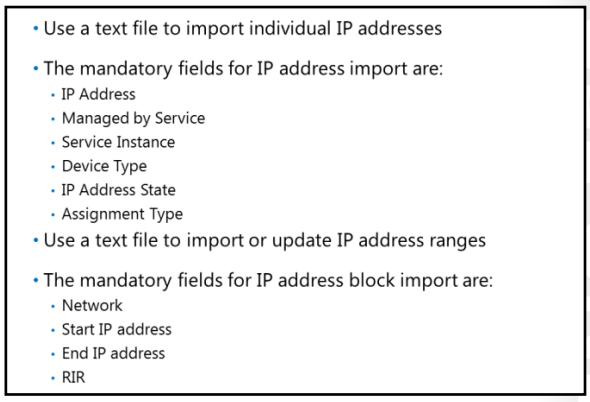

You can create IP address objects when you import IP address information into IPAM using a text file. When you import information for a file, you must include the required fields for the address type as you do when you add addresses through the console. The file you create is a comma delimited file with the field names in the first row.

You can import information into custom fields; however, they must be defined prior to importing the data, and the defined field name must be included in the first line. The fields are not required to be in any particular order; however, the data must be in the same order as the fields.

When you create test files, the following rules apply to the data:

- Field names and Data can be enclosed quotes.
- Field names and Data can contain spaces.
- Field names and Data are not case sensitive.
- Data must be valid for the field into which it is being imported.

For example, the you can use the following to import two addresses into the IPAM database managing a DHCP server named DHCP1.adatum.com:

```
“IP Address”,”Managed by Service”,”Service Instance”,”Device Type”,”IP Address
State”,”Assignment Type”
10.10.0.25,ms dhcp,dhcp1.adatum.com,host,in-use,static
10.10.0.26,ms dhcp,dhcp1.adatum.com,host,in-use,static
```

For IP address blocks, subnets and ranges, the network ID and network prefix length are combined in a single field named Network. For example, to import an IP Address block of 65.52.0.0/14 assigned by the ARIN regional authority, use the following in a text file:

```
“Network”,“Start IP address”,“End IP address”,RIR
65.52.0.0/14,65.52.0.0,65.52.255.255,ARIN
```

If a required field is missing or you try to import the wrong data type for a field, an error report is created in the user’s Documents folder. The mandatory fields for importing data are as follows:

- IP address block import: Network, Start IP address, End IP address, RIR
- IP address subnet import: Name, Network
- IP address range import: Network, Start IP address, End IP address, Managed by Service, Service Instance, Assignment Type, Utilization Calculation
- IP address import: IP address, Managed by Service, Service Instance, Device Type, IP Address State, Assignment Type

Import and Update IP Address Ranges

You can import and update data for IP address ranges. Updating the data will delete ranges that are no longer present. You do not have to perform the update step, as performing the import step only will create new ranges as appropriate. The import and update process is specific for a defined Managed by Service and Service Instance pair.

Finding, Allocating, and Reclaiming IP Addresses

There are two operations that can be performed on an IP Address Range: You can Find and Allocate Available IP Address or Reclaim IP Addresses.

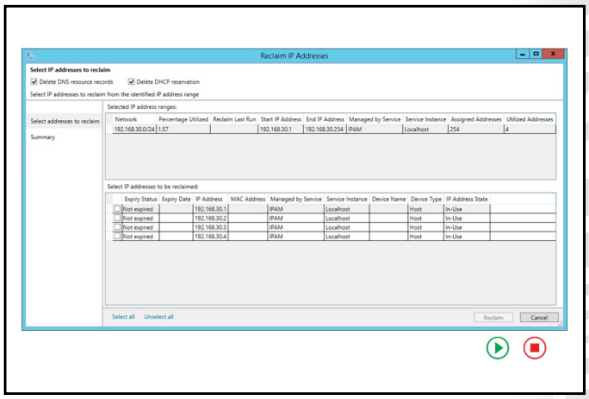

Find and Allocate Available IP Address

To find available IP addresses in a range, in the IP ADDRESS SPACE, change the current view to IP Address Ranges. The Find and Allocate Available IP Addresses task is available by right-clicking the desired IP address range. Choosing this operation opens the Find and Allocate Available IP Addresses dialog box. IPAM will search the range starting with the first unassigned IP address. If the address range is part of a managed DHCP scope, IPAM ignores all reserved or excluded IP addresses. The search includes a ping of the address and a DNS PTR query; If both methods fail to get a response, the address is available. You can then press the Find Next button to move to the next unassigned IP address.

Once you have completed the searches, you can allocate one of the addresses that you found. By default the last address you find is highlighted, you can allocate that address or select a different address is found. Choose each section to configure the IP address as needed, the Basic Configurations section is filled in with the selected IP address and the default values for the mandatory fields.

If you configure the DHCP Reservation and DNS Record sections, only the IPAM database is affected by default. If you want to configure a DHCP reservation for a managed IP address, complete the DHCP Reservation section and check the Automatically create DHCP reservation for this IP address check box. To create DNS records for the selected IP Address, complete the DNS Record section and check the Automatically create DNS records for this IP address check box.

Reclaim IP Addresses

When manually added IP Addresses are no longer in use you need to reclaim them to make them available for use with other devices. Additionally, the reclaim operations cleans DHCP reservations and DNS records on managed DNS and DHCP servers. There are two ways to reclaim IP addresses:

To reclaim IP addresses in a range, in the IP ADDRESS SPACE, change the current view to IP Address Ranges. The Reclaim IP Addresses task is available if you right-click the desired IP address range. If you choose this operation, it opens the Reclaim IP Addresses dialog box.

The Reclaim IP Addresses dialog box displays all the utilized IP addresses for the range, the IP Address State, and additional information such as the Device Name and Device Type. Once you have determined the IP addresses that you want to reclaim, check the select check box next to the IP addresses, and click the Reclaim button. By default, this operation removes the DNS resource records and DHCP reservations.

Maintaining IP Address Inventory in IPAM

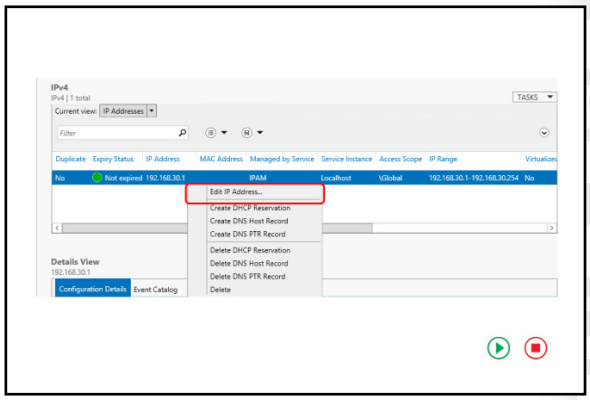

You can manage individual addresses through the IPAM console as necessary. When you want to manage an individual address, there are two locations from where this can be accomplished. In the IP ADDRESS SPACE, you change the current view to IP Addresses or the IP Address Inventory pane. The management options are available when you right-click the desired IP address. The context menu includes three types of operations, Edit, Create, and Delete.

Edit IP Address

The Edit IP Address dialog box allows you to add information to an IP address or change information that was previously configured. You can modify all aspects of the IP address information.

Create Operations

There are three options available for creating records for an IP Address. These include:

- Create DHCP Reservation. This option creates a DHCP reservation in the appropriate IP Address Range.
- Create DNS Host Record. This option creates a DNS record on the appropriate DNS server or servers for the IP Address Range.
- Create DNS PTR Record. This option creates a DNS OTR record on the appropriate DNS server or servers for the IP Address Range.

Delete Operations

There are four options available for deleting IP Addresses or the information associated with them. These include:

- Delete. The delete option will remove the IP address from the IPAM database. By default, this will remove the DNS records and DHCP reservations if they exist.
- Delete DHCP Reservation. The option will remove any DHCP reservations created for the IP address, without removing the IP address from the IPAM database.
- Delete DNS Host Record. The option will remove any DNS Host Records for the IP address, without removing the IP address from the IPAM database.
- Delete DNS PTR Record. The option will remove any DNS PTR Records for the IP address, without removing the IP address from the IPAM database.

Demonstration: Using IPAM to Manage IP Addressing

In this demonstration, you will see how to use IPAM to manage IP addressing.

Demonstration Steps

1. On LON-SVR2, add an IP address block in the IPAM console with the following parameters:
   - Network ID: **172.16.0.0**
   - Prefix length: **16**
   - Description: **Head Office**
2. Add IP addresses for the network router by adding to the IP Address Inventory with the following parameters:
   - IP address: **172.16.0.1**
   - MAC address: **112233445566**
   - Device type: **Routers**
   - Description: **Head Office Router**
3. Use the IPAM console to create a DHCP reservation as follows:
   - IP address: **172.16.0.101**
   - MAC address: **223344556677**
   - Device type: **Host**
   - Client ID: **Associate MAC to Client ID** check box
   - Reservation server name: **LON-DC1.Adatum.com**
   - Reservation name: **Webserver**
   - Reservation type: **Both**
4. Use the IPAM console to create the DNS host record as follows:
   - Device name: **Webserver**
   - Forward lookup zone: **Adatum.com**
   - Forward lookup primary server: **LON-DC1.adatum.com**
   - **Automatically create DNS records for this IP address**
5. On LON-DC1, open the DHCP console and confirm that the reservation was created in the 172.16.0.0 scope.
6. On LON-DC1, open the DNS Manager console and confirm that the DNS host record was created.

IPAM Monitoring

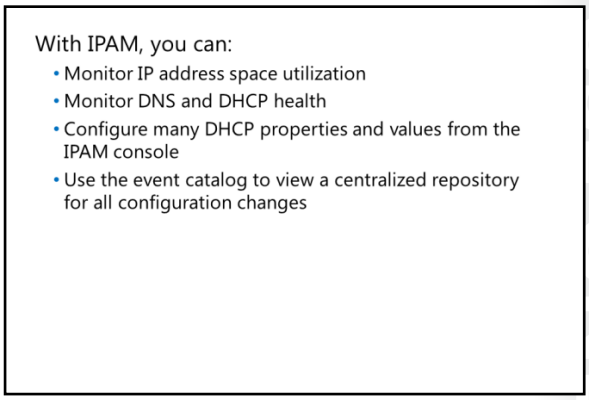

The IPAM address space management feature allows you to efficiently view, monitor, and manage the IP address space on the network. Address space management supports IPv4 public and private addresses, and IPv6 global and unicast addresses. Using the MONITOR AND MANAGE section and the DNS and DHCP, DHCP Scopes, DNS Zone Monitoring, and Server Groups views, you can view and monitor health and configuration of all the DNS and DHCP servers that are being managed by IPAM. IPAM uses scheduled tasks to periodically collect data from managed servers. You can also retrieve data on demand by using the Retrieve All Server Data option.

Utilization Monitoring

Utilization data is maintained for IP address ranges, IP address blocks, and IP range groups within IPAM. You can configure thresholds for the percentage of the IP address space that is utilized, and then use those thresholds to determine under-utilization and over-utilization.

You can perform utilization trend building and reporting for IPv4 address ranges, blocks, and range groups. The utilization trend window allows you to view trends over time periods such as daily, weekly, monthly, or annually, or you can view trends over custom date ranges. Utilization data from managed DHCP scopes is auto-discovered, and you can view this data.

Monitoring DHCP and DNS

Using IPAM, you can monitor DHCP and DNS servers from any physical location of the enterprise. One of the primary benefits of IPAM is its ability to simultaneously manage multiple DHCP servers or DHCP scopes that are spread across one or more DHCP servers.

The IPAM monitoring view allows you to view the status and health of selected sets of Microsoft DNS and DHCP servers from a single console. IPAM's monitoring view displays the basic health of servers and recent configuration events that occurred on these servers. The monitoring view also allows you to organize the managed servers into logical sever groups.

For DHCP servers, the server view allows you to track various server settings, server options, the number of scopes, and the number of active leases that are configured on the server. For DNS servers, this view allows you to track all zones that are configured on the server, along with details of the zone type. The view also allows you to see the total number of zones that are configured on the server, and the overall zone health status as derived from the zone status of individual zones on the server.

DHCP Server Management

From the IPAM console, you can manage DHCP servers and perform the following actions:

- Edit DHCP server properties.
- Edit DHCP server options.
- Create DHCP scopes.
- Configure predefined options and values.
- Configure the user class across multiple servers simultaneously.
- Create and edit new and existing user classes across multiple servers simultaneously.
- Configure the vendor class across multiple servers simultaneously.

- Start the management console for a selected DHCP server.
- Retrieve server data from multiple servers.

DNS Server Management

You can start the DNS management console for any managed DNS server from a central console in the IPAM server. Once you start the DNS management console, you can retrieve server data from the selected set of servers. The DNS Zone Monitoring view displays all the forward lookup and reverse lookup zones on all the DNS servers that IPAM is currently managing. For the forward lookup zones, IPAM also displays all the servers that are hosting the zone, the aggregate health of the zone across all these servers, and the zone properties.

The Event Catalog

The IPAM event catalog provides a centralized repository for auditing all configuration changes that are performed on DHCP servers that are managed from a single IPAM management console. The IPAM configuration events console gathers all of the configuration events. These configuration event catalogs allows you to view, query, and generate reports of the consolidated configuration changes, along with details specific to each record.

Demonstration: Using IPAM Monitoring

In this demonstration, you will see how to use the IPAM console to monitor DNS and DHCP.

Demonstration Steps

1. On LON-SVR2, review the information displayed in the **DNS and DHCP Servers** pane in the IPAM console.
2. Review the information in the **DHCP Scopes** pane.
3. Review the information in the **DNS Zone Monitoring** pane.
4. Review the information in the **Server Groups** pane.

Lab: Implementing Advanced Network Services

Scenario

A. Datum Corporation has grown rapidly over the last few years. The company has deployed several new branch offices, and it has significantly increased the number of users in the organization. Additionally, it has expanded the number of partner organizations and customers that are accessing A. Datum websites and applications. Because of this expansion, the complexity of the network infrastructure has increased, and the organization now needs to be much more aware of network-level security.

As one of the senior network administrators at A. Datum, you are responsible for implementing some of the advanced networking features in Windows Server 2012 R2 to manage the networking infrastructure. You need to implement new features in DHCP and DNS, with the primary goal of providing higher levels of availability while increasing the security of these services. You also need to implement IPAM so that you can simplify and centralize the management of the IP address usage and configuration in an increasingly complex network.

Objectives

In this lab, you will see how to:

- Configure advanced DHCP settings.
- Configure advanced DNS settings.
- Configure IP address management.

Lab Setup

Estimated Time: 70 minutes

Virtual machines	20412C-LON-DC1 20412C-LON-SVR1 20412C-LON-SVR2 20412C-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 **20412C-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 to 4 for **20412C-LON-SVR1** and **20412C-LON-SVR2**. Do not start **20412C-LON-CL1** until directed to do so.

Exercise 1: Configuring Advanced DHCP Settings

Scenario

With the expansion of the network, and the increased availability and security requirements at A. Datum Corporation, you need to implement additional DHCP features. Because of the recent business expansion, the main office DHCP scope is almost completely utilized, which means that you need to configure a superscope. Additionally, you need to configure DHCP Name Protection and DHCP failover.

The main tasks for this exercise are as follows:

- Configure a superscope.
- Configure DHCP name protection.
- Configure and verify DHCP failover.

▶ Task 1: Configure a superscope

1. On LON-DC1, configure a scope named **Scope1**, with a range of **192.168.0.50 – 192.168.0.100**, and with the following settings:
    - Subnet mask: **255.255.255.0**
    - Router: **192.168.0.1**
    - DNS Suffix: **Adatum.com**
    - Choose to activate the scope later.
2. Configure a second scope named **Scope2** with a range of **192.168.1.50 – 192.168.1.100**, and with the following settings:
    - Subnet mask: **255.255.255.0**
    - Router: **192.168.1.1**
    - DNS Suffix: **Adatum.com**
    - Choose to activate the scope later.
3. Create a superscope called **AdatumSuper** that has **Scope1** and **Scope2** as members.
4. Activate the **AdatumSuper** superscope.

▶ Task 2: Configure DHCP name protection

- Switch to the DHCP console on LON-DC1, and enable **DHCP Name Protection** found on the DNS tab of the IPv4 node.

▶ Task 3: Configure and verify DHCP failover

1. On LON-SVR1, start the DHCP console and observe the current state of DHCP. Note that the server is authorized, but that no scopes are configured.
2. On LON-DC1, in the DHCP console, launch the Configure Failover Wizard.
3. Configure failover replication with the following settings:
    - Partner server: **172.16.0.21**
    - Relationship Name: **Adatum**
    - Maximum Client Lead Time: **15 minutes**
    - Mode: **Load balance**
    - Load Balance Percentage: **50%**

   - State Switchover Interval: **60 minutes**
   - Message authentication shared secret: **Pa$$w0rd**
4. Complete the Configure Failover Wizard.
5. On LON-SVR1, refresh the IPv4 node. Notice that the IPv4 node is active, and that Scope Adatum is configured.
6. Start **20412C-LON-CL1**, and sign in as **Adatum\Administrator**.
7. Configure LON-CL1 to obtain an IP address from the DHCP server.
8. Open a command prompt window, and record the IP address.
9. Switch to LON-DC1, and stop the DHCP server service.
10. Switch back to LON-CL1, and renew the IP address.
11. On LON-DC1, in the Services console, start the DHCP server service.
12. Close the Services console.

Results: After completing this exercise, you will have configured a superscope, configured DHCP Name Protection, and configured and verified DHCP failover.

Exercise 2: Configuring Advanced DNS Settings

Scenario

To increase the level of security for the DNS zones at A. Datum, you need configure DNS security settings such as DNSSEC, DNS socket pool, and cache locking. A. Datum has a business relationship with Contoso, Ltd and will host the Contoso.com DNS zone. A. Datum clients use an application that accesses a server named App1 in the Contoso.com zone by using its NetBIOS name. You need to ensure that these applications can resolve the names of the required servers correctly. You will employ a GlobalNames zone to achieve this.

The main tasks for this exercise are as follows:

- Configure DNSSEC.
- Configure the DNS socket pool.
- Configure DNS cache locking.
- Configure a GlobalNames zone.

▶ Task 1: Configure DNSSEC

1. On LON-DC1, start the DNS Manager.
2. Use the DNSSEC Zone Signing Wizard to sign the **Adatum.com** zone.
3. Choose to customize zone-signing parameters.
4. Ensure that DNS server **LON-DC1** is the Key Master.
5. Add the **Key Signing Key** by accepting the default values for the new key.
6. Add the **Zone Signing Key** by accepting the default values for the new key.
7. Choose to use **NSCE3** with the default values.
8. Choose to enable the distribution of trust anchors for this zone.

9. Accept the default values for **signing and polling**.
10. Verify that the DNSKEY resource records have been created in the **Trust Points zone**.
11. Minimize the DNS console.
12. Use the Group Policy Management Console, in the Default Domain Policy object, to configure the Name Resolution Policy Table.
13. Create a rule that enables DNSSEC for the Adatum.com suffix, and that requires DNS clients to verify that the name and address data were validated.

▶ Task 2: Configure the DNS socket pool

1. On LON-DC1, start Windows PowerShell.
2. Run the following command to view the current size of the socket pool:

```
Get-DNSServer
```

3. Run the following command to change the socket pool size to 3,000:

```
dnscmd /config /socketpoolsize 3000
```

4. Restart the DNS service.
5. Run the following command to confirm the new socket pool size:

```
Get-DnsServer
```

▶ Task 3: Configure DNS cache locking

1. Run the following command to view the current cache lock size:

```
Get-DnsServer
```

2. Run the following command to change the cache lock value to 75 percent:

```
Set-DnsServerCache –LockingPercent 75
```

3. Restart the DNS service.
4. Run the following command to confirm the new cache lock value:

```
Get-DnsServer
```

▶ Task 4: Configure a GlobalNames zone

1. Create an Active Directory-integrated forward lookup zone named **Contoso.com**, by running the following command:

```
Add-DnsServerPrimaryZone –Name Contoso.com –ReplicationScope Forest
```

2. Run the following command to enable support for GlobalName zones:

```
Set-DnsServerGlobalNameZone –AlwaysQueryServer $true
```

3. Create an Active Directory-integrated forward lookup zone named **GlobalNames** by running the following command:

```
Add-DnsServerPrimaryZone –Name GlobalNames –ReplicationScope Forest
```

4. Open the DNS Manager console, and add a new host record to the Contoso.com domain named **App1** with the IP address of **192.168.1.200**.
5. In the GlobalNames zone, create a new alias named **App1** using the FQDN of **App1.Contoso.com**.
6. Close DNS Manager, and close the command prompt.

Results: After completing this exercise, you will have configured DNSSEC, the DNS socket pool, DNS cache locking, and the GlobalName zone.

Exercise 3: Configuring IPAM

Scenario

A. Datum Corporation is evaluating solutions for simplifying IP address management. Since implementing Windows Server 2012, you have decided to implement IPAM.

The main tasks for this exercise are as follows:

- Install the IPAM feature.
- Configure IPAM–related GPOs.
- Configure IP management server discovery.
- Configure managed servers.
- Configure and verify a new DHCP scope with IPAM.
- Configure IP address blocks, record IP addresses, and create DHCP reservations and DNS records.
- To prepare for the next module.

▶ Task 1: Install the IPAM feature

- On LON-SVR2, install the **IP Address Management (IPAM) Server** feature by using the Add Roles and Features Wizard in Server Manager.

▶ Task 2: Configure IPAM–related GPOs

1. On LON-SVR2, in the Server Manager, in the IPAM Overview pane, provision the IPAM server using Group Policy.
2. Enter **IPAM** as the GPO name prefix, and provision IPAM using the Provision IPAM Wizard.

▶ Task 3: Configure IP management server discovery

1. In the IPAM Overview pane, configure server discovery for the Adatum domain.
2. In the IPAM Overview pane, start the server discovery process. Discovery may take 5 to 10 minutes to run. The yellow bar will indicate when discovery is complete.

▶ Task 4: Configure managed servers

1. In the IPAM Overview pane, add the servers that you need to manage. Verify that IPAM access is currently blocked for both LON-DC1 and LON-SVR1.
2. Use Windows PowerShell to grant the IPAM server permission to manage by running the following command:

```
Invoke-IpamGpoProvisioning –Domain Adatum.com –GpoPrefixName IPAM –IpamServerFqdn
LON-SVR2.adatum.com –DelegatedGpoUser Administrator
```

3. For both LON-DC1 and LON-SVR1, set the manageability status to **Managed**.
4. Switch to LON-DC1, and force the update of Group Policy using **gpupdate /force**.
5. Switch to LON-SVR1, and force the update of Group Policy using **gpupdate /force**.
6. Return to LON-SVR2 and refresh the server access status for LON-DC1 and LON-SVR1 and the Server Manager console view. It may take up to 10 minutes for the status to change. If necessary, repeat both refresh tasks as needed until a green check mark displays next to LON-DC1 and the IPAM Access Status displays as **Unblocked**.
7. In the IPAM Overview pane, right click LON-SVR1 and **Retrieve All Server Data**.
8. In the IPAM Overview pane, right-click LON-DC1 and **Retrieve All Server Data**.

▶ Task 5: Configure and verify a new DHCP scope with IPAM

1. On LON-SVR2, use IPAM to create a new DHCP scope with the following parameters:
   - Scope Name: **TestScope**
   - Scope start address: **10.0.0.50**
   - Scope end address: **10.0.0.100**
   - Subnet mask: **255.0.0.0**
   - Default gateway: **10.0.0.1**
2. Use IPAM to configure failover for the TestScope on LON-DC1 with the following parameters:
   - Partner server: LON-SVR1.adatum.com
   - Relationship name: TestFailover
   - Shared secret: Pa$$w0rd
   - Maximum client lead time: 15 minutes
   - Mode: Load balance
   - Load balance percentage: 50%
   - State Switchover Interval: 60 minutes
3. On LON-DC1, verify the scope in the DHCP MMC.
4. On LON-SVR1, verify the scope in the DHCP MMC.

▶ Task 6: Configure IP address blocks, record IP addresses, and create DHCP reservations and DNS records

1. On LON-SVR2, add an IP address block in the IPAM console with the following parameters:
   - Network ID: **172.16.0.0**
   - Prefix length: **16**
   - Description: **Head Office**
2. Add IP addresses for the network router by adding to the IP Address Inventory with the following parameters:
   - IP address: **172.16.0.1**
   - MAC address: **112233445566**

- Device type: **Routers**
- Description: **Head Office Router**

3. Use the IPAM console to create a DHCP reservation as follows:
   - IP address: **172.16.0.10**
   - MAC address: **223344556677**
   - Device type: **Host**
   - Client ID: **Associate MAC to Client ID checkbox**
   - Reservation server name: **LON-DC1.Adatum.com**
   - Reservation name: **Webserver**
   - Reservation type: **Both**
4. Use the IPAM console to create the DNS host record as follows:
   - Device name: **Webserver**
   - Forward lookup zone: **Adatum.com**
   - Forward lookup primary server: **LON-DC1.adatum.com**
   - **Automatically create DNS records for this IP address**
5. On LON-DC1, open the DHCP console and confirm that the reservation was created in the 172.16.0.0 scope.
6. On LON-DC1, open the DNS Manager console and confirm that the DNS host record was created.

▶ Task 7: To prepare for the next module

1. On the host computer, start the Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20412C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20412C-LON-SVR1**, **20412C-LON-SVR2**, and **20412C-LON-CL1**.

Results: After completing this exercise, you will have installed IPAM and configured IPAM with IPAM-related GPOs, IP management server discovery, managed servers, a new DHCP scope, IP address blocks, IP addresses, DHCP reservations, and DNS records.

Question: Will client computers immediately stop communicating on the network if there is no functioning DHCP server?

Question: What is the default size of the DNS socket pool?

Question: What value does the DNS cache lock use to determine when to update an IP address in the DNS cache?

Module Review and Takeaways

Best Practices

- Implement DHCP failover to ensure that client computers can continue to receive IP configuration information in the event of a server failure.
- Ensure that there are at least two DNS servers hosting each zone.
- Use IPAM to control IP address distribution and static address assignments.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Users can no longer access a vendor's website that they have historically been able to access.	
Managed servers are unable to connect to the IPAM server.	

Review Question

Question: What is one of the drawbacks of using IPAM?

Real-world Issues and Scenarios

Some network clients are receiving incorrect DHCP configuration. What tool should you use to begin the troubleshooting process?

Answer: The **IPConfig /All** command will report to you if the client is receiving DHCP configuration, and if so, the IP address of the DHCP server from which the configuration came.

What are some possible causes of the incorrect configurations?

Answer: There may be a rogue DHCP server on the network. Common things to look for will be gateway devices—such as cable modems or Private Branch Exchange (PBX) boxes—that have a DHCP component enabled. Another possibility is that someone has manually configured the IP address on the client.

Tools

Tool	Use	Location
Dnscmd	Configure all aspects of DNS management	%systemroot%\System32\dnscmd.exe
DHCP console	Control all aspects of DHCP management from a user interface	%systemroot%\System32\dhcpmgmt.msc
DNS console	Control all aspects of DNS management from a user interface	%systemroot%\System32\dnsmgmt.msc
IPAM management console	Control all aspects of IPAM management	Server Manager

Module 2

Implementing Advanced File Services

Contents:

Module Overview

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

In addition to minimizing disk space, another storage concern is the connection between the storage and the remote disks. Internet SCSI (iSCSI) storage in Windows Server 2012 is a cost-effective feature that helps create a connection between the servers and the storage. To implement iSCSI storage in Windows Server 2012, you must be familiar with the iSCSI architecture and components. In addition, you must be familiar with the tools that are provided in Windows Server to implement an iSCSI-based storage.

In organizations with branch offices, you have to consider slow links and how to use these links efficiently when sending data between your offices. The Windows BranchCache® feature in Windows Server 2012 helps address the problem of slow connectivity. This module explains the BranchCache feature, and how to configure it.

Objectives

After completing this module, you will be able to:

- Configure iSCSI storage.
- Configure BranchCache.
- Optimize storage usage.
- Implement advanced file services.

Lesson 1
Configuring iSCSI Storage

iSCSI storage is an inexpensive and simple way to configure a connection to remote disks. Many application requirements dictate that remote storage connections must be redundant in structure to ensure fault tolerance or high availability. In addition, many companies already have fault-tolerant networks, which can be made redundant relatively inexpensively when compared to using storage area networks (SANs). In this lesson, you will learn how to create a connection between servers and iSCSI storage. You will perform these tasks by using IP-based iSCSI storage. You will also learn how to create both single and redundant connections to an iSCSI target. You will practice this by using the iSCSI initiator software that is available in Windows Server 2012.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe iSCSI and its components.
- Describe the iSCSI target server and the iSCSI initiator.
- Implement high availability for iSCSI.
- Describe iSCSI security options.
- Configure the iSCSI target.
- Connect to iSCSI storage.
- Describe considerations for implementing the iSCSI storage solution.

What Is iSCSI?

iSCSI is a protocol that supports access to remote, small computer system interface (SCSI)-based storage devices over a TCP/IP network. iSCSI carries standard SCSI commands over IP networks to facilitate data transfers over intranets, and to manage storage over long distances. You can use iSCSI to transmit data over local area networks (LANs), wide area networks (WANs), or even over the Internet.

iSCSI transmits SCSI commands over IP networks

Component	Description
IP network	Provides high performance and redundancy
iSCSI targets	Run on the storage device and enable access to the disks
iSCSI initiators	A software component or host adapter on the server that provides access to iSCSI targets
IQN	A globally unique identifier used to address initiators and targets on an iSCSI network

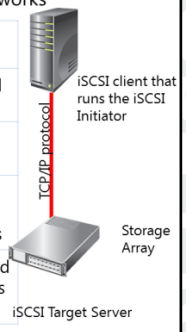

iSCSI relies on standard Ethernet networking architecture. Specialized hardware such as host bus adapters (HBA) or network switches are optional. iSCSI uses TCP/IP (typically, TCP port 3260). This means that iSCSI simply enables two hosts to negotiate tasks—for example, session establishment, flow control, and packet size—and then exchange SCSI commands by using an existing Ethernet network. By doing this, iSCSI uses a popular, high performance, local storage bus subsystem architecture, and emulates it over LANs and WANs to create a storage area networks (SANs). Unlike some SAN technologies, iSCSI requires no specialized cabling. You can run it over the existing switching and IP infrastructure. However, you can increase the performance of an iSCSI SAN deployment by operating it on a dedicated network or subnet, as best practices recommend.

Note: Although you can use a standard Ethernet network adapter to connect the server to the iSCSI storage device, you can also use dedicated iSCSI HBAs.

An iSCSI SAN deployment includes the following:

- TCP/IP network. You can use standard network interface adapters and standard Ethernet protocol network switches to connect the servers to the storage device. To provide sufficient performance, the network should provide speeds of at least 1 gigabit per second (Gbps), and should provide multiple paths to the iSCSI target. As a best practice, use a dedicated physical and logical network to achieve fast, reliable throughput.
- iSCSI targets. This is another method of gaining access to storage. iSCSI targets present, or advertise storage, similar to controllers for hard disk drives of locally attached storage. However, this storage is accessed over a network instead of locally. Many storage vendors implement hardware-level iSCSI targets as part of their storage device's hardware. Other devices or appliances, such as Windows Storage Server 2012 devices, implement iSCSI targets by using a software driver together with at least one Ethernet adapter. Windows Server 2012 provides the iSCSI target server—which is effectively a driver for the iSCSI protocol—as a role service.
- iSCSI initiators. The iSCSI target displays storage to the iSCSI initiator (also known as the *client*), which acts as a local disk controller for the remote disks. All versions of Windows Server beginning with Windows Server® 2008 include the iSCSI initiator, and can connect to iSCSI targets.
- iSCSI Qualified Name (IQN). IQNs are globally unique identifiers that are used to address initiators and targets on an iSCSI network. When you configure an iSCSI target, you must configure the IQN for the iSCSI initiators that will be connecting to the target. iSCSI initiators also use IQNs to connect to the iSCSI targets. However, if name resolution on the iSCSI network is a possible issue, iSCSI endpoints (both target and initiator) can be identified by their IP addresses.

  Question: Can you use your organization's internal TCP/IP network to provide iSCSI?

iSCSI Target Server and iSCSI Initiator

The iSCSI target server and the iSCSI initiator are described below.

iSCSI Target Server

The iSCSI target server role service provides for software-based and hardware-independent iSCSI disk subsystems. You can use the iSCSI target server to create iSCSI targets and iSCSI virtual disks. You can then use Server Manager to manage these iSCSI targets and virtual disks.

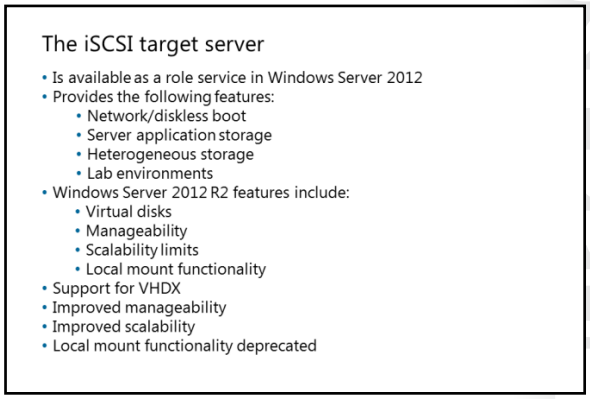

The iSCSI target server that is included in Windows Server 2012 provides the following functionality:

- Network/diskless boot. By using boot-capable network adapters or a software loader, you can use iSCSI targets to deploy diskless servers quickly. By using differencing virtual disks, you can save up to 90 percent of the storage space for the operating system images. This is ideal for large deployments of identical operating system images, such as a Hyper-V server farm, or for high-performance computing (HPC) clusters.

- Server application storage. Some applications such as Hyper-V and Microsoft® Exchange Server require block storage. The iSCSI target server can provide these applications with continuously available block storage. Because the storage is remotely accessible, it can also combine block storage for central or branch office locations.
- Heterogeneous storage. iSCSI target server supports iSCSI initiators that are not based on the Windows operating system, so you can share storage on Windows servers in mixed environments.
- Lab environments. The iSCSI target server role enables your Windows Server 2012 computers to be network-accessible block storage devices. This is useful in situations in which you want to test applications before deploying them on SAN storage.

Windows Server 2012 R2 Enhancements

Windows Server® 2012 R2 provides the following new or updated features for iSCSI target server:

Feature/Functionality	New/Updated	Description
Virtual disks	New	Windows Server 2012 R2 provides support for the new VHDX format, which has a much larger storage capacity than the older VHD format. Windows Server 2012 R2 also provides data corruption protection during power failures.
Manageability	Updated	Windows Server 2012 R2 uses the Storage Management Initiative Specification (SMI-S) provider with Microsoft System Center 2012-Virtual Machine Manager to manage iSCSI Target Server in a hosted and/or private cloud. There are also new Windows PowerShell cmdlets that allow the export and import of configuration files.
Scalability limits	Updated	The maximum number of sessions per target server is increased to 544, and the maximum number of logical units per target server is increased to 256.
Local mount functionality	Updated	Local mount functionality for snapshots has been deprecated; however, you can use the loopback initiator to access exported snapshots.

iSCSI target servers that provide block storage utilize your existing Ethernet network; no additional hardware is required. If high availability is an important criterion, consider setting up a high availability cluster. With a high availability cluster, you will need shared storage for the cluster—either hardware Fibre Channel storage, or a Serial Attached SCSI storage array. The iSCSI target server integrates directly into the failover cluster feature as a cluster role.

iSCSI Initiator

The iSCSI initiator service has been a standard component installed by default since Windows Server 2008 and Windows Vista®. To connect your computer to an iSCSI target, you simply start the Microsoft iSCSI Initiator service, and then configure it.

The features in Windows Server 2012 include:

- Authentication. You can enable Challenge Handshake Authentication Protocol (CHAP) to authenticate initiator connections, or you can enable reverse CHAP to allow the initiator to authenticate the iSCSI target.
- Query initiator computer for ID. This is only supported with Windows 8.1 or Windows Server 2012.

Options for locating iSCSI targets

iSCSI initiators can discover iSCSI targets using four different methods, as outlined in the following table:

Discovery mechanism	Description
SendTargets	• The iSCSI Initiator performs an iSCSI discovery login and then a SendTargets operation on portals (where a portal is a target's IP and TCP port number pair) that are statically configured in the iSCSI Initiator properties on the Discovery tab or by using the **iscsicli AddTargetPortal** command. Discovery occurs under three conditions: • When a target portal is added. • When the service starts. • When a management application requests it.
Internet Storage Name Service (iSNS)	• You need to set a static address for the iSNS server by using the **iscsicli AddiSNSServer** command. The iSCSI initiator gets a list of targets under three conditions: • When the service starts. • When an application requests it. • When the iSNS server sends a state notification change.
Manually configured targets	iSCSI targets can be manually configured in the iSCSI Initiator properties on the Targets tab or by using the **iscsicli AddTarget** command. Manually configured targets are persistent. Those targets that are not configured as Hidden are available whenever the service restarts.
HBA discovery	iSCSI HBAs that conform to interfaces in iSCSI Initiator can perform target discovery through the interface between the HBA and the iSCSI Initiator service. The HBA will respond to requests from the iSCSI Initiator to send a list of targets. The iSCSI Initiator will make this request when the service starts or when the HBA tells the iSCSI Initiator that the list has changed.

Introduction of iSCSI Target in Windows Server 2012

http://go.microsoft.com/fwlink/?LinkId=270038

Question: When would you consider implementing diskless booting from iSCSI targets?

Implementing High Availability for iSCSI

In addition to configuring the basic iSCSI target server and iSCSI initiator settings, you can integrate these services into more advanced configurations.

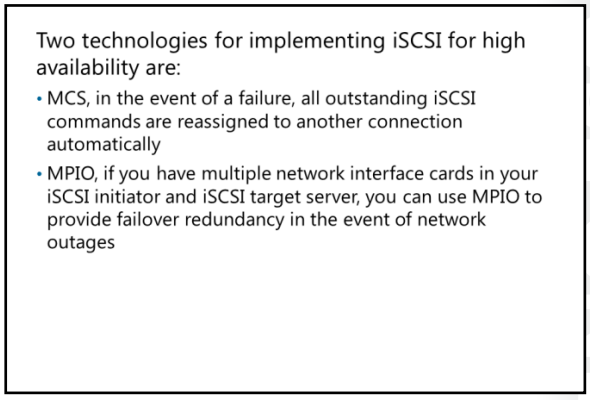

Configuring iSCSI for High Availability

Creating a single connection to iSCSI storage makes that storage available. However, it does not make that storage highly available. If iSCSI loses the connection, the server loses access to its storage. Therefore, most iSCSI storage connections are made redundant through one of two high availability technologies: Multiple Connected Session (MCS) and Multipath I/O (MPIO).

Although similar in results they achieve, these two technologies use different approaches to achieve high availability for iSCSI storage connections.

MCS is a iSCSI protocol feature that:

- Enables multiple TCP/IP connections from the initiator to the target for the same iSCSI session.
- Supports automatic failover. If a failure occurs, all outstanding iSCSI commands are reassigned to another connection automatically.
- Requires explicit support by iSCSI SAN devices, although the Windows Server 2012 iSCSI target server role supports it.

MPIO provides redundancy differently, as follows:

- If you have multiple network interface cards in your iSCSI initiator and iSCSI target server, you can use MPIO to provide failover redundancy during network outages.
- MPIO requires a device-specific module (DSM) if you want to connect to a third-party SAN device that is connected to the iSCSI initiator. The Windows operating system includes a default MPIO DSM that is installed as the MPIO feature within Server Manager.
- MPIO is widely supported. Many SANs can use the default DSM without any additional software, while others require a specialized DSM from the manufacturer.
- MPIO is more complex to configure, and is not as fully automated during failover as MCS is.

iSCSI Security Options

iSCSI is a protocol that provides access to storage devices over a TCP/IP network. Therefore, it is imperative that you secure your iSCSI solution to protect it from malicious users or attacks. Because iSCSI employs the existing network infrastructure, it is subject to the same types of security issues to which any network traffic is susceptible. iSCSI supports a comprehensive set of security features, including access control lists (ACLs), IP security protocol (IPsec), and CHAP.

Mitigate security risks to iSCSI by the following best practices:

- Segregate the iSCSI SAN channel
- Secure management consoles
- Disable unneeded services
- Use CHAP authentication
- Use IPsec authentication
- Use IPsec encryption

iSCSI Security Best Practices

You can mitigate risks to your iSCSI solution by using the following best practices for iSCSI security:

Best Practice	Description
Segregate the iSCSI SAN channel	Perhaps the most important security measure you can take is to segregate iSCSI traffic from other network traffic. This prevents exposure of the iSCSI storage to the open LAN. This segregation can be physical, by establishing network paths using separate network equipment, or it segregation can be logical through the use of virtual local area networks (VLANs) and ACLs at the network layer.
Secure management consoles	The management consoles that control access to data and storage allocation are often web-based and have well-known default passwords. Use dedicated systems to access these consoles.
Disable unneeded services	Services not directly related to the iSCSI implementation should not be running on systems involved in the iSCSI configuration.
Use CHAP authentication	CHAP has two possible configurations: one-way CHAP or mutual CHAP authentication. In one-way CHAP, the iSCSI target on the storage array authenticates the initiator on the server. In mutual CHAP, the target and initiator authenticate each other using separate secrets for each direction of the connection. iSCSI CHAP authentication can be configured using group policy. These policies can found in Computer Configuration, Administrative Templates, System iSCSI, and iSCSI Security. CHAP authentication can be configured on the Configuration tab of the iSCSI Initiator.
Use IPsec authentication	Although CHAP authentication is effective, it is subject to vulnerabilities such as dictionary attacks. IPsec provides much stronger authentication. IPsec operates at the IP packet layer and can use pre-shared keys, Kerberos or certificate-based authentications. IPsec may not be supported on all storage platforms. Consult the vendor about support for IPsec.

Best Practice	Description
Use IPsec encryption	IPsec also supports IP traffic encryption to provide the highest levels of security, but that additional security affects performance. The encryption and decryption process requires much more processing power and should only be employed where the need is paramount, such as across untrusted networks. IPsec Tunnel Mode encryption can be configured on the Configuration tab of the iSCSI Initiator.

 **Note:** In Windows Server 2008 R2, the Storage Explorer snap-in for Microsoft Management Console (MMC) could be used to configure many iSCSI security settings for iSCSI initiators. That feature has been removed in Windows Server 2012 R2

Demonstration: Configuring an iSCSI Target

In this demonstration, you will see how to:

- Add the iSCSI target server role service.
- Create two iSCSI virtual disks and an iSCSI target.

Demonstration Steps

Add the iSCSI target server role service

1. On LON-DC1, open the Server Manager.
2. In the Add Roles and Features Wizard, install the following roles and features on the local server, and accept the default values:
   - **File And Storage Services (Installed)\File and iSCSI Services\iSCSI Target Server**

Create two iSCSI virtual disks and an iSCSI target

1. On LON-DC1, in the Server Manager, in the navigation pane, click **File and Storage Services**, and then click **iSCSI**.
2. In the iSCSI VIRTUAL DISKS pane, click **TASKS**, and then in the TASKS drop-down list box, click **New iSCSI Virtual Disk**.
3. Create a virtual disk with the following settings:
   - Name: **iSCSIDisk1**
   - Disk size: **5 GB**
   - iSCSI target: **New**
   - Target name: **LON-SVR2**
   - Access servers: **172.16.0.22**
4. On the **View results** page, wait until creation completes, and then close the **View Results** page.
5. In the iSCSI VIRTUAL DISKS pane, click **TASKS**, and then in the **TASKS** drop-down list, click **New iSCSI Virtual Disk**.

6. Create a virtual disk that has these settings:
    - Name: **iSCSIDisk2**
    - Disk size: **5 GB**
    - iSCSI target: **LON-SVR2**
7. On the **View results** page, wait until creation completes, and then close the **View Results** page.

Demonstration: Connecting to the iSCSI Storage

In this demonstration, you will see how to:

- Connect to the iSCSI target.
- Verify the presence of the iSCSI drive.

Demonstration Steps

Preparation steps

Before you start this demonstration, perform the following steps:

1. On the host computer, in the Hyper-V Manager console, open the settings for 20412C-LON-SVR2 virtual machine.
2. In the **Settings for 20412C-LON-SVR2** window, ensure that both legacy network adapters are connected to Private Network.
3. If a legacy network adapter has a status of Not connected, connect it to the virtual network named Private Network.

Connect to the iSCSI target

1. Sign in on 20412C-LON-SVR2 with username **Adatum\Administrator** and the password **Pa$$w0rd**.
2. Open Server Manager, and on the **Tools** menu, open **iSCSI Initiator**.
3. In the **iSCSI Initiator Properties** dialog box, configure the following:
    - Quick Connect: **LON-DC1**
    - Discover targets: **iqn.1991-05.com.microsoft:lon-dc1-lon-svr2-target**

Verify the presence of the iSCSI drive

1. In Server Manager, on the Tools menu, open **Computer Management**.
2. In the Computer Management console, under Storage node, access **Disk Management**. Notice that the new disks are added. However, they all are currently offline and not formatted.
3. Close the Computer Management console.

Considerations for Implementing iSCSI Storage

When you design your iSCSI storage solution, consider the following best practices:

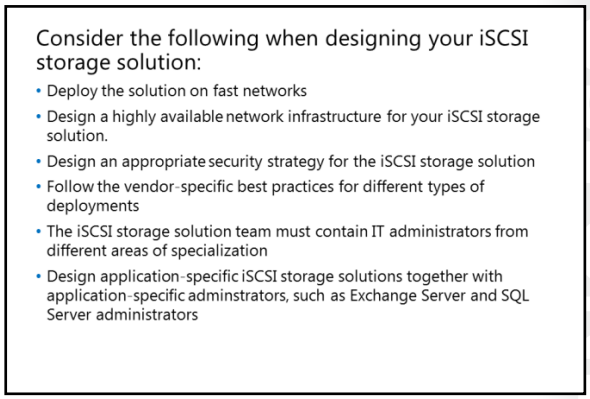

- Deploy the iSCSI solution on networks with a speed of at least 1 Gbps.
- High availability design for network infrastructure is crucial because data from servers to iSCSI storage is transferred through network devices and components.
- Design an appropriate security strategy for the iSCSI storage solution.
- Read the vendor-specific best practices for different types of deployments and applications that will use iSCSI storage solution, such as Exchange Server and Microsoft SQL Server®.
- IT personnel who will design, configure, and administer the iSCSI storage solution must include IT administrators with different areas of specialization, such as Windows Server 2012 administrators, network administrators, storage administrators, and security administrators. This is necessary so that the iSCSI storage solution has optimal performance and security, and has consistent management and operations procedures.
- When designing an iSCSI storage solution, the design team should also include application-specific administrators, such as Exchange Server administrators and SQL Server administrators, so that you can implement the optimal configuration for the specific technology or solution.

Lesson 2
Configuring BranchCache

Branch offices have unique management challenges. A branch office typically has slow connectivity to the enterprise network and limited infrastructure for securing servers. In addition, you need to back up data that you maintain in your remote branch offices, which is why organizations prefer to centralize data where possible. Therefore, the challenge is to provide efficient access to network resources for users in branch offices. BranchCache helps you overcome these problems by caching files so they do not have to be transferred repeatedly over the network.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe how BranchCache works.
- Describe BranchCache requirements.
- Configure BranchCache server settings.
- Configure BranchCache client settings.
- Configure BranchCache.
- Monitor BranchCache.

How Does BranchCache Work?

The BranchCache feature that was introduced with Windows Server 2008 R2 and Windows 7 reduces the network use on WAN connections between branch offices and headquarters by locally caching frequently used files on computers in the branch office. BranchCache improves the performance of applications that use one of the following protocols:

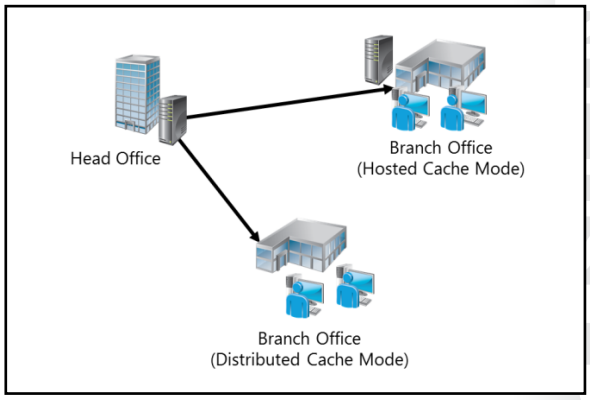

- HTTP or HTTPS protocols. These protocols are used by web browsers and other applications.
- Server Message Block (SMB), including signed SMB traffic protocol. This protocol is used for accessing shared folders.
- Background Intelligent Transfer Service (BITS). A Windows component that distributes content from a server to clients by using only idle network bandwidth. BITS is also a component that Microsoft System Center Configuration Manager uses.

When the client requests the data, BranchCache retrieves data from a server. Because BranchCache is a passive cache, it will not increase WAN use. BranchCache only caches the read requests, and will not interfere when a user saves a file.

BranchCache improves the responsiveness of common network applications that access intranet servers across slow WAN links. Because BranchCache does not require additional infrastructure, you can improve the performance of remote networks by deploying Windows 7 or newer client computers, and by deploying Windows Server 2008 R2 or newer servers, and then enabling the BranchCache feature.

BranchCache maintains file and folder permissions to ensure that users only have access to files and folders for which they have permission.

BranchCache works seamlessly with network security technologies, including Secure Sockets Layer (SSL), SMB signing, and end-to-end IPsec. You can use BranchCache to reduce network bandwidth use and to improve application performance, even if the content is encrypted.

You can configure BranchCache to use hosted cache mode or distributed cache mode, which are described below:

- Hosted cache mode. This mode operates by deploying a server that is running Windows Server 2008 R2 or newer versions as a hosted cache server in the branch office. Client computers locate the server so that they can retrieve content from the hosted cache when the hosted cache is available. If the content is not available in the hosted cache, the content is retrieved from the content server by using a WAN link, and then is provided to the hosted cache so that the successive client requests can retrieve it from there.
- Distributed cache mode. For smaller remote offices, you can configure BranchCache in the distributed cache mode without requiring a server. In this mode, local client computers running Windows 7 or Windows 8.1 maintain a copy of the content and make it available to other authorized clients that request the same data. This eliminates the need to have a server in the branch office. However, unlike the hosted cache mode, this configuration works per subnet only. In addition, clients who hibernate or disconnect from the network cannot reliably provide content to other requesting clients.

 **Note:** When using BranchCache, you may use both modes in your organization, but you can configure only one mode per branch office.

BranchCache functionality in Windows Server 2012 R2 has the following improvements:

- To allow for scalability, BranchCache allows for more than one hosted cache server per location.
- A new underlying database uses the Extensible Storage Engine (ESE) database technology from Exchange Server. This enables a hosted cache server to store significantly more data (even up to terabytes).
- A simpler deployment means that you do not need a Group Policy Object (GPO) for each location. To deploy BranchCache, you only need a single GPO that contains the settings. This also enables clients to switch between hosted cache mode and distributed mode when they are traveling between locations, without needing to use site-specific GPOs, which should be avoided in multiple scenarios.

How Client Computers Retrieve Data by Using BranchCache

When BranchCache is enabled on both a client computer and a server, and when the client computer is using the HTTP, HTTPS, or SMB protocol, the client computer performs the following process to retrieve data:

1. The client computer that is running Windows 8.1 connects to a content server in the head office that is running Windows Server 2012, and the content is initially requested the same way it would be without using BranchCache.
2. The content server in the head office authenticates the user and verifies that the user is authorized to access the data.
3. Instead of sending the content itself, the content server in the head office returns hashes as identifiers of the requested content to the client computer. The content server sends that data over the same connection that the content would have been sent over typically.

4. Using retrieved identifiers, the client computer does the following:
   - If you configure the client computer to use distributed cache, then it multicasts on the local subnet to find other client computers that have already downloaded the content.
   - If you configure the client computer to use hosted cache, then it searches for the content on the configured hosted cache.
5. If the content is available in the branch office, either on one or more clients or on the hosted cache, the client computer retrieves the data from the branch office. The client computer also ensures that the data is updated and has not been tampered with or corrupted.
6. If the content is not available in the remote office, then the client computer retrieves the content directly from the server across the WAN link. The client computer then either makes it available on the local network to other requesting client computers (distributed cache mode), or sends it to the hosted cache, where it is made available to other client computers.

BranchCache Requirements

BranchCache optimizes traffic flow between head offices and branch offices. Windows Server 2008 R2 and newer, and client computers running Windows 7 Enterprise or Windows Vista Ultimate versions and Windows 8.1 Enterprise can benefit from using BranchCache. Earlier versions of Windows operating systems do not benefit from this feature. You can use BranchCache to cache only the content that is stored on file servers or web servers that are running Windows Server 2008 R2 or newer.

Requirements for using BranchCache	Requirements for the modes
• Install the BranchCache feature or the BranchCache for Network Files feature on the server that is hosting the content • Configure client computers by using either Group Policy or the **netsh** command	• In distributed cache mode, the content cache at a branch office is distributed between client computers • In hosted cache mode, the content cache at the branch office is hosted on one or more server computers that are hosted cache servers

• Content versions:
  • V1 uses fixed file segment sizes
  • V2 uses smaller, variable size segments	

Requirements for Using BranchCache

To use BranchCache for file services, you must perform the following tasks:

- Install the BranchCache feature or the BranchCache for Network Files role service on the host server that is running Windows Server 2012 R2.
- Configure client computers either by using Group Policy or the **netsh branchcache set service** command.

If you want to use BranchCache to cache content from the file server, you must perform following tasks:

- Install BranchCache for the Network Files role service on the file server.
- Configure hash publication for BranchCache.
- Create BranchCache-enabled file shares.

If you want to use BranchCache for caching content from the web server, you must install the BranchCache feature on the web server. You do not need additional configurations.

BranchCache is supported on the full installation and Server Core installation of Windows Server 2008 R2 or newer. By default, BranchCache is not installed on Windows Server 2012 R2. Server Core installations of Windows Server 2008 R2 Enterprise with Hyper-V or Windows Server 2012 R2 Datacenter with Hyper-V cannot be used as BranchCache content servers.

Requirements for Distributed Cache Mode and Hosted Cache Mode

In the distributed cache mode, BranchCache works without a dedicated server. BranchCache works between clients on the same site. If client computers are configured to use the distributed cache mode, any client computer can use a multicast protocol called WS-Discovery to search locally for the computer that has already downloaded and cached the content. You should configure the client firewall to enable incoming traffic, HTTP, and WS-Discovery.

Windows 8.1 clients, however, will search for a hosted cache server, and if they discover one, will automatically self-configure as hosted cache mode clients. In the hosted cache mode, the client computers automatically self-configure as hosted cache mode clients, and they will search for the host server so that they can retrieve content from the hosted cache. Furthermore, you can use Group Policy so that you can use the fully qualified domain name (FQDN) of the hosted cache servers or enable automatic hosted cache discovery by service connection points. You must configure a firewall to enable incoming HTTP traffic from the hosted cache server.

In both cache modes, BranchCache uses the HTTP protocol for data transfer between client computers and the computer that is hosting the cached data.

Content Versions

Content cached on Windows Server 2008 R2 and Windows 7 is named version 1 (V1), whereas content that is cached on Windows Server 2012 R2 and Windows 8.1 is version 2 (V2). V1 content is a fixed file-segment size and is larger than in V2. Because of these large sizes, when a user makes a change that modifies the file length, the segment that changed and all the following segments of the file are re-sent over the WAN link. V2 content is more tolerant to changes within the file. Only the changed content will be re-sent, and it will use less WAN bandwidth.

When you have content servers and hosted cache servers that are running Windows Server 2012, they use the content version that is appropriate based on the operating system of the BranchCache client that requests content. When computers running Windows Server 2012 and Windows 8.1 request content, the content and hosted cache servers use V2 content; when computers running Windows Server 2008 R2 and Windows 7 request content, the content and hosted cache servers use V1 content.

When you deploy BranchCache in distributed cache mode, clients that use different content information versions do not share content with each other.

Configuring BranchCache Server Settings

You can use BranchCache to cache web content, which is delivered by HTTP or HTTPS. You can also use BranchCache to cache shared folder content, which is delivered by the SMB protocol.

The following table lists the servers that you can configure for BranchCache.

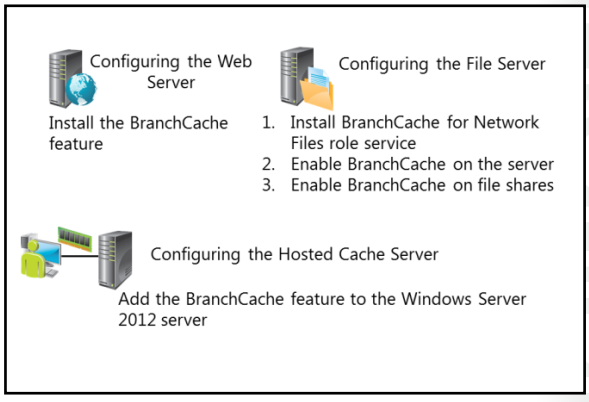

Server	Description
Web server or BITS server	To configure a Windows Server 2012 R2 web server or an application server that uses the BITS protocol, install the BranchCache feature. Ensure that the BranchCache service has started. Then, configure clients that will use the BranchCache feature. No additional web server configuration is required.
File server	You must install the BranchCache for the Network Files role service of the File Services server role before you enable BranchCache for any file shares. After you install the BranchCache for the Network Files role service, use Group Policy to enable BranchCache on the server. You must then configure each file share to enable BranchCache.
Hosted cache server	For the hosted cache mode, you must add the BranchCache feature to the Windows Server 2012 R2 server that you are configuring as a hosted cache server. To help secure communication, client computers use Transport Layer Security (TLS) when communicating with the hosted cache server. By default, BranchCache allocates five percent of the disk space on the active partition for hosting cache data. However, you can change this value by using Group Policy or by running the **netsh branchcache set cachesize** command.

Configuring BranchCache Client Settings

You do not have to install the BranchCache feature on client computers, because BranchCache is already included if the client is running Windows 7 or Windows 8.1. However, BranchCache is disabled by default on client computers. To enable and configure BranchCache, you must perform the following steps:

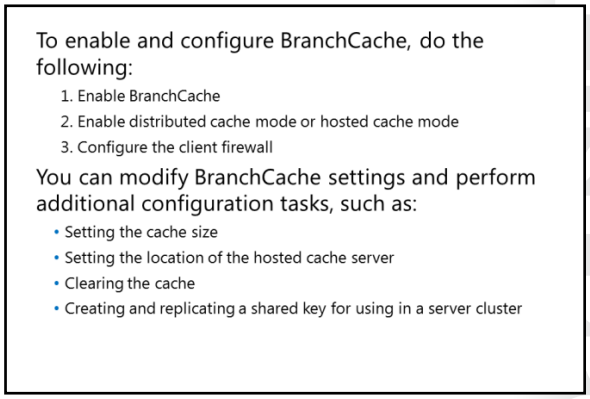

1. Enable BranchCache.
2. Enable the distributed cache mode or the hosted cache mode. Windows 8.1 clients can use either mode dynamically.
3. Configure the client firewall to enable BranchCache protocols.

Enabling BranchCache

You can enable the BranchCache feature on client computers by using Group Policy, Windows® PowerShell, or the **netsh branchcache set service** command.

To enable BranchCache settings by using Group Policy, perform the following steps for a domain-based GPO:

1. Open the Group Policy Management Console.
2. Create a GPO that will be linked to the organizational unit (OU) where the branch office client computers are located.

3. In a GPO, browse to **Computer Configuration\Policies\Administrative Templates: Policy definitions (ADMX files) retrieved from the local computer\Network**, and then click **BranchCache**.
4. Enable the Turn on BranchCache setting in the GPO.

Enabling the Distributed Cache Mode or Hosted Cache Mode

You can configure the BranchCache mode on client computers by using Group Policy, Windows PowerShell, or the **netsh branchcache set service** command.

To configure BranchCache mode by using Group Policy, perform the following steps for a domain-based GPO:

1. Open the Group Policy Management Console.
2. Create a GPO that will be linked to the OU where client computers are located.
3. In a GPO, browse to **Computer Configuration\Policies\Administrative Templates: Policy definitions (ADMX files) retrieved from the local computer\Network**, and then click **BranchCache**.
4. Select either the distributed cache mode or the hosted cache mode. You may also enable both the distributed cache mode and automatic hosted cache discovery by Service Connection Point policy settings. The client computers will operate in distributed cache mode unless they find a hosted cache server in the branch office. If they find a hosted cache server in the branch office, they will work in hosted cache mode.

 **Note:** A number of the GPO settings in the BranchCache container require at least Windows Server 2012, Windows 8, or Windows RT.

To enable BranchCache with Windows PowerShell, use the **Enable-BCDistributed** or **Enable-BCHostedServer** cmdlets. You can also use **Enable-BCHostedClient** cmdlet to configure BranchCache to operate in hosted cache client mode.

For example, the following cmdlet enables hosted cache client mode using the SRV1.adatum.com computer as a hosted cache server for Windows 7 clients and HTTPS.

```
Enable-BCHostedClient –ServerNames SRV1.adatum.com –UseVersion Windows7
```

The following cmdlet enables hosted cache mode and register service connection point in Active Directory® Domain Services (AD DS).

```
Enable-BCHostedServer –RegisterSCP
```

The following cmdlet enables distributed cache mode on the server.

```
Enable-BCDistributed
```

To configure BranchCache settings by using the **netsh branchcache set service** command, open a command prompt window, and perform the following steps:

1. Type the following netsh syntax for the distributed cache mode:

```
netsh branchcache set service mode=distributed
```

2. Type the following netsh syntax for the hosted mode:

```
netsh branchcache set service mode=hostedclient location=<hosted cache server>
```

Configuring the Client Firewall to Enable BranchCache Protocols

In the distributed cache mode, BranchCache clients use the HTTP protocol for data transfer between client computers, and the WS-Discovery protocol for cached content discovery. You should configure the client firewall to enable the following incoming rules:

- BranchCache–Content Retrieval (use HTTP)
- BranchCache–Peer Discovery (use WS–Discovery)

In hosted cache mode, BranchCache clients use the HTTP protocol for data transfer between client computers, but this mode does not use the WS-Discovery protocol. In the hosted cache mode, you should configure the client firewall to enable the incoming rule, BranchCache–Content Retrieval (use HTTP).

Additional Configuration Tasks for BranchCache

After you configure BranchCache, clients can access the cached data in BranchCache-enabled content servers, which are available locally in the branch office. You can modify BranchCache settings and perform additional configuration tasks, such as:

- Setting the cache size.
- Setting the location of the hosted cache server.
- Clearing the cache.
- Creating and replicating a shared key to use in a server cluster.

Demonstration: Configuring BranchCache

In this demonstration, you will see how to:

- Add BranchCache for the Network Files role service.
- Configure BranchCache in Local Group Policy Editor.
- Enable BranchCache for a file share.

Demonstration Steps

Add BranchCache for the Network Files role service

1. On LON-DC1 open the Server Manager.
2. In the Add Roles and Features Wizard, install the following roles and features to the local server:
    - **File and Storage Services (Installed)\File and iSCSI Services\BranchCache for Network Files**

Enable BranchCache for the server

1. On the Start screen, type **gpedit.msc**, and then press Enter.
2. Browse to **Computer Configuration\Administrative Templates\Network\Lanman Server**, and do the following:
    - Enable **Hash Publication for BranchCache**
    - Select **Allow hash publication only for shared folder on which BranchCache is enabled**

Enable BranchCache for a file share

1. Open a File Explorer window, and on drive C, create a folder named **Share**.
2. Configure the **Share** folder properties as follows:
   - Enable **Share this folder**
   - Check **Enable BranchCache** in **Offline Settings**

Monitoring BranchCache

After the initial configuration, you want to verify that BranchCache is configured correctly and functioning correctly. You can use the **netsh branchcache show status all** command to display the BranchCache service status. You can also use the **Get-BCStatus** cmdlet to provide BranchCache status and configuration information. The client and hosted cache servers display additional information, such as the location of the local cache, the size of the local cache, and the status of the firewall rules for HTTP and WS-Discovery protocols that BranchCache uses.

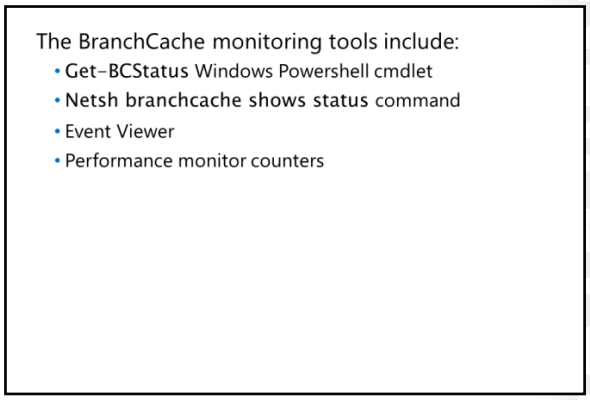

You can also use the following tools to monitor BranchCache:

- Event Viewer. Use this tool to monitor the BranchCache events that are recorded in both the Application log and the Operational log. The Application log is located in the Windows Logs folder, and the Operational log is located in the Application and Service Logs\Microsoft\Windows\BranchCache folder.
- Performance counters. Use this tool to monitor BranchCache performance monitor counters. BranchCache performance monitor counters are useful debugging tools for monitoring BranchCache effectiveness and health. You can also use BranchCache performance monitoring to determine the bandwidth savings in the Distributed Cache mode or in the hosted cache mode. If you have implemented Microsoft System Center 2012-Operations Manager in the environment, you can use the Windows BranchCache Management Pack for Operations Manager 2012.

Lesson 3
Optimizing Storage Usage

Every organization stores data on different storage systems. As storage systems process more and more data at higher speeds, the demand for disk space for storing the data has increased. The large volume of files, folders, and information, and the way they are stored, organized, managed, and maintained, becomes a challenge for organizations. Furthermore, organizations must satisfy requirements for security, compliance, and data leakage prevention for a company's confidential information.

Windows Server 2012 introduces many technologies that can help organizations respond to the challenges of managing, maintaining, securing, and optimizing data that is stored on different storage devices. The technologies include the File Server Resource Manager (FSRM), File Classification Infrastructure, and data deduplication, each of which provides new features as compared to Windows Server 2008 R2.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the FSRM.
- Describe File Classification Management.
- Describe file classification properties.
- Describe file classification rules.
- Configure classification management.
- Describe considerations for using file classification options for storage optimization in Windows Server 2012
- Describe storage optimization options.
- Configure data deduplication.
- Describe tiered storage.

What Is FSRM?

You can use the FSRM to manage and classify data that is stored on file servers. FSRM includes the following features:

You can use the FSRM to manage and classify data that is stored on file servers	
FSRM features	New FSRM features
• File classification infrastructure • File management tasks • Quota management • File screening management • Storage reports	• DAC • Manual classification • Access-denied assistance • File management tasks • Automatic classification

- File Classification Infrastructure. This feature automates the data classification process. You can apply access policies dynamically to files based on their classification. Example policies include Dynamic Access Control (DAC) for restricting access to files, file encryption, and file expiration. You can classify files automatically by using file classification rules, or manually by modifying the properties of a selected file or folder. You can modify file properties automatically based on the application type or content of the file, or by manually setting options on the server that will trigger file classification.

- File management tasks. You can use this feature to apply a conditional policy or action to files, based on their classification. The conditions of a file management task include the file location, the classification properties, the date the file was created, the last modified date of the file, and the last time that the file was accessed. The actions that a file management task can take include the ability to expire files, encrypt files, and run a custom command.
- Quota management. You can use this feature to limit the space that is allowed for a volume or folder. You can apply quotas automatically to new folders that are created on a volume. You can also define quota templates that you can apply to new volumes or folders.
- File screening management. You can use this feature to control the types of files that users can store on a file server. You can limit the extension that can be stored on your file shares. For example, you can create a file screen that disallows files with an MP3 extension from being stored in personal shared folders on a file server.
- Storage reports. You can use this feature to identify trends in disk usage, and identify how your data is classified. You can also monitor attempts by users to save unauthorized files.

You can configure and manage the FSRM by using the File Server Resource Manager MMC snap-in, or by using the Windows PowerShell command-line interface.

The following FSRM features are new with Windows Server 2012:

- Integration with DAC. DAC can use a File Classification Infrastructure to help you centrally control and audit access to files on your file servers.
- Manual classification. Manual classification enables users to classify files and folders manually without the need to create automatic classification rules.
- Access-denied assistance. You can use access-denied assistance to customize the access denied error message that displays for users in Windows 8.1 when they do not have access to a file or a folder.
- File management tasks. The updates to file management tasks include AD DS and Active Directory Rights Management Services (AD RMS) file management tasks, continuous file management tasks, and dynamic namespace for file management tasks.
- Automatic classification. The updates to automatic classification increase the level of control you have over how data is classified on your file servers, including continuous classification, Windows PowerShell for custom classification, updates to the existing content classifier, and dynamic namespace for classification rules.

What's New in File Server Resource Manager in Windows Server 2012

http://go.microsoft.com/fwlink/?LinkId=270039

Question: Are you currently using the FSRM in Windows Server 2008 R2? If yes, for what areas do you use it?

What Is File Classification Management?

The File Classification Infrastructure in Windows Server 2012 R2 allows administrators to classify files and create and apply policies based on those classifications. This can lower costs and help to manage the security risks involved with managing data. When File Classification Infrastructure is implemented, the storage layout is unaffected by data management requirements, and the organization can adapt more easily to a changing business and regulatory environment.

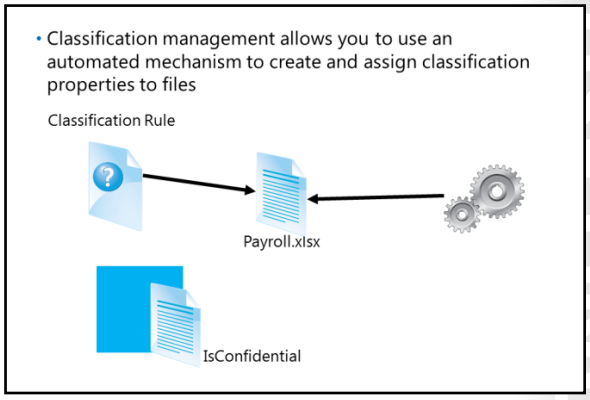

File Classification Management is a tool designed to ease the burden and management of data that is spread out in your organization. Using Classification Management, you can classify files in many different ways. In most scenarios, you perform classification manually. In Windows Server 2012, the File Classification Infrastructure feature allows organizations to automate the way in which data is classified. You can then assign file management policies based on a file's classification, and you can enforce corporate requirements based on the business value of the data. You can also modify the policies easily by using the built-in tools. You can use file classification to perform the following actions:

1. Define classification properties and values, which can be assigned to files by running classification rules.
2. Create, update, and run classification rules. Each rule assigns a single predefined property and value to files within a specified directory, based on installed classification plug-ins.

When you run a classification rule, you can reevaluate files that are already classified. You can choose to overwrite existing classification values, or add the value to properties that support multiple values.

What Are File Classification Properties?

File classification properties are used to assign values to files. There are many built-in property types from which you can choose. You can define these properties based on the needs of your organization. Classification properties are assigned to files that use classification rules, which are discussed in the next topic.

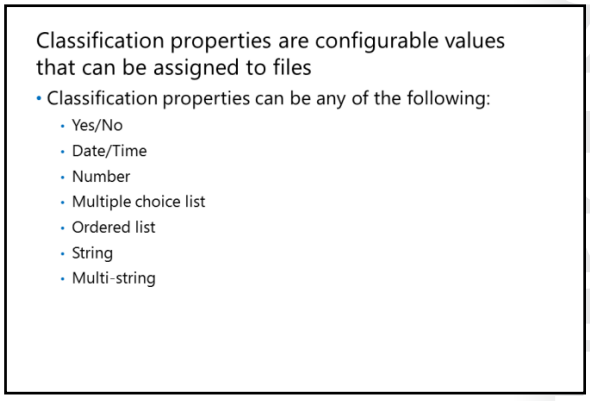

The following table defines the available property types, and the policy that is applied when a file is reclassified:

Property type	Description
Yes/No	A Boolean property that can have a value of either YES or NO. When multiple values are combined, a NO value overwrites a YES value.
Date-Time	A simple date and time property. When multiple values are combined, conflicting values prevent reclassification.
Number	A simple number property. When multiple values are combined, conflicting values prevent reclassification.
Multiple choice list	A list of values that can be assigned to a property. More than one value can be assigned to a property at a time. When multiple values are combined, each value in the list is used.
Ordered list	A list of fixed values. Only one value can be assigned to a property at a time. When multiple values are combined, the value highest in the list is used.
String	A simple string property. When multiple values are combined, conflicting values prevent reclassification.
Multi-string	A list of strings that can be assigned to a property. More than one value can be assigned to a property at a time. When multiple values are combined, each value in the list is used.

What Is a File Classification Rule?

A file classification rule assigns a classification property to a file system object. A classification rule includes information detailing when to assign a classification property to a file.

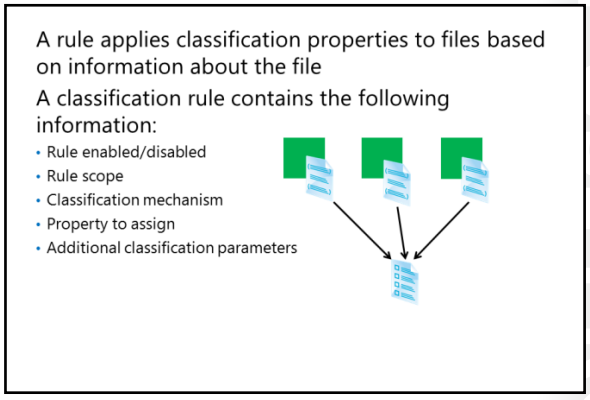

Key Classification Rule Properties

To define the behavior of a file classification rule, ask yourself the following questions:

- Is the rule enabled? On the classification rule Properties page, on the Rule Settings tab, the Enabled check box allows you to specifically disable or enable the classification rule.
- What is the scope of the rule? On the Rule Settings tab, the Scope parameter allows you to select a folder or folders to which the classification rule will apply. When the rule is run, it processes and attempts to classify all file system objects within this location.
- What classification mechanism will the rule use? On the classification rule Properties page, on the rule's Classification tab, you must choose a classification method that the rule will use to assign the classification property. By default, there are two methods from which you can choose:
    - Folder classifier. The folder classifier mechanism assigns properties to a file based on the file's folder path.

  - Content classifier. The content classifier searches for strings or regular expressions in files. This means that the content classifier classifies a file based on the textual contents of the file, such as whether it contains a specific word, phrase, numeric value, or type.

- What property will the rule assign? The main function of the classification rule is to assign a property to a file object based on how the rule applies to that file object. On the Classification tab, you must specify a property and the specific value that the rule will assign to that property.
- What additional classification parameters will be used? The core of the rule's logic lies in the additional classification parameters. Click the Advanced button on the Classification tab will open the Additional Classification Parameters window. Here, you can specify additional parameters—including strings or regular expressions—that, if found in the file system object, will cause the rule to apply itself. For example, this parameter could be the phrase "Social Security Number" or any number with the format 000-00-000. If this parameter is found, then the classification parameter will apply a YES value for a Confidential classification property to the file. This classification could then be leveraged to perform some tasks on the file system object, such as moving it to a secure location.

A classification parameter can be one of the following three types:

- RegularExpression. Match a regular expression by using the Microsoft .NET syntax. For example, **\d\d\d** will match any three-digit string.
- StringCaseSensitive. Match a case-sensitive string. For example, Confidential will only match Confidential, and not confidential or CONFIDENTIAL.
- String. Match a string, regardless of case. Confidential will match Confidential, confidential, and CONFIDENTIAL.

Classification Scheduling

You can run classification rules in two ways, on-demand, or based on a schedule. Either way you choose, each time you run classification, it uses all rules that you have left in the enabled state.

Configuring a schedule for classification allows you to specify a regular interval at which file classification rules will run, ensuring that your server's files are regularly classified and up to date with the latest classification properties.

Demonstration: How to Configure Classification Management

In this demonstration, you will see how to:

- Create a classification property.
- Create a classification rule.
- Modify the classification schedule.

Demonstration Steps

Create a classification property

1. Open **File Server Resource Manager**, and expand the **Classification Management** node.
2. Using the **Classification Properties** node, create a new **Classification Property** named **Confidential**, with the **Yes**/**No** property type.

Create a classification rule

1. Using the **Classification Rules** node, create a new **Classification Rule** named **Confidential Payroll Documents**.
2. Configure the rule to classify documents with a value of Yes for the **Confidential** classification property, if the file contains the string expression PAYROLL.
3. Configure the scope of the rule to apply to E:\Labfiles\Data

Modify the classification schedule

1. Create a classification schedule that runs every **Sunday** at **8:30 A.M**.
2. Using the **Classification Rule** node, manually run **Classification With All Rules Now**, and view the report. Ensure that **File3.txt** is listed at the bottom of the report.
3. Navigate to E:\Labfiles\Data and view the files to ensure that they were correctly classified.
4. Keep the virtual machines running for the next demonstration.

Considerations for Using File Classification Options for Storage Optimization in Windows Server 2012

Although Classification Management provides a powerful mechanism to catalog, categorize, and classify your file system objects, you should consider the following factors when you implementing Classification Management:

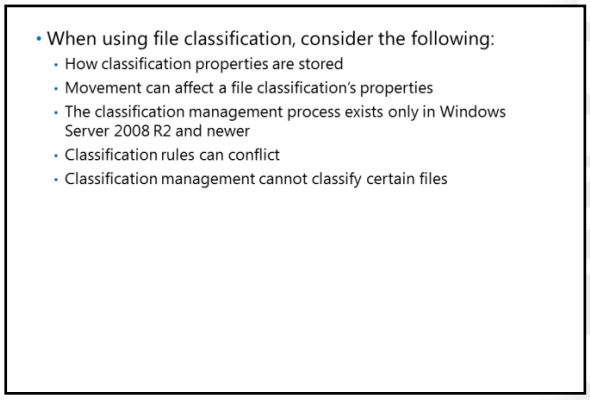

- How classification properties are stored. Classification properties are stored in an alternate data stream, which is a feature of the New Technology File System (NTFS). If a file moves within the NTFS, the alternate data streams move with the file, but they do not appear in the file's contents. In Microsoft Office applications, the classification properties are also stored within file formats as custom document properties or server document properties.
- How movement affects classification properties. When you move a file from one NTFS file system to another, if you use a standard mechanism such as Copy or Move, the file retains its classification properties. However, if you move a file to a non-NTFS file system, regardless of how you move the file, file classification properties are not retained. If the file is the product of a Microsoft Office application, then the classification properties remain attached, regardless of how the file is moved. File classification is currently not supported on the Resilient File System (ReFS).
- Classification Management process in Windows Server 2012 and Windows Server 2008. Classification properties are available only to servers that run Windows Server 2008 R2 or newer versions. However, Microsoft Office documents will retain classification property information in Document Properties, which is viewable regardless of the operating system being used.
- Conflicting classification rules. At times, classification rules can conflict. When this happens, the File Classification Infrastructure will attempt to combine properties. The following behaviors will occur when conflicting classification rules arise:
    - For Yes or No properties, a YES value takes priority over a NO value.

  - For ordered list properties, the highest property value takes priority.
  - For multiple choice properties, the property sets are combined into one set.
  - For multiple string properties, a multistring value is set that contains all the unique strings of the individual property values.
  - For other property types, an error occurs.
- Classification Management cannot classify certain files. File Classification Infrastructure will not identify individual files within a container, files such as a .zip or .vhd/vhdx file. In addition, File Classification Infrastructure will not allow content classification for the contents of encrypted files.

Options for Storage Optimization

Windows Server 2012 R2 has enhanced storage features such as File Access Auditing and network file system (NFS) data stores and data deduplication. This topic provides an overview of these features.

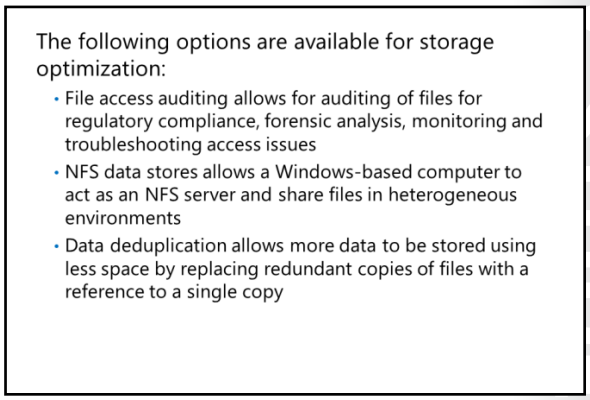

File Access Auditing

Audit policies can help you achieve a number of goals, such as regulatory compliance, monitoring, forensic analysis, and troubleshooting. The advanced security auditing capabilities in Windows Server 2012 R2 can be used with DAC to extend your overall security auditing strategy.
DAC enables you to create targeted audit policies by using expressions based on user, computer, and resource claims. For example, you can create an audit policy to track all Read and Write operations on files classified as high-business-impact (HBI) by employees who do not have a high-security clearance. Expression-based audit policies can be created directly for a file or folder, or centrally through the Advanced Audit Policy Configuration settings in Group Policy.

NFS Data Stores

NFS storage is supported in Windows Server 2012 R2 through the server for NFS role service. This allows a Windows-based computer to act as an NFS server and share files in a mixed environment of computers, operating systems, and networks. Once the role service is installed, NFS sharing can be configured on folders in the folder's properties on the NFS Sharing tab. Users on computers running NFS client software can access directories (shares) on computers running server for NFS role service by mounting those shares to their computers. Files on mounted NFS shares are accessed in the same manner as local files are accessed.

Data Deduplication

In an enterprise with a large volume of shared files, there are often duplicate files, or parts of files that are duplicates. This occurs particularly when there are source files for applications or virtual disk files. To address this, data deduplication can be used to segment files into small chunks and identify duplicate chunks. Those duplicates are replaced by a pointer to a single stored copy. Those chunks are also compressed to save even more space. All of this is transparent to the user.

Data Deduplication, a role service of File and Storage Services, can be installed using the Add Roles and Features Wizard or by using Windows PowerShell to **Add-WindowsFeature** cmdlet to add the FS-Data-Deduplication feature. Once installed, the feature must be enabled on a data volume by using the Server

Manager dashboard or by using the Windows PowerShell **Enable-DedupVolume** cmdlet. Once enabled, the data deduplication job can be run on demand or scheduled to run at regular intervals.

Demonstration: Configuring Data Deduplication

In this demonstration, you will see how to:

- Add the Data Deduplication role service.
- Enable data deduplication.
- Test data deduplication.

Demonstration Steps

Add the Data Deduplication role service

1. Sign in on LON-SVR2 as **Adatum\Administrator** using the password **Pa$$w0rd**.
2. Open the Server Manager.
3. In the Add Roles and Features Wizard, install the following roles and features to the local server, and accept the default values:
   - **File And Storage Services (Installed)\File and iSCSI Services\Data Deduplication**

Enable Data Deduplication

1. In Server Manager, in the navigation pane, click **File and Storage Services**, and then click **Volumes**.
2. In the Volumes pane, right-click drive **E:**, and then click **Configure Data Deduplication**.
3. Configure Data Deduplication with the following settings:
   - Enable Data Deduplication: **General purpose file server**
   - Deduplicate files older than (in days): **3**
   - Set Deduplication Schedule: **Enable throughput optimization**
   - Start time: **2 A.M.**

Test Data Deduplication

1. On LON-SVR2, copy the **Group Policy Preferences.docx** file from the root folder of the E: drive to the **E:\LabFiles** folder.
2. Verify the file size in both locations.
3. In Windows PowerShell, on LON-SVR2, type the following cmdlet to start the deduplication job in optimization mode:

```
Start-DedupJob –Type Optimization –Volume E:
```

4. When the job completes, verify the size of the **Group Policy Preferences.docx** file on the root folder in drive E:.
5. In the **Properties** dialog box of the **Group Policy Preferences.docx** file, note the values for Size and Size on Disk. The size on the disk should be much smaller than it was previously.

What Is Tiered Storage?

- Tiered storage uses a virtual disk that consists of at least one SSD and one hard disk drive
- Two tiers are created:
  - Faster Tier (all the SSDs)
  - Standard Tier (all the hard disk drives)
- Heavily accessed files are automatically moved to the faster tier
- Parallelized repair allows data to be rebuilt after a disk failure
- The virtual disk may be fixed or thin provisioned
- Thin provisioning and trim is enabled by default

Windows Server 2012 introduced the concept of storage pools, a technology that allows you to virtualize storage by grouping multiple physical disks into pools of storage, and then create virtual disks from that pool. Then you can create volumes on those virtual disks. In Windows Server 2012 R2, the storage pools' abilities have been extended to support tiered storage. This topic provides an overview of tiered storage.

Tiered storage allows you to create a storage pool that consists of traditional hard disks with large capacity, Serial ATA (SATA) or SCSI for example, and Solid State Drive (SSDs) that have lower capacity but better performance. Tiered storage allows you to take advantage of the SSD performance for your heavily accessed files by automatically moving these files to the SSD. Windows Server 2012 R2 runs a daily scheduled task to analyze the data usage and identify files that have the heaviest usage and move them to the SSD. Administrators can also choose to permanently assign a file to the faster tier.

Configuring Tiered Storage

You must first create a storage pool that has at least one SSD and one mechanical disk. Then you use the New Virtual Disk Wizard to create a virtual disk with tiered storage by providing the following information:

1. Select the storage pool you where you want to create the virtual disk.
2. Provide a name for the virtual disk and check the check box to **Create storage tiers on this virtual disk**. Two tiers will be created: one named Faster Tier that contains all of the SSDs, and one named Standard Tier that contains all the remaining hard disk drives.
3. Select the storage layout: **Simple**, or **Mirror**, or **Parity** (Parity requires three or more physical disks).
4. Specify the provisioning type: Thin or Fixed.
5. Specify the size of the virtual disk. You must configure how much of the SSD space (Faster Tier) and how much of the hard disk drive space (Standard Tier) will be used by the virtual disk.
6. You may then select to launch the New Volume Wizard to create volumes on the virtual disk.

Parallelized Repair

Windows Server 2012 R2 supports parallelized repair. A typical Redundant Array of Independent Disks (RAID) array often involves using a hot spare to replace a failed disk. That method can be slow and may require human intervention. With parallelized repair, when a disk in the storage space fails, the remaining healthy disks will rebuild the data that was stored on the failed disk. This provides a much faster recovery and involves no human intervention. It is recommended to add disks that are active in the storage space but contain no data. That way they are available for the parallelized repair process.

Thin Provisioning and Trim

Thin provisioning, also referred to as 'Just-in-Time', allocation of storage space allows you to create a volume that allocates more space than physically exists on the virtual disk. Trim support allows you to claim back, or trim, some of the space that is not being used on the volume. Thin provisioning and trim is enabled by default.

Thin provisioning allows you to set volume sizes that reflect the expected future data growth. For example, you may use thin provisioning to size a volume at 5 terabytes (TB), even though you only have 2

TB of physical space today. Over time, as demand for space increases, you can add disks to the storage pool to actually have 5 TB of disk space without having to worry about increasing the logical size of the volume.

Thin provisioning provides notification when physical storage use thresholds are reached so that more physical disks can be added.

Conversely, if you over-allocate the physical space to a volume, you can trim back the volume to reclaim that space. Windows Server 2012 R2 provides a new application program interface (API) that lets applications return storage when it is no longer needed. NTFS issues trim notifications are issued as part of storage consolidation which is performed on a regular scheduled basis.

Lab A: Implementing Advanced File Services

Scenario

As the A. Datum Corporation has expanded, the requirements for managing storage and shared file access have also expanded. Although the cost of storage has decreased significantly over recent years, the amount of data produced by the A. Datum business groups has increased even faster. The organization is considering alternate ways to decrease the cost of storing data on the network, and is considering options for optimizing data access in the new branch offices. The organization would also like to ensure that data that is stored on the shared folders is limited to company data, and that it does not include unapproved file types.

As a senior server administrator at A. Datum, you are responsible for implementing the new file storage technologies for the organization. You will implement iSCSI storage to provide a less complicated option for deploying large amounts of storage.

Objectives

After completing this lab, the students will be able to:

- Configure iSCSI storage.
- Configure the File Classification Infrastructure.

Lab Setup

Estimated Time: 75 minutes

Virtual machines	20412C-LON-DC1 20412C-LON-SVR1 20412C-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, open **Hyper-V Manager**.
2. In the Hyper-V Manager console, right-click **20412C-LON-SVR2** and then click **Settings**.
3. In the **Settings for 20412C-LON-SVR2** window, in the left pane, ensure that the first network adapter is connected to Private Network, and change the second network adapter to also be connected to Private Network.

 **Note:** You can only perform the previous step if LON-SVR2 has not been started. If LON-SVR2 is already started, shut down LON-SVR2 and then perform these steps.

4. In the Hyper-V Manager, click **20412C-LON-DC1**, and in the Actions pane, click **Start**.
5. In the Actions pane, click **Connect**. Wait until the virtual machine starts.

6. Sign in using the following credentials:
   - User name: **Adatum\Administrator**
   - Password: **Pa$$w0rd**
7. On LON-DC1 right-click the **Start** button and click **Network Connections**.
8. Right-click **Ethernet 2** and click **Enable**.
9. Close the Network Connections dialog box.
10. Repeat steps four through six for 20412C-LON-SVR1 and 20412C-LON-SVR2.

Exercise 1: Configuring iSCSI Storage

Scenario

To decrease the cost and complexity of configuring centralized storage, A. Datum has decided to use iSCSI to provide storage. To get started, you will install and configure the iSCSI target, and then configure access to the target by configuring the iSCSI initiators.

The main tasks for this exercise are as follows:

1. Install the iSCSI target feature
2. Configure the iSCSI targets
3. Configure MPIO
4. Connect to and configure the iSCSI targets

▶ Task 1: Install the iSCSI target feature

1. Sign in on LON-DC1 with the username **Adatum\Administrator** and the password **Pa$$w0rd**.
2. Open the Server Manager.
3. In the Add Roles and Features Wizard, install the following roles and features to the local server, and accept the default values:
   - **File And Storage Services (Installed)\File and iSCSI Services\iSCSI Target Server**

▶ Task 2: Configure the iSCSI targets

1. On LON-DC1, in the Server Manager, in the navigation pane, click **File and Storage Services**, and then click **iSCSI**.
2. Create a virtual disk with the following settings:
   - Storage location: **C:**
   - Disk name: **iSCSIDisk1**
   - Size: **5 GB**
   - iSCSI target: **New**
   - Target name: **lon-dc1**
   - Access servers: **172.16.0.22** and **131.107.0.2**
3. On the **View results** page, wait until the creation completes, and then click **Close**.

4. Create a New iSCSI virtual disk with the following settings:
   - Storage location: **C:**
   - Disk name: **iSCSIDisk2**
   - Size: **5 GB**
   - iSCSI target: **lon-dc1**
5. Create a New iSCSI virtual disk with the following settings:
   - Storage location: **C:**
   - Disk name: **iSCSIDisk3**
   - Size: **5 GB**
   - iSCSI target: **lon-dc1**

▶ Task 3: Configure MPIO

1. Sign in on LON-SVR2 with the username **Adatum\Administrator** and the password **Pa$$w0rd**.
2. On LON-SVR2, from the Server Manager, open the Routing and Remote Access console.
3. Right-click **LON-SVR2**, and then click **Disable Routing and Remote Access**. Close the Routing and Remote Access console.

Note: Normally, you do not disable Routing and Remote Access (RRAS) before configuring MPIO. You do it here because of lab requirements.

4. In the Server Manager, start the Add Roles and Features Wizard and install the **Multipath I/O** feature.
5. In Server Manager, on the **Tools** menu, open **iSCSI Initiator**, and configure the following:
   - Enable the **iSCSI Initiator** service
   - **Quick Connect** to target: **LON-DC1**
6. In Server Manager, on the **Tools** menu, open **MPIO**, and configure the following:
   - Enable **Add support for iSCSI devices** on **Discover Multi-paths**
7. After the computer restarts, sign in on LON-SVR2 with the username **Adatum\Administrator** and password **Pa$$w0rd**.
8. In the Server Manager, on the **Tools** menu, click **MPIO**, and verify that **Device Hardware ID MSFT2005iSCSIBusType_0x9** is added to the list.

▶ Task 4: Connect to and configure the iSCSI targets

1. On LON-SVR2, in the Server Manager, on the **Tools** menu, open **iSCSI Initiator**.
2. In the **iSCSI Initiator Properties** dialog box, perform the following steps:
   - **Disconnect** all **Targets**.
   - **Connect** and **Enable multi-path**.
   - Set **Advanced** options as follows:
     - Local Adapter: **Microsoft iSCSI Initiator**
     - Initiator IP: **172.16.0.22**
     - Target Portal IP: **172.16.0.10 / 3260**

- **Connect** to another target, **enable multi-path**, and configure the following **Advanced** settings:
  - Local Adapter: **Microsoft iSCSI Initiator**
  - Initiator IP: **131.107.0.2**
  - Target Portal IP: **131.107.0.1 / 3260**

3. In the **Targets** list, open **Devices** for **iqn.1991-05.com.microsoft:lon-dc1-lon-dc1-target**, access the MPIO information, and then verify that in **Load balance policy**, **Round Robin** is selected. Verify that two paths are listed by looking at the IP addresses of both network adapters.

Results: After completing this exercise, you will have configured and connected to iSCSI targets.

Exercise 2: Configuring the File Classification Infrastructure

Scenario

A. Datum has noticed that many users are copying corporate documentation to their mapped drives on the users' or departmental file servers. As a result, there are many different versions of the same documents on the network. To ensure that only the latest version of the documentation is available for most users, you need to configure a file classification system that will delete specific files from user folders.

The main tasks for this exercise are as follows:

1. Create a classification property for corporate documentation
2. Create a classification rule for corporate documentation
3. Create a classification rule that applies to a shared folder
4. Create a file management task to expire corporate documents
5. Verify that corporate documents are expired
6. Prepare for the next lab

▶ Task 1: Create a classification property for corporate documentation

1. On LON-SVR1, from the Server Manager, start the File Server Resource Manager.
2. In File Server Resource Manager, under Classification Management, create a local property with the following settings:
   - Name: **Corporate Documentation**
   - Property Type: **Yes/No**
3. Leave the File Server Resource Manager console open.

▶ Task 2: Create a classification rule for corporate documentation

1. In the File Server Resource Manager console, create a classification rule with following settings:
   - **General** tab, Rule name: **Corporate Documents Rule**, and ensure that the rule is enabled.
   - **Scope** tab: **E:\Labfiles\Corporate Documentation** folder
   - **Classification** tab:
     - Classification method: **Folder Classifier**
     - Property, Choose a property to assign to files: **Corporate Documentation**

- Property, Specify a value: **Yes**
- Evaluation type tab: **Re-evaluate existing property values** and **Aggregate the values**

2. Select both **Run the classification with all rules** and **Wait for classification to complete**.
3. Review the Automatic Classification report that displays in Internet Explorer, and ensure that the report lists the same number of classified files as in the Corporate Documentation folder.
4. Close Internet Explorer, but leave the File Server Resource Manager console open.

▶ Task 3: Create a classification rule that applies to a shared folder

1. In the File Server Resource Manager console, create a local property with following settings:
   - Name: **Expiration Date**
   - Property Type: **Date-Time**
2. In the File Server Resource Manager console, create a classification rule with the following settings:
   - General tab, Rule name: **Expiration Rule**, and ensure that the rule is enabled
   - Scope tab: **E:\Labfiles\Corporate Documentation**
   - Classification tab, Classification method: **Folder Classifier**
   - Property, Choose a property to assign to files: **Expiration Date**
   - Evaluation type tab: **Re-evaluate existing property values** and **Aggregate the values**
3. Select both **Run the classification with all rules** and **Wait for classification to complete**.
4. Review the Automatic classification report that appears in Internet Explorer, and ensure that report lists the same number of classified files as the Corporate Documentation folder.

   Close Internet Explorer, but leave the File Server Resource Manager console open.

▶ Task 4: Create a file management task to expire corporate documents

- In the File Server Resource Manager, create a file management task with following settings:
  - General tab: Task name: **Expired Corporate Documents**, and ensure that the task is enabled
  - Scope tab: **E:\Labfiles\Corporate Documentation**
  - Action tab, Type: **File expiration** is selected,
  - Expiration directory: **E:\Labfiles\Expired**
  - Notification tab: **Event Log** and **Send warning to event log**
  - Condition tab: Days since the file was last modified: **1**

Note: This value is for lab purposes only. In a real scenario, the value would be 365 days or more, depending on each company's policy

  - Schedule tab: Weekly and Sunday
  - Leave the File Server Resource Manager console open.

▶ Task 5: Verify that corporate documents are expired

1. In the File Server Resource Manager, click **Run File Management Task Now**, and then click **Wait for the task to complete**.
2. Review the file management task report that displays in Internet Explorer, and ensure that the report lists the same number of classified files as the Corporate Documentation folder.
3. Start the Event Viewer, and in the Event Viewer console, open the **Application** event log.
4. Review events with numbers **908** and **909**. Notice that 908 – FSRM started a file management job, and 909 – FSRM finished a file management job.
5. Close open Windows.

▶ Task 6: Prepare for the next lab

When you finish the lab, revert 20412C-LON-SVR1. To do this, complete the following steps.

1. On the host computer, start Hyper-V Manager.
2. On the **Virtual Machines** list, right-click **20412C-LON-SVR1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.

Keep all other virtual machines running for the next lab.

Results: After completing this exercise, you will have configured a File Classification Infrastructure so that the latest version of the documentation is always available to users.

Question: Why would you implement MPIO together with iSCSI? What problems would you solve with this approach?

Question: Why must you have the iSCSI initiator component?

Question: Why would you configure file classification for documents located in a folder such as a Corporate Documentation folder?

Lab B: Implementing BranchCache

Scenario

The A. Datum Corporation has deployed a new branch office, which has a single server. To optimize file access in branch offices, you must configure BranchCache. To reduce WAN use out to the branch office, you must configure BranchCache to cache data retrieved from the head office. You will also implement FSRM to assist in optimizing file storage.

Objectives

After completing this lab, the students will be able to:

- Configure the main office servers for BranchCache.
- Configure the branch office servers for BranchCache.
- Configure client computers for BranchCache.
- Monitor BranchCache.

Lab Setup

Estimated Time: 40 Minutes

Virtual machines	20412C-LON-DC1 20412C-LON-SVR2 20412C-LON-CL1 20412C-LON-CL2
User name	**Adatum\Administrator**
Password	**Pa$$w0rd**

For this lab, you will continue to use the 20412C-LON-DC1 and 20412C-LON-SVR2 virtual machines that are still running from the last lab. You will also use 20412C-LON-CL1 and 20412-LON-CL2, but do not start 20412C-LON-CL1 and 20412C-LON-CL2 until instructed to do so in the lab steps.

Exercise 1: Configuring the Main Office Servers for BranchCache

Scenario

Before you can configure the BranchCache feature for your branch offices, you must configure the network components.

The main tasks for this exercise are as follows:

1. Configure LON-DC1 to use Windows BranchCache
2. Simulate a slow link to the branch office
3. Enable a file share for BranchCache
4. Configure client firewall rules for BranchCache

▶ Task 1: Configure LON-DC1 to use Windows BranchCache

1. Switch to LON-DC1.
2. Open the Server Manager, and install the **BranchCache for network files** role service.

3. Open the Local Group Policy Editor (**gpedit.msc**).
4. Navigate to and open **Computer Configuration/Administrative Templates/Network/Lanman Server/Hash Publication for BranchCache**.
5. Enable the BranchCache setting, and then select **Allow hash publication only for shared folders on which BranchCache is enabled**.

▶ Task 2: Simulate a slow link to the branch office

1. In the Local Group Policy Editor console, navigate to **Computer Configuration\Windows Settings\Policy-based QoS**.
2. Create a new policy with the following settings:
   - Name: **Limit to 100 Kbps**
   - Specify Outbound Throttle Rate: **100**
   - Accept default values on other wizard pages

▶ Task 3: Enable a file share for BranchCache

1. On LON-DC1, in the File Explorer window, create a new folder named **C:\Share**.
2. Share this folder with the following properties:
   - Share name: **Share**
   - Permissions: **default**
   - Caching: **Enable BranchCache**
3. Copy **C:\Windows\System32\mspaint.exe** to the **C:\Share** folder.
4. Close all the command prompt and File Explorer.

▶ Task 4: Configure client firewall rules for BranchCache

1. On LON-DC1, open the Group Policy Management console.
2. Navigate to **Forest: Adatum.com\Domains\Adatum.com\Default Domain Policy**, and then open the policy for editing.
3. Navigate to **Computer Configuration\Policies\Windows Settings\Security Settings\Windows Firewall with Advanced Security\Windows Firewall with Advanced Security\Inbound Rules**.
4. Create a new inbound firewall rule with the following properties:
   - Rule type: **predefined**
   - Use **BranchCache – Content Retrieval (Uses HTTP)**
   - Action: **Allow**
5. Create a new inbound firewall rule with the following properties:
   - Rule type: **predefined**
   - Use **BranchCache – Peer Discovery (Uses WSD)**
   - Action: **Allow**

6. Close the Group Policy Management Editor and Group Policy Management console.
7. Update the group policy.

Results: At the end of this exercise, you will have deployed BranchCache, configured a slow link, and enabled BranchCache on a file share.

Exercise 2: Configuring the Branch Office Servers for BranchCache

Scenario

The next step you must perform is to configure a file server for the BranchCache feature. You will install and start the BranchCache feature.

The main tasks for this exercise are as follows:

1. Install the BranchCache feature on LON-SVR2
2. Start the BranchCache host server

▶ Task 1: Install the BranchCache feature on LON-SVR2

- On LON-SVR2 from Server Manager, add the **BranchCache for Network Files** role service and the **BranchCache** feature.

▶ Task 2: Start the BranchCache host server

1. On LON-SVR2, open Windows PowerShell, and run the following cmdlet:

```
Enable-BCHostedServer
-RegisterSCP
```

2. On LON-SVR1, in Windows PowerShell, run the following cmdlet:

```
Get-BCStatus
```

3. Ensure that Branch Cache is enabled and running. Note in the DataCache section, the current active cache size is zero.
4. Update group policy

Results: At the end of this exercise, you will have enabled the BranchCache server in the branch office.

Exercise 3: Configuring Client Computers for BranchCache

Scenario

After configuring the network components, you must ensure that the client computers are configured correctly. This is a preparatory task for using BranchCache.

The main tasks for this exercise are as follows:

- Configure client computers to use BranchCache in hosted cache mode

▶ Task 1: Configure client computers to use BranchCache in hosted cache mode

1. On LON-DC1, open the Server Manager, and then open Group Policy Management.
2. Create a new OU named **Branch**.

3. Create and link a new GPO named **BranchCache** to Branch OU.
4. Edit the **BranchCache** GPO.
5. In the Group Policy Management Editor, browse to **Computer Configuration\Policies\Administrative Templates\Network\BranchCache**, and configure the following:
   - Turn on BranchCache: **Enabled**
   - Enable Automatic Hosted Cache Discovery by Service Connection Point: **Enabled**
   - Configure BranchCache for network files**: Enabled**
   - Type the maximum round trip network latency (milliseconds) after which caching begins: **0**
6. Open the Server Manager, and then open **Active Directory Users and Computers**.
7. Move **LON-CL1** and **LON-CL2** from the Computers container to the Branch OU.
8. Start **20412C-LON-CL1**, sign in as Administrator.
9. At the command prompt, type **netsh branchcache show status all**, and then press Enter.
10. Verify that the BranchCache Status is **Running**. If the status is **Stopped**, restart the client computer.
11. Start the **20412C-LON-CL2**, sign in as Adatum\Administrator and open the command prompt window.
12. At the command prompt, type **netsh branchcache show status all**, and then press Enter.
13. Verify that the BranchCache status is **Running**. If the status is **Stopped**, restart the client computer.

Results: At the end of this exercise, you will have configured the client computers for BranchCache.

Exercise 4: Monitoring BranchCache

Scenario

Lastly, you must test and verify that the BranchCache feature is working as expected.

The main tasks for this exercise are as follows:

1. Configure Performance Monitor on LON-SVR1
2. Configure performance statistics on LON-CL1
3. Configure performance statistics on LON-CL2
4. Test BranchCache in the hosted cache mode
5. Prepare for the next module

▶ Task 1: Configure Performance Monitor on LON-SVR1

1. On LON-SVR2, open the Performance Monitor.
2. In the Performance Monitor console, in the navigation pane, under **Monitoring Tools**, click **Performance Monitor**.
3. Remove existing counters, change to report view, and then add the **BranchCache** object counters to the report.

▶ Task 2: Configure performance statistics on LON-CL1

1. Switch to LON-CL1, and open the Performance Monitor.
2. In the navigation pane of the Performance Monitor console, under **Monitoring Tools**, click **Performance Monitor**.
3. In Performance Monitor, remove existing counters, change to a report view, and then add the **BranchCache** object to the report.

▶ Task 3: Configure performance statistics on LON-CL2

1. Switch to LON-CL2, and open the Performance Monitor.
2. In the Performance Monitor console, in the navigation pane, under **Monitoring Tools**, click **Performance Monitor**.
3. In the Performance Monitor, remove existing counters, change to a report view, and then add the **BranchCache** object to the report.

▶ Task 4: Test BranchCache in the hosted cache mode

1. Switch to LON-CL1.
2. Open **\\LON-DC1.adatum.com\share**, and copy **mspaint.exe** to the local desktop. This could take several minutes because of the simulated slow link.
3. Read the performance statistics on LON-CL1. This file was retrieved from LON-DC1 (Retrieval: Bytes from Server). After the file was cached locally, it was passed up to the hosted cache. (Retrieval: Bytes Served).
4. Switch to LON-CL2.
5. Open **\\LON-DC1.adatum.com\share**, and copy **mspaint.exe** to the local desktop. This should not take as long as the first file copy, because the file is cached.
6. Read the performance statistics on LON-CL2. This file was obtained from the hosted cache (Retrieval: Bytes from Cache).
7. Read the performance statistics on LON-SVR2. This server has offered cached data to clients (Hosted Cache: Client file segment offers made).
8. On LON-SVR2, in Windows PowerShell, run the following cmdlet:

```
Get-BCStatus
```

 **Note:** In the DataCache section, the current active cache size is no longer zero, it is 6560896.

▶ Task 5: Prepare for the next module

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

1. On the host computer, start **Hyper-V Manager**.
2. On the **Virtual Machines** list, right-click **20412C-LON-DC1**, and then click **Revert**.

3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps two and three for **20412C-LON-SVR2**, **20412C-LON-CL1**, and **20412C-LON-CL2**.

Results: At the end of this exercise, you will have verified that BranchCache is working as expected.

Question: When would you consider implementing BranchCache into your own organization?

Module Review and Takeaways

Review Questions

Question: How does BranchCache differ from the Distributed File System (DFS) to branch offices?

Question: Why would you choose to implement BranchCache in hosted cache mode instead of distributed cache mode?

Question: Can you configure Data Deduplication on a boot volume?

Question: Why would you implement a File Classification Infrastructure?

Real-world Issues and Scenarios

Your organization is considering deploying an iSCSI solution. You are a Windows Server 2012 administrator who is responsible for designing and deploying the new solution. This new solution will be used by different types of technologies, such as Windows Server 2012 file server, Exchange Server, and SQL Server. You face that of designing an optimal iSCSI solution, but at the same time you are not sure whether the solution you are going to propose to your organization will meet the requirements of all technologies that will be accessing the iSCSI storage. What should you do?

Answer: You should include on the team that will design and deploy the iSCSI solution experts from different areas of specialization. Team members who will be involved in the project should include Windows Server 2012 administrators, network administrators, storage administrators, and security administrators. This is necessary so that the iSCSI storage solution has optimal performance and security, and has consistent management and operations procedures.

Your organization is considering deploying a BranchCache solution. You are a Windows Server 2012 administrator in your organization, and you are responsible for designing and deploying the new solution. The organization's business managers are concerned about security of the data that will be stored in the branch offices. They are also concerned about how the organization will address security risks such as data tampering, information disclosure, and denial of service attacks. What should you do?

Answer: You should include a security expert on your design team. You should also consider the defense-in-depth model of analyzing security risks. BranchCache addresses the security risks as follows:

- Data tampering. The BranchCache technology uses hashes to confirm that during the communication, the client computer and the server did not alter the data.
- Information disclosure. BranchCache sends encrypted content to clients, but they must have the encryption key to decrypt the content. Because potential malicious users would not have the encryption key, if an attacker attempts to monitor the network traffic to access the data while it is in transit between clients, the attempt will not be successful.
- Denial of service. If an attacker tries to overload the client with requests for data, BranchCache technology includes queue management counters and timers to prevent clients from being overloaded.

Your organization is using large amounts of disk space for data storage and faces the challenge of organizing and managing the data. Furthermore, your organization must satisfy requirements for security, compliance, and data leakage prevention for company confidential information. What should you do?

Answer: You should deploy the File Classification Infrastructure. Based on file classification, you can configure file management tasks that will enable you to manage groups of files based on various file and folder attributes. You can also automate file and folder maintenance tasks, such as cleaning up stale data or protecting sensitive information.

Tools

Tool	Use	Where to find it
iSCSI target server	Configure iSCSI targets	In Server Manager, under File and Storage Servers
iSCSI initiator	Configure a client to connect to an iSCSI target virtual disk	In Server Manager, in the Tools drop-down list box
Deduplication Evaluation tool (DDPEval.exe)	Analyze a volume to find out the potential savings when enabling Data deduplication	Available from the command prompt and stored in C:\windows\system32
File Server Resource Manager	A set of features that allow you to manage and classify data that is stored on file servers	In Server Manager, in the Tools drop-down list box

Best Practice:

- When you consider an iSCSI storage solution for your organization, spend most of the time on the design process. The design process is crucial because it allows you to optimize the solution for all technologies that will be using iSCSI storage, such as file services, Exchange Server, and SQL Server. The design should also accommodate future growth of your organization's business data. Successful design processes help guarantee a successful deployment of the solution that will meet your organization's business requirements.
- When you plan for BranchCache deployment, ensure that you work closely with your network administrators so that you can optimize network traffic across the WAN.
- When you plan for file classifications, ensure that you start with your organization's business requirements. Identify the classifications that you will apply to documents, and then define a method that you will use to identify documents for classification. Before you deploy the File Classification Infrastructure, create a test environment. Then test the scenarios to ensure that your solution will result in a successful deployment, and that your organization's business requirements will be met.

Module 3

Implementing Dynamic Access Control

Contents:

Module Overview

The Windows Server® 2012 operating system introduces new features for enhancing access control for file-based and folder-based resources, as well as features for accessing your work data from various locations. These features, named Dynamic Access Control (DAC) and Work Folders, extend traditional access control.

DAC enables administrators to use claims, resource properties, policies, and conditional expressions to manage access. Work Folders, specific to Windows Server® 2012 R2, enable users to have more flexible data access. In this module, you will learn about DAC and Work Folders, and how to plan and implement these technologies.

Objectives

After completing this module, you will be able to:

- Describe DAC.
- Implement DAC components.
- Implement DAC for access control.
- Implement access-denied assistance.
- Implement and manage Work Folders.

Lesson 1
Overview of DAC

DAC is a new Windows Server 2012 feature that you can use to enable more functional and flexible access management. DAC offers a new way to secure and control access to resources. Before you implement this feature, you should understand how it works and the components it uses. This lesson presents an overview of DAC.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the limitations of current access management methods.
- Describe DAC.
- Describe claims.
- Describe Resource Properties.
- Describe how access resources with DAC.
- Describe the requirements for DAC implementation.

Limitations of Current Access Management Methods

In previous versions of Windows Server, such as Windows Server 2008, the basic mechanism for file and folder access control was the configuration of the NTFS file system permissions. When you use the NTFS file system permissions and their access control lists (ACLs), administrators can control access to resources based on the user security identifiers (SIDs) or group membership SIDs, and the level of access, such as Read-only, Change, and Full Control.

- NTFS file system permissions and ACLs provide access control that is based on a user's SID or group membership SID
- AD RMS provides greater protection for documents by controlling how applications use them, and also works with user or group SID
- NTFS file system permissions cannot use AND between conditions
- In NTFS file system permissions, you cannot build your own conditions for access control

When you implement Active Directory Rights Management Services (AD RMS), you can establish an additional level of file access control. Unlike NTFS file system permissions, which are not application-aware, AD RMS sets a policy that can control access inside the application that the user uses to open a document. When you implement AD RMS, you enable users to protect documents within applications. For example, you can let someone read a file but prevent that person from copying, printing, or saving the file. However, this technology also relies only on user or group SIDs to manage access control.

In previous Windows versions, it was not possible to set conditional expression-based access control to files by using NTFS file system or AD RMS. For example, NTFS file system permissions cannot be set in a way that a user can access documents only if he or she is a member of two security groups at the same time. Or, you cannot set NTFS file system permissions to allow users access if their **employeeType** AD DS attribute is set to a Full Time Employee (FTE) value. In addition, there are other limitations. For example, you cannot set permissions so that only users who have an AD DS department attribute populated with the same value as the department attribute for the file can access the content.

These limitations can be generalized in the following way: The NTFS file system-based approach for access management does not allow you to use conditional expressions as a way to manage access, and you

cannot use AND between access conditions. In real life, this means you cannot build your own conditions for access control, and you cannot set two different conditions to apply at the same time.

In Windows Server 2012, DAC technology solves these issues. You can use DAC to take into account Active Directory Domain Services (AD DS) attribute values of users or resource objects when providing or denying access.

What Is DAC?

DAC in Windows Server 2012 is a new access control mechanism for file system resources. It enables administrators to define central file access policies that can apply to every file server in an organization. DAC implements a safety net over file servers and any existing Share and NTFS file system permissions. It also ensures that regardless of how the Share and NTFS file system permissions might change, this central overriding policy is still enforced.

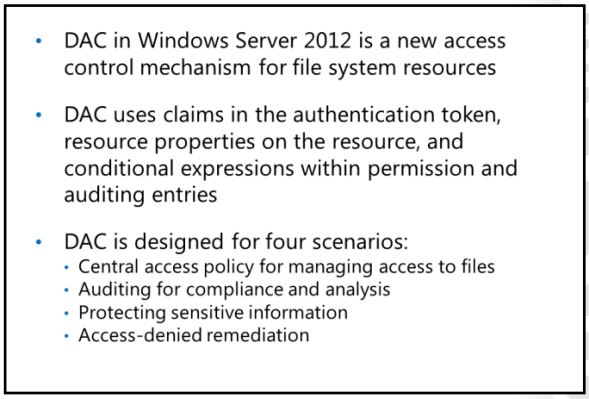

DAC combines multiple criteria into access decisions. This augments the NTFS file system ACL so that users need to satisfy Share permissions, NTFS file system ACL, and the central access policy to gain access to a file. However, DAC also can work independently from NTFS file system permissions.

DAC provides a flexible way to apply, manage, and audit access to domain-based file servers. DAC uses claims in the authentication token, the Resource Properties on the resource, and the conditional expressions within permission and the auditing entries. With this combination of features, you now can grant and audit access to files and folders based on AD DS attributes.

DAC is primarily used to control file access in a much more flexible way than NTFS file system and Share permissions. It also can be used for auditing file access and can provide optional AD RMS protection integration.

DAC is designed for four scenarios:

- Central access policy for managing access to files. This enables organizations to set safety net policies that reflect business and regulatory compliance.
- Auditing for compliance and analysis. This enables targeted auditing across file servers for compliance reporting and forensic analysis.
- Protecting sensitive information. DAC identifies and protects sensitive information within the Windows Server 2012 environment and, if integrated with AD RMS, it protects the information when it leaves the Windows Server 2012 environment.
- Access-denied remediation. This improves the access-denied experience to reduce help desk load and the incident time for troubleshooting. This technology puts control to the files closer to the people who are responsible for those files. Access-denied remediation can inform a different owner for each folder, with additional information why the access was denied, to allow him to make an educated decision how this should be fixed.

What Is the Claim?

Previous Windows Server versions such as Windows Server 2008 used *claims* in Active Directory Federation Services (AD FS). In this context, claims are statements made about users, which may include name, identity, key, group, privilege, or capability. These claim statements are understood by the partners in an AD FS federation. AD FS also provides AD DS–based claims and the ability to convert the data from these claims into Security Assertions Markup Language format. Windows Server 2012 now allows you to read and use any attribute directly from AD DS.

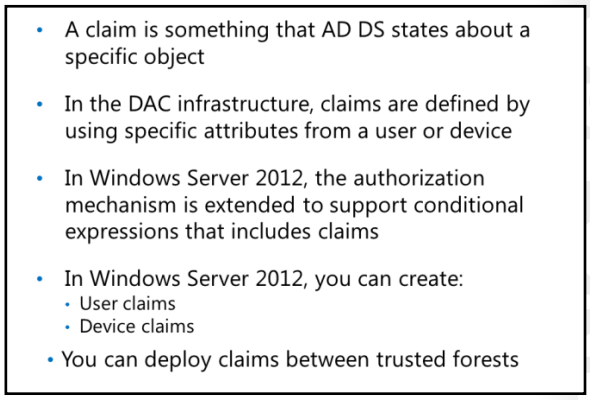

By definition, a claim is something that AD DS states about a specific object, usually a user or a computer. A claim provides information from the trusted source about an entity. Some examples of claims include a user's department, a user's security clearance, and the health state of a computer that is used in Network Access Protection. All these claims state something about a specific object in context of DAC; it is always about the user or device. When you configure resource access, you can use any combination of claims to authorize access to the resources.

In a DAC infrastructure, claims are defined by using specific AD DS attributes of a user or a computer. By defining a claim, you actually tell AD DS which attributes can be used in conditional expressions that DAC uses. You cannot define any conditional expressions or access rules when using DAC until you define at least one claim.

In Windows Server 2012, the authorization mechanism is extended to support conditional expressions. You now can use user and device claims for file and folder authorization, in addition to NTFS file system permissions that are based on a user's SID or group SIDs. By using claims, you now can base your access control decision on SID and other attribute values. This means that Windows Server 2012 still supports using group membership for authorization decisions.

User Claims

A user claim is information that is provided by a Windows Server 2012 domain controller about a user. Windows Server 2012 domain controllers can use most AD DS user attributes as claim information. This provides administrators with a wide range of possibilities for configuring and using claims for access control. Before defining a user claim, you should populate the user attributes that you want to use for access control with appropriate values

Device Claims

A device claim, which is often called a *computer claim*, is information that is provided by a Windows Server 2012 domain controller about a device that is represented by a computer account in AD DS. As with user claims, device claims can use most of the AD DS attributes that are applicable to computer objects. Unlike NTFS file system permissions, DAC also can take into account the device that a user is using when trying to access a resource. Device claims are used to represent device attributes that you want to use for access control.

Claims Across Forests

AD DS maintains a claims dictionary in each forest in Windows Server 2012. All claim types that are defined and used within the forest are defined at the AD DS forest level. However, there are scenarios where a user or device security principal may need to traverse a trust boundary to access resources in a trusted forest. Cross-forest claims transformation in Windows Server 2012 enables you to transform

incoming and outgoing claims that traverse forests so that the claims are recognized and accepted in the trusting and trusted forests. By default, a trusted forest allows all outgoing claims to pass, and a trusting forest drops all incoming claims that it receives

What Is Resource Property?

- Resource properties define attributes of the resource that you want to use
- Resource properties are grouped in resource property lists
- When creating a resource property, you can specify the property type and the allowed or suggested values

When you use claims or security groups to control access to files and folders, you also have the ability to provide additional information for those resources. The information you provide about accessing the resource can be used in DAC rules for access management.

Similar to user or device claims, you have to define the attributes of the resource that you want to use. You do this by configuring the *resource properties*, also known as resource property objects. These objects define additional properties (or attributes) that can be assigned to files and folders. Usually these properties are assigned during file classification tasks. Windows Server 2012 can use these properties for authorization purposes.

For example, these properties can represent the value of a classification for a file or folder, such as Confidential or Internal Use. Other properties can represent the value of a file, such as which department owns the information, or to which project the file is associated, such as, Research, Project X or similar.

You manage resource properties in the resource properties container, which is displayed in the DAC node in the Active Directory Administrative Center.

You can create your own resource properties, or you can use one of the preconfigured properties, such as **Project**, **Department**, and **Folder Usage**. All predefined resource property objects are disabled by default, so if you want to use any of them, you should enable them first. If you want to create your own resource property object, you can specify the property type and the allowed or suggested values of the object.

When you create resource property objects, you can select the properties to include in the files and folders. When evaluating file authorization and auditing, the Windows operating system uses the values in these properties along with the values from user and device claims.

Accessing Resources with DAC

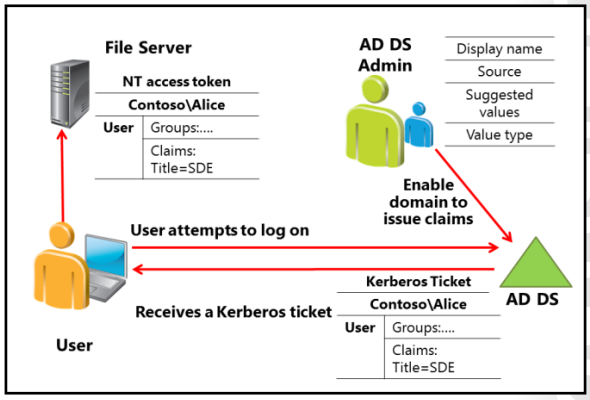

DAC is the new authorization and auditing mechanism that brings required extensions to AD DS. These extensions build the Windows claim dictionary, which is where Windows operating systems store claims for an Active Directory forest. Claims authorization also relies on the Kerberos version 5 protocol Key Distribution Center (KDC).

When NTFS file system is used to manage access control, the user's access token contains the user's SID and the SIDs of all groups that have that user as a member. When the user tries to access the

resource, the ACL of that resource is evaluated. If at least one SID from the user's token is matched to the SID on the ACL, the appropriate rights are assigned to the users.

DAC, however, does not just use SIDs to manage access. Claims also are used to define some of the additional properties that a user or device can have. This means that a user's access token should not only have information about SIDs, it also should have information about the user's claims and information about claims from the device that the user is using to access the resource.

The Windows Server 2012 KDC contains Kerberos protocol enhancements that are required to transport the claims within a Kerberos ticket and to use compound identity. Windows Server 2012 KDC also includes an enhancement to support *Kerberos armoring*. Kerberos armoring is an implementation of Flexible Authentication Secure Tunneling, which provides a protected channel between the Kerberos client and the KDC. Claims are stored in Kerberos Privilege Account Certificate, and they don't increase classical token size.

After you have configured user and device claims and Resource Properties, you then must protect files and folders by using conditional expressions. Conditional expressions evaluate user and device claims against constant values or values within Resource Properties. You can do this in the following three ways:

- If you want to include only specific folders, you can use the Advanced Security Settings Editor to create conditional expressions directly in the security descriptor.
- If you want to include some or all file servers, you can create Central Access Rules, and then link those rules to the central access policy objects. You then can use Group Policy to apply the central access policy objects to the file servers, and then configure the share to use the central access policy object. Using these Central Access Policies is the most efficient and preferred method for securing files and folders. This is discussed further in the next topic.
- When managing access with DAC, you can use file classifications to include certain files with a common set of properties across various folders or files.

Windows Server 2012 and Windows 8 support one or more conditional expressions within a permission entry. Conditional expressions simply add another applicable layer to the permission entry. The results of all conditional expressions must evaluate to TRUE for a Windows operating system to grant the permission entry for authorization. For example, suppose that you define a claim named **Department**, with a source attribute department, for a user and that you define a Resource Property object named **Department**. You now can define a conditional expression that says that the user can access a folder, with the applied Resource Property objects, only if the user's attribute **Department** value is equal to the value of property **Department** on the folder. Note that if the **Department** Resource Property object has not been applied to the file or files in question, or if **Department** is a null value, then the user will be granted access to the data.

Requirements for DAC Implementation

Your servers must meet certain prerequisites before implementing DAC. Claims-based authentication requires the following infrastructure:

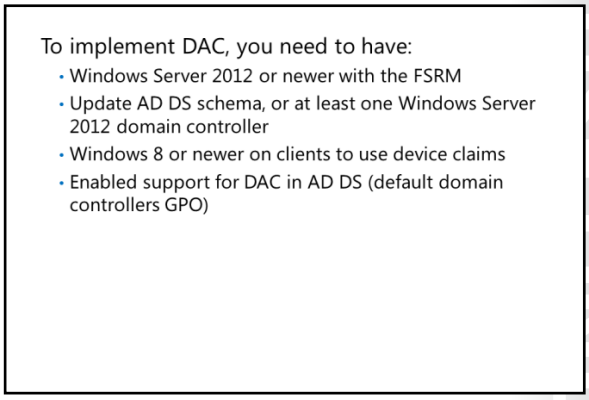

- Windows Server 2012 or newer with the File Server Resource Manager (FSRM) role service enabled. This must be installed on the file server that will host the resources that DAC will protect. The file server that will host the share must be a Windows Server 2012 file server so that it can do the following: read the claims and device authorization data from a Kerberos version 5 ticket, translate the SIDs and the claims from the ticket into an authentication token, and then compare the authorization data in the token against the conditional expressions in the security descriptor.
- At least one Windows Server 2012 domain controller to store the central definitions for the resource properties and policies. User claims are not required for security groups. If you use the user claims, then at least one Windows Server 2012 domain controller in the user domain should be accessible by the file server so that the file server can retrieve the claims on behalf of the user. If you use device claims, then all the client computers in the AD DS domain need to use the Windows 8 operating system.

Note: Only Windows 8 or newer devices use device claims.

Prerequisites for using claims are as follows:

- If you use claims across a forest trust, you must have the Windows Server 2012 domain controllers in each domain, exclusively.
- If you use device claims, then you must have a Windows 8 client. Earlier Windows operating systems do not support device claims.

Although a Windows Server 2012 domain controller is required when using user claims, there is no requirement for having a Windows Server 2012 domain and a forest functional level, unless you want to use the claims across a forest trust. This means that you also can have domain controllers that run Windows Server 2008 and Windows Server 2008 R2 with the forest functional level set to Windows Server 2008. However, if you want to always provide claims to users and devices, by configuring that in Group Policy, you should raise your domain and forest functional level to Windows Server 2012. This is discussed in following paragraph.

Enabling Support for DAC in AD DS

After fulfilling software requirements for enabling DAC support, you must enable claim support for the Windows Server 2012 KDC. Kerberos protocol support for DAC provides a mechanism for including user claim and device authorization information in a Windows authentication token. Access checks performed on resources such as files or folders use this authorization information to verify identity.

You should first use a Group Policy to enable AD DS for DAC. This setting is specific to domain controllers, so you can create a new Group Policy Object (GPO) and then link the setting to the domain controllers' organizational unit (OU), or by editing the Default Domain Controllers GPO that is already linked to that OU.

Whichever method you choose, you should open the **Group Policy Object Editor**, expand **Computer Configuration**, expand **Policies**, expand **Administrative Templates**, expand **System**, and then expand **KDC**. In this node, open a setting called **Support Dynamic Access Control and Kerberos Armoring**.

To configure the **Support Dynamic Access Control and Kerberos Armoring** policy setting, you can choose one of the four listed options:

1. **Do not support Dynamic Access Control and Kerberos Armoring**.
2. **Support Dynamic Access Control and Kerberos Armoring**.
3. **Always provide claims and FAST RFC behavior**.
4. **Also fail unarmored authentication requests**.

Claims and Kerberos armoring support are disabled by default, which is equivalent to the policy setting of not being configured or being configured as **Do not support Dynamic Access Control and Kerberos Armoring**.

The **Support Dynamic Access Control and Kerberos Armoring** policy setting configures DAC and Kerberos armoring in a mix-mode environment, when there is a mixture of Windows Server 2012 domain controllers and domain controllers running older versions of the Windows Server operating system.

The remaining policy settings are used when all the domain controllers are Windows Server 2012 domain controllers and the domain functional level is configured to Windows Server 2012. The **Always provide claims and FAST RFC behavior** and the **Also fail unarmored authentication requests** policy settings enable DAC and Kerberos armoring for the domain. However, the later policy setting requires all Kerberos authentication service and ticket granting service communication to use Kerberos armoring. Windows Server 2012 domain controllers read this configuration while other domain controllers ignore this setting.

 **Note:** Implementing DAC in an environment with multiple forests has additional setup requirements.

Lesson 2
Implementing DAC Components

Before you put DAC into production, you have to configure several components. First, you need to define the claims and the resource properties, and then you have to build access control rules. By using access control rules, you build access policies that are applied on the file servers. By using the classification mechanism, you can make DAC usage even more efficient. In this lesson, you will learn how to configure and implement the building blocks for DAC.

Lesson Objectives

After completing this lesson, you will be able to:

- Create and manage claims.
- Create and manage resource properties and resource property lists.
- Create and manage access control rules.
- Create and manager access policies.
- Configure claims, resource properties, and rules.
- Implement and manage file classifications.
- Configure classification rules.

Creating and Managing Claims

To create and configure claims, you primarily use the Active Directory Administrative Center. You use the Active Directory Administrative Center to create attribute-based claims, which are the most common type of claim. However, you also can use the Active Directory module for Windows PowerShell® to create certificate-based claims. All claims are stored within the configuration partition in AD DS. Because this is a forest-wide partition, all domains within the forest share the claim dictionary, and the domain controllers from the domain issue claims information during user and computer authentication.

- Use the AD CS to create attribute-based claims
- Use the Active Directory module for Windows PowerShell to create certificate-based claims
- Claims are stored within the configuration partition in AD DS
- Attributes are used to source values for claims
- Make sure that you configure attributes for your computer and user accounts in AD DS with the information that is correct for the respective user or computer

To create attribute-based claims in the Active Directory Administrative Center, navigate to the DAC node, and then open the Claim Types container. By default, no claim types are defined here. In the Actions pane, you can click Create Claim Type to view the list of attributes. These attributes are used to source values for claims. When you create a claim, you associate the claim with the specific attribute. The value of that attribute is populated as a claim value. Therefore, it is crucial that the information contained within the AD DS attributes that are used to source claim types contain accurate information or remain blank, as this will be used for security access control. You can select each claim type if it applies to user object, computer object or both.

When you select the attribute that you want to use to create a claim, you also must provide a name for the claim. The suggested name for the claim is always the same as the selected attribute name. However, you also can provide an alternate or more meaningful name for the claim. Optionally, you also can provide suggested values for a claim. This is not mandatory, but we recommend it because it can reduce

the possibility of making mistakes, when you create conditional expressions that contain claim values. By default, no suggested values are provided.

You also can specify the claim identification (ID). This value is generated automatically, but you might want to specify the claim ID if you define the same claim for multiple forests and want the ID to be identical.

Note: Claim types are sourced from AD DS attributes. For this reason, you must configure attributes for your computer and the user accounts in AD DS with information that is correct for the respective user or computer. Windows Server 2012 domain controllers do not issue a claim for an attribute-based claim type when the attribute for the authenticating principal is empty. Depending on the configuration of the data file's Resource Property object attributes, a null value in a claim might result in the user being denied access to DAC-protected data.

Creating and Managing Resource Properties and Resource Property Lists

Besides defining claims, you also must define *resource properties* to create efficient conditional expressions. Resource properties describe resources that you protect with DAC, and they help you to better define the scope for DAC implementation. Also, they can help with file classification. You create and manage resource properties by using the Active Directory Administrative Center.

- Resource properties describe resources that you protect with DAC
- Several resource properties are already predefined in Windows Server 2012
- All predefined resource properties are disabled
- When creating a new resource property, you have to set its name, and value type
- In Windows Server 2012 R2, you also can create reference resource properties
- Resource properties are grouped in resource property lists

Unlike claim types, which are not defined by default, several resource properties are already included in Windows Server 2012. For example, predefined resource properties include **Company**, **Confidentiality**, **Folder Usage**, **Department**, and so on. All predefined resource properties are disabled, and if you want to use any of them in conditional expressions or in file classifications, you have to enable them first.

If you do not want to use predefined resource properties, you can define your own. When creating a new Resource Property, you have to set up its name and select the value type.

Resource Property value types can be :

- Date/Time
- Multi-valued Choice
- Multi-valued Text
- Number
- Ordered List
- Single-valued Choice
- Text
- Yes/No

Similar to claims, you can set the ID for a resource property to be used in a trusted forest.

While suggested values are not mandatory for claims, you must provide at least one suggested value for each Resource Property you define.

In Windows Server 2012 R2, you also can create reference resource properties. A reference resource property is a resource property that uses an existing claim type that you created before for its suggested value. If you want to have claims and resource properties with the same suggested values, then you should use the reference resource properties.

Note: Access is controlled not by the claim, but by the resource property object. The claim must provide the correct value corresponding to the requirements set by the resource property object. If the resource property object does not involve a particular attribute, then additional or extra claim attributes associated with the user or device are ignored.

Resource properties are grouped in resource property lists. A global resource property list is predefined, and it contains all resource properties that applications can use. You also can create your own resource property lists if you want to group some specific resource properties.

Creating and Managing Access Control Rules

After you have configured user claims, device claims, and resource properties, you then must protect files and folders by using a conditional expression that evaluates user and device claims against constant values or values within resource properties.

- A central access rule contains one or multiple criteria that the Windows operating system uses when evaluating access
- You create and configure central access rules in the Active Directory Administrative Center
- To create a new central access rule, you should:
  - Provide a name and description for the rule
  - Configure the target resources
  - Configure permissions

A central access rule contains one or more criteria that Windows operating systems use when evaluating access. For example, a central access rule can use conditional expressions to target specific files and folders. Each central access rule has a condition that determines which information the rule targets and the multiple permission entry lists that you use to manage the rule's current or proposed permission entries. You also can revert the rule's current permission entry list to its last known list of permission entries. Each central access rule can be a member of one or more central access policy objects.

Configuring Central Access Rules

You typically create and configure central access rules in the Active Directory Administrative Center. However, you also can use the Windows PowerShell command-line interface to perform the same tasks.

To create a new central access rule, do the following:

- Provide a name and description for the rule. You also should choose to protect the rule against accidental deletion.
- Configure the target resources. In the Active Directory Administrative Center, use the Target Resources section to create a scope for the access rule. You create the scope by using resource properties within one or more conditional expressions. You want to create a target condition based on the business requirement that drives this rule. For example, Resource.Compliancy Equals HIPAA. To simplify the process, you can keep the default value (All resources), but usually you apply some resource filtering. You can build the conditional expressions by using many logical and relational operators. You can use the following operators when building conditional expressions:

  Logical: AND, OR, NOT and Exists (resource properties)

Relational: =, != , <, >, <=, >=, Member_of, Device_Member of, Member_of_Any, Device_Member_of_Any, Any_of, Contains and NOT*
Also, you can join multiple conditional expressions within one rule by using AND or OR.

- In addition, you can group conditional expressions together to combine the results of two or more joined conditional expressions. The Target Resources section displays the currently configured conditional expression that is being used to control the rule's applicability.
- Configure permissions with either of the following options:
    - Use the following permissions as proposed permissions. Select this option to add the entries in the permissions list to the list of proposed permissions entries for the newly created central access rule. You can combine the proposed permissions list with file system auditing to model the effective access that users have to the resource, without having to change the entries in the current permissions list. Proposed permissions generate a special audit event to the event log that describes the proposed effective access for the users.

 **Note:** Proposed permissions do not apply to resources; they exist for simulation purposes only.

    - Use the following permissions as current permissions. Select this option to add the entries in the permissions list to the list of the current permissions entries for the newly created central access rule. The current permissions list represents the additional permissions that the Windows operating system considers when you deploy the central access rule to a file server. Central access rules do not replace existing security. When making authorization decisions, the Windows operating system evaluates the permission entries from the central access rule's current permissions list, from NTFS file system, and from the share permissions lists.

Once you are satisfied with your proposed permissions, you can convert them to current permissions. Alternatively, you can use current permissions in a test environment and effectively test access as specified in the Advanced Security tab to model how the policy applies to different users.

Creating and Managing Access Policies

Central access policies enable you to manage and deploy consistent authorization throughout an organization by using central access rules and central access policy objects. These policies act as a safety net that an organization applies across its servers. You use Group Policy to deploy a central access policy, and you manually apply the policies to all Windows Server 2012 file servers that will use DAC. A central access policy enables you to deploy a consistent configuration to multiple file servers. In addition, you can use the Data Classification Toolkit to apply a central access policy shared across multiple servers that reports on which central access policies are applied to shares.

- Central access policies enable you to manage and deploy consistent authorization throughout an organization
- The main component of a central access policy is a central access rule
- Central access policies act as a security net that an organization applies across its servers
- Group Policy is used to deploy a central access policy
- Manually apply the policies to all Windows Server 2012 file servers

The main component of a central access policy is central access rule. Central access policy objects represent a collection of central access rules. Before you create a central access policy, you should create a central access rule because polices are made up of rules.

You create central access policy objects in the Active Directory Administrative Center. To create a central access policy, you must have at least one central access rule created. After you create a central access policy, you have to publish it by using Group Policy. By doing this, you make the central access policy visible to the file servers in your organization. However, you still have to apply the central access policy manually to each folder that you want to protect with DAC.

Demonstration: Configuring Claims, Resource Properties, and Rules

In this demonstration, you will see how to:

- Configure claims.
- Configure resource properties.
- Configure access rules.

Demonstration Steps

1. In the Active Directory Administrative Center, in the navigation pane, click **Dynamic Access Control**.
2. Open the **Claim Types** container, and then create a new claim type for users and computers by using the following settings:
    - Source Attribute: **department**
    - Display name: **Company Department**
3. In the Active Directory Administrative Center, in the Tasks pane, click **New**, and then click **Claim Type**.
4. Create a new claim type for computers by using the following settings:
    - Source Attribute: **description**
    - Display name: **description**
5. In the Active Directory Administrative Center, click **Dynamic Access Control**, and then open the **Resource properties** container.
6. Enable the **Department** and **Confidentiality** Resource properties.
7. Open Properties for the **Department** property.
8. Add **Research** as a suggested value.
9. Open the **Global Resource Property List**, ensure that **Department** and **Confidentiality** are included in the list, and then click **Cancel**.
10. Click **Dynamic Access Control**, and then open the **Central Access Rules** container.
11. Create a new central access rule with the following values:
    - Name: **Department Match**
    - Target Resource: **use condition Resource-Department-Equals-Value-Research**
    - Current Permissions:
    - Remove **Administrators**
    - Add **Authenticated Users**,
    - Modify, with condition **User-Company Department-Equals-Resource-Department**

12. Create another central access rule with the following values:
    - Name: **Access Confidential Docs**
    - Target Resource: **use condition Resource-Confidentiality-Equals-Value-High**
    - Current Permissions:
    - Remove **Administrators**
    - Add **Authenticated Users**
    - Modify, and set first condition to: **User-Group-Member** of **each-Value-Managers**
13. Permissions: Set second condition to: **Device-Group-Member of each-Value-ManagersWKS**.

Implementing and Managing File Classifications

When you plan your DAC implementation, you should include file classifications. Although file classifications are not mandatory for DAC, they can enhance the automation of the entire process. For example, if you require that security-critical documents be accessible to top management only and classified with the attribute Confidentiality set to High, regardless of the server on which the documents exist, you should ask yourself how you identify these documents, and how to classify them appropriately.

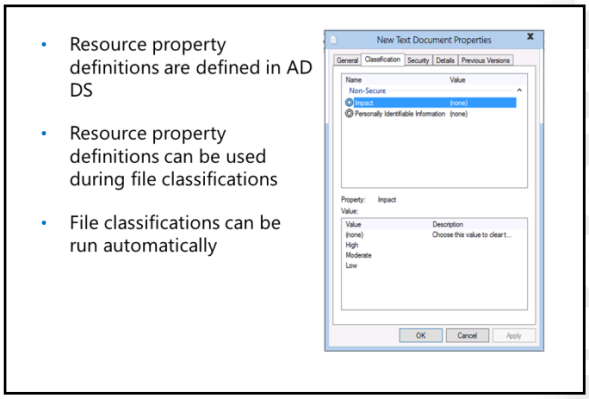

The file classification infrastructure (FCI) uses location-based classification, such as For this Folder structure Confidentiality is High. Also, you can use automatic classification rules to scan files automatically and then classify them according to the contents of the file. Classification and resource properties are defined centrally in AD DS so that these definitions can be shared across the file servers in an organization. You can create classification rules that scan files for a standard string or for a string that matches a pattern, or a regular expression. When a configured classification pattern is found in a file, that file is classified as configured in the classification rule.

To implement DAC effectively, you must have well-defined claims and resource properties. Although claims are defined by attributes for a user or a device, resource properties are most often manually created and defined. File classifications enable administrators to define automatic procedures for defining a desired property on the file, based on conditions specified in a classification rule. For example, you can set the Confidentiality property to High on all documents with contents that contain the word "secret." You then could use this property in DAC to specify that only employees with their employeeType attributes set to Manager can access those documents.

In Windows Server 2008 R2 and Windows Server 2012, classification management and file management tasks enable administrators to manage groups of files based on various file and folder attributes. With these tasks, you can automate file and folder maintenance tasks, such as cleaning up outdated data or protecting sensitive information.

Classification management is designed to ease the burden and management of data that is spread out in the organization. You can classify files in a variety of ways. In most scenarios, you classify files manually. The FCI in Windows Server 2012 enables organizations to convert these manual processes into automated policies. Administrators can specify file management policies based on a file's classification and then apply corporate requirements for managing data based on a business value.

You can use file classification to perform the following actions:

- Define classification properties and values, so you then can assign them to files by running classification rules.
- Classify a folder so that all the files within the folder structure inherit the classification.
- Create, update, and run classification rules. Each rule assigns a single predefined property and value to the files within a specified directory, based on installed classification add ins.

When you run a classification rule, reevaluate the files that are classified already. You can choose to overwrite existing classification values, or add the value to properties that support multiple values. You also can declassify files that are no longer in the classification criteria.

Demonstration: Configuring Classification Rules

This demonstration shows how to classify files by using a file classification mechanism.

Demonstration Steps

1. On LON-SVR1, open **File Server Resource Manager**.
2. Refresh Classification Properties, and then verify that the **Confidentiality** and **Department** properties are listed.
3. Create a classification rule with following values:
    - Name: Set Confidentiality
    - Scope: **C:\Docs**
    - Classification method: **Content Classifier**
    - Property: Confidentiality
    - Value: **High**
    - Classification Parameters: **String "secret"**
    - Evaluation Type: Re-evaluate existing property values, and then click Overwrite the existing value
4. Run the classification rule.
5. Open File Explorer, browse to the **C:\Docs** folder, and then open the Properties window for files Doc1.txt, Doc2.txt, and Doc3.txt.
6. Verify values for Confidentiality. Doc1.txt and Doc2.txt should have confidentiality set to High.

Lesson 3
Implementing DAC for Access Control

After you have configured the DAC building blocks, you must plan and implement DAC policies to actually control the resource access. Also, you have to learn how to manage and evaluate the effects of DAC. In this lesson, you will learn how to implement DAC.

Lesson Objectives

After completing this lesson, you will be able to:

- Plan Central Access Policies for file servers.
- Create and deploy Central Access Policies.
- See how access check works when DAC is in use.
- Manage and monitor DAC.
- Evaluate and manage DAC.

Planning Central Access Policies for File Servers

Implementing a central access policy is not mandatory for DAC. However, for a consistent configuration of access control on all file servers, you should implement at least one central access policy. By doing so, you enable all file servers within a specific scope to use a central access policy when protecting content in shared folders.

When planning deployment of central access policies, you should:
- Identify the resources that you want to protect
- Define the authorization policies
- Translate the authorization policies that you require into expressions
- Identify attributes for access filtering

Before you implement a central access policy, create a detailed plan as follows:

1. Identify the resources that you want to protect. If all these resources are on one file server or in just one folder, then you might not have to implement a central access policy. Instead, you can configure conditional access on the folder's ACL. However, if resources are distributed across several servers or folders, then you might benefit from deploying a central access policy. Data that might require additional protection might include payroll records, medical history data, employee personal information, and company customer lists. You also can use targeting within the central access rules to identify resources to which you want to apply a central access policy.
2. Define the authorization policies. These policies usually are defined from your business requirements. Some examples are:
    - All documents that have property Confidentiality set to High must be available only to managers.
    - Marketing documents from each country should be writable only by marketing people from the same country.
    - Only full-time employees should be able to access technical documentation from previous projects.
3. Translate the authorization policies that you require into *expressions*. In the case of DAC, expressions are attributes that are associated with both the resources, such as files and folders, and the users or

devices that seek access to these resources. These expressions state additional identification requirements that must be met to access protected data. Values that are associated with any expressions on the resource obligate the user or the device to produce the same value.

4. Lastly, you should break down the expressions that you have created to determine what claim types, security groups, resource properties, and device claims you must create to deploy your policies. In other words, you must identify the attributes for access filtering.

 **Note:** You are not required to use user claims to deploy central access policies. You can use security groups to represent user identities. We recommend that you start with security groups because it simplifies the initial deployment requirements.

Demonstration: Creating and Deploying Central Access Policies

This demonstration shows how to create and deploy central access policy.

Demonstration Steps

1. On LON-DC1, in the Active Directory Administrative Center, create a new central access policy with following values:
   - Name: **Protect confidential docs**
   - Rules included: **Access Confidential Docs**
2. Create another Central Access Policy with following values:
   - Name: **Department Match**
   - Rules included: **Department Match**.
3. On LON-DC1, from the Server Manager, open the Group Policy Management Console.
4. Create new GPO named **DAC Policy**, and in the Adatum.com domain, link it to **DAC-Protected** OU.
5. Edit the DAC Policy, browse to **Computer Configuration /Policies/Windows Settings/Security Settings/File System**, and then right-click **Central Access Policy**.
6. Click **Manage Central Access Policies**, click both **Department Match** and **Protect confidential docs**, click **Add**, and then click **OK**.
7. Close both the Group Policy Management Editor and the Group Policy Management Console.
8. On LON-SVR1, use Windows PowerShell to refresh Group Policy on LON-SVR1.
9. Open File Explorer, and then browse to the **C:\Docs** folder.
10. Apply the **Protect confidential docs** central policy to the **C:\Docs** folder.
11. Browse to the **C:\Research** folder.
12. Apply the **Department Match** Central Policy to the **C:\Research** folder.

How Does Access Check Work When DAC Is in Use

When planning and performing deployment of DAC, it is important that you know how DAC works with other access management technologies. In most environments, you will not be setting up DAC as the only access management technology. In most scenarios, you will deploy DAC on top of existing access management methods. In some cases, a company will decide to migrate access control entirely to DAC, while others might decide to keep it in coexistence with technologies such as NTFS file system permissions or AD RMS.

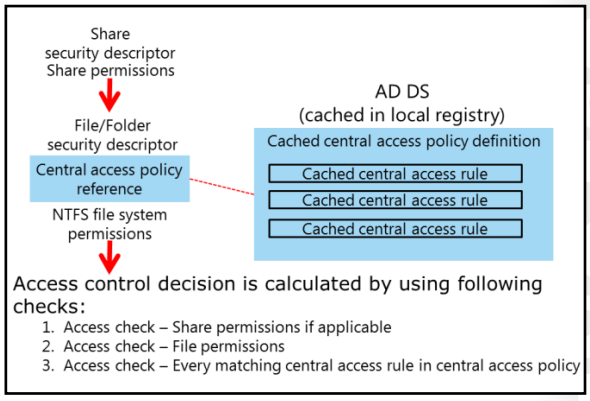

Because of this, it is important to know how DAC works with Share and NTFS file system permissions when they are applied on the same resource, such as a shared folder. This is similar to the scenario when you combine Share and NTFS file system permissions, because adding a central access policy keeps the same processing algorithm. Most restrictive access permission will always take effect if the same identity is listed on more than one ACL or access control rule. If that is not the case, the access control mechanism will apply share permissions, and then will apply NTFS file system permissions. Then it will process all Central Access Rules associated with the resource based on the assigned access policy. Similar to NTFS file system permissions, a central access policy will be applied when a user accesses the resource locally or remotely.

If, in some scenario, you want to control access to the resource by using DAC only, you can set Share and NTFS file system permissions to let all authenticated users access the resource, and then define precise permissions by using the access control rules.

Managing and Monitoring DAC

Once DAC is implemented, you might have to make some changes. For example, you might have to update the conditional expressions, or you might want to change the claims. However, you must plan carefully any changes that you make to DAC components.

Changing a central access policy can affect access drastically. For example, a change potentially could grant more access than desired, or it could restrict a policy too much, resulting in an excessive number of help desk calls. As a best practice, you should test changes before implementing a central access policy update.

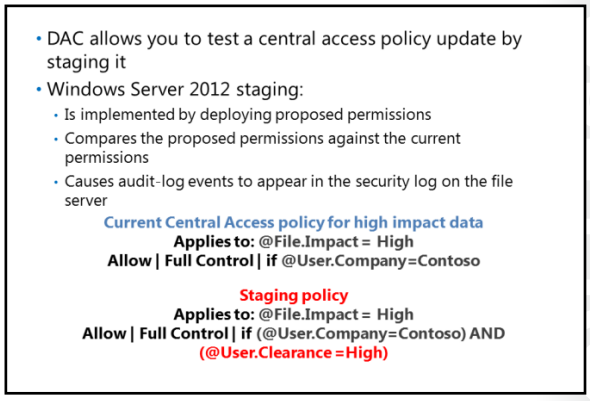

For this purpose, Windows Server 2012 introduces the concept of *staging*. Staging enables users to verify their proposed policy updates before enforcing them. To use staging, you deploy the proposed policy along with the enforced policies, but you do not actually grant or deny permissions. Instead, the Windows operating system logs an audit event, event 4818, any time the result of the staged policy differs from the result of the access check that uses the enforced policy.

You must first configure Group Policy to use staging. You should open the Group Policy Management Editor and navigate to: **Computer Configuration\Policies\Windows Settings\Security Settings\Advanced Audit Policy Configuration\Audit Policies\Object Access**. In this location you should enable Success and Failure auditing for the Audit Central Access Policy Staging and Audit File System policies.

Demonstration: Evaluating and Managing DAC

This demonstration shows how to evaluate and manage DAC.

Demonstration Steps

1. On LON-DC1, open the Group Policy Management Console.
2. Open the Group Policy Management Editor for **DAC Policy**.
3. Browse to **Computer Configuration\Policies\Windows Settings\Security Settings \Advanced Audit Policy Configuration\Audit Policies**, and then click **Object Access**.
4. Double-click **Audit Central Access Policy Staging**, select all three check boxes, and then click **OK**.
5. Double-click **Audit File System**, select all three check boxes, and then click **OK**.
6. Close the Group Policy Management Editor and the Group Policy Management Console.
7. On LON-DC1, open Active Directory Administrative Center, and then open the Properties for the **Department Match** central access rule.
8. In the **Proposed permissions** section, configure the condition for **Authenticated Users** as **User-Company Department-Equals-Value-Marketing**.
9. On LON-SVR1, refresh group policy settings.

Lesson 4
Implementing Access Denied Assistance

One of the most common causes for help desk support calls is the inability to access resources. When a user attempts to access a resource and does not have proper permissions, in most cases, a generic access-denied error will be presented. In Windows Server 2012 and Windows 8, there is a new feature called Access Denied Assistance. Administrators can use Access Denied Assistance to customize access-denied messages and to enable users to request access.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe Access Denied Assistance.
- Configure Access Denied Assistance.
- Implement Access Denied Assistance.

What Is Access Denied Assistance?

One of the most common errors that users receive when they try to access a file or folder on a remote file server is an access-denied error. Typically, this error occurs when a user tries to access a resource without having the proper permissions to do so, or because of incorrectly configured permissions on the resource ACLs. Using DAC can create further complications if it is implemented incorrectly. For example, users who should have permission to access a resource will not be granted access if a relevant attribute value used by DAC in their account is misspelled.

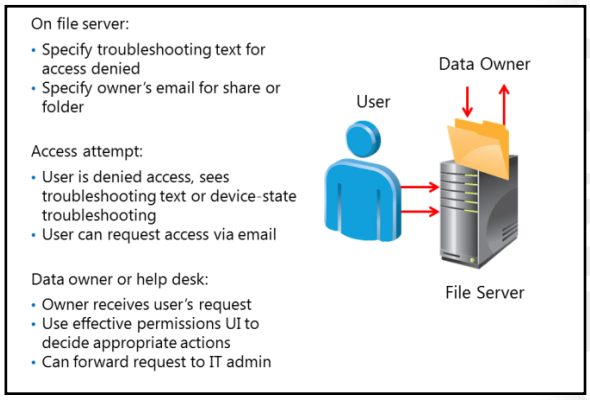

When users receive this error, they typically try to contact the administrator to obtain access. However, administrators usually do not approve access to resources, so they redirect users to someone else for approval.

In Windows Server 2012, there is a new feature to help users and administrators in such situations. This feature is called Access Denied Assistance. Access Denied Assistance helps users respond to access-denied issues without involving IT staff. It does this by providing information to users about the problem, and by directing them to the proper person.

Access Denied Assistance is a feature in the Windows 8 operating system that helps users to notify administrators when they are unable to access a resource. It allows an IT staff to properly diagnose a problem, and then to implement a resolution. Windows Server 2012 enables you to customize messages about denied access and provide users with the ability to request access without contacting the help desk or IT team. Combined with DAC, Access Denied Assistance can inform a file administrator of user and resource claims, and enable the administrator to make educated decisions about how to adjust policies or fix user attributes.

Note: Only Windows 8 or newer versions and Windows Server 2012 or newer support Access Denied Assistance.

Configuring Access Denied Assistance

When you plan for Access Denied Assistance, you should include the following:

- When implementing Access Denied Assistance:
  - Define messages that users will receive when they attempt to access resources
  - Determine whether users should be able to send a request for access
  - Determine recipients for the access-request email messages
  - Consider target operating systems
- Use Group Policy to enable and configure Access Denied Assistance
- Decide about the method for remediation

- Define messages that users will receive when they attempt to access resources for which they do not have access rights. The message should be informal and easy to understand.
- Determine whether users should be able to send a request for access via email, and if so, optionally configure the text that will be added to their email messages.
- Determine the recipients for access-request email messages. You can choose to send email to folder owners, file server administrators, or any other specified recipient. Messages should always be directed to the proper person. If you have a help desk tool or monitoring solution that allows email messages, you also can direct those messages to generate user requests automatically in your help desk solution.
- Decide on target operating systems. Access Denied Assistance only works with Windows 8 or Windows Server 2012 or newer versions.

The Access Denied Assistance feature provides three ways to troubleshoot issues with access-denied errors:

- Self-remediation. Administrators can create customized access-denied messages that are authored by the server administrator. By using the information in these messages, users can try to self-remediate access-denied cases. The message also can include URLs that direct users to self-remediation websites that are provided by the organization.
- Remediation by the data owner. Administrators can define owners for shared folders. This enables users to send email messages to data owners to request access. For example, if a user is left off a security group membership accidentally, or the user's department attribute value is misspelled, the data owner might be able to add the user to the group. If the data owner does not know how to grant access to the user, the data owner can forward this information to the appropriate IT administrator. This is helpful because the number of user support requests that escalate to the support desk should be limited to specialized cases, or cases that are difficult to resolve.
- Remediation by the help desk and file server administrators. If users cannot self-remediate issues, and if data owners cannot resolve the issue, then administrators can troubleshoot issues by accessing the UI to view the effective permissions for the user. Examples of when an administrator should be involved are cases where claims attributes or resource object attributes are defined incorrectly or contain incorrect information, or when the data itself seems to be corrupted.

You use Group Policy to enable the Access Denied Assistance feature. Open the Group Policy Object Editor and navigate to **Computer Configuration\Policies\Administrative Templates\System\Access-Denied Assistance**. In the Access Denied Assistance node, you can enable Access Denied Assistance, and you also can provide customized messages for users. Alternatively, you can use the FSRM console to enable Access Denied Assistance. However, if Access Denied Assistance is enabled in Group Policy, the appropriate settings in the FSRM console are disabled for configuration.

You also can use the FSRM Management Properties page to configure a customized Access Denied Assistance message for a particular folder tree within the server—for example, a per share message.

Demonstration: Implementing Access Denied Assistance

This demonstration shows how to configure and implement Access Denied Assistance.

Demonstration Steps

1. On LON-DC1, open the Group Policy Management Console and browse to **Group Policy objects**.
2. Edit the DAC Policy.
3. Under the Computer Configuration node, browse to **Policies\Administrative Templates\System**, and then click **Access-Denied Assistance**.
4. In the details pane, double-click **Customize Message for Access Denied errors**.
5. In the Customize Message for Access Denied errors window, click **Enabled**.
6. In the **Display the following message to users who are denied access** text box, type **You are denied access because of permission policy. Please request access.**
7. Select the **Enable users to request assistance** check box, and then click **OK**.
8. Double-click **Enable access-denied assistance on client for all file types**, enable it, and then click **OK**.
9. Close the Group Policy Management Editor and the Group Policy Management Console.
10. Switch to LON-SVR1, and refresh Group Policy.

Lesson 5
Implementing and Managing Work Folders

In today's business environment, it has become more and more common for people to use their own computers, tablets, and smart phones while they are at work. Users are always using the same UI, without need to change the UI each day. That is the result of the Bring Your Own Device (BYOD) approach that a lot of companies have adopted over the last few years. BYOD means the policy of permitting employees to bring personally owned mobile devices—laptops, tablets, and smart phones—to the workplace, and permitting them to use those devices to access privileged company information and applications.

To help users access business data on all their devices, Microsoft has implemented Work Folders technology. In this lesson, you will learn about how to implement Work Folders.

Lesson Objectives

After completing this lesson, you will be able to :

- Describe Work Folders.
- Configure Work Folders.
- Implement Work Folders.

What Are Work Folders?

For various reasons, storing user data on a local hard drive of a computer or a tablet is unsecure and inefficient. Since users commonly use more than one device, it is hard to keep these devices synchronized with business data, and it is hard to back up and protect this data efficiently.

- Work Folders enable users to access business data securely at any location and on any device
- Work Folders are managed by administrators
- Currently supported on Windows 8.1 devices, and support also is planned for iOS-based devices

As a result, users often use services such as Microsoft SkyDrive to store their data and to keep all their devices synchronized. However, these services are made for consumer data—not business data. Administrators cannot control the behavior of services such as SkyDrive on a user's private computer, which makes it difficult to implement in business environments.

On the other hand, users who have mobile computers or laptops that are members of a company's AD DS domain often need to access company data while they are offline. To date far, Offline Files have been used mostly to keep important data available locally on a user's computer, even when the computer was not connected to the network. However, Offline Files were synchronized only when the user connected to the company's local network. If the user was offline for a long time, there was a possibility that the use was working on old copies of data.

To overcome these problems, in Windows Server 2012 R2 and Windows 8.1, Microsoft has implemented a new technology named Work Folders. This technology enables users to access their business data independently of their location, and it enables administrators to manage the data and settings of this technology.

The main purpose of Work Folders is to provide access to the latest data, no matter where the user is located, internally or externally. Also, by using Work Folders, administrators can manage data and a user's connections to Work Folders. The administrator can enforce the encryption of Work Folders and can

control which users can use this functionality. The administrator also can enforce some security settings on the device that uses Work Folders, even if it is not a domain member.

Users can use Work Folders on various types of devices while they are in a local network, but also when they are out of the network—for example, while they are at home or traveling. Work Folders can be published to the Internet by using the Web Application Proxy functionality, also specific to Windows Server 2012 R2, which allows users to synchronize their data whenever they have an Internet connection.

Note: Currently, Work Folders are available only for Windows 8.1 client operating systems. However, by the time Windows 8.1 is available globally, Work Folders should be available for Windows 7 and iOS-based devices such as the iPad.

The following table shows the comparison between similar technologies for managing and accessing user data.

Technology	Personal data	Individual work data	Team/group work data	Personal devices	Data location
SkyDrive	Yes			Yes	Public cloud
SkyDrive Pro		Yes	Yes	Yes	Microsoft SharePoint/ Microsoft Office 365
Work Folders		Yes		Yes	File server
Folder Redirection / Client-side caching		Yes			File server

Configuring Work Folders

To use Work Folders, you should have at least one Windows Server 2012 R2 file server and at least one Windows Server 2012 R2 domain controller in your network. Work Folders is a role service of the File and Storage Services server role and can be installed easily by using Server Manager. It is a best practice, but is not mandatory, that you also install FSRM and Data Deduplication functionality if you want to manage user data more efficiently.

To use Work Folders, you should:

- Have at least one Windows Server 2012 R2 file server
- Have at least one Windows Server 2012 R2 domain controller
- Install Work Folders functionality on file server
- Provision a share where users' data will be stored
- Run New Sync Share Wizard to create Work Folders structure
- Configure clients to use Work Folders by using Group Policy or manually

Note: When you install Work Folders functionality, Internet Information Services (IIS) Hostable Web Core and IIS Management tools also will be installed. You do not have to configure any IIS settings, but you must assign a trusted Secure Sockets Layer (SSL) certificate to your file server in IIS Console and bind it to port 443 on the Default Web Site. The certificate should have both a file server name and the name under which you plan to publish your Work Folders, if those are different.

After you install Work Folders functionality, you should provision a share where users' data will be stored. A share can be stored on any location, such as folder on the local or iSCSI storage that is accessible and controlled by the file server where you installed Work Folders. When you create a root share, we recommend that you leave Share and NTFS file system permissions on their default values and that you enable access-based enumeration.

After you create a root share where users' Work Folders will be located, you should start New Sync Share Wizard to create the Work Folders structure. You should select the root folder that you provisioned as a share, and you also should choose the format for the subfolders naming. It can be a user alias, or alias@domain. If you have more than one domain in your AD DS forest, it is recommended that you choose the alias@domain naming format.

Sync Access can be controlled by explicitly listing users who will be able to use the Work Folders structure that you created, or by specifying a group. We recommend that you specify a group for later, easier administration. Also, we recommend that you disable permission inheritance for Work Folders so that each user has exclusive access to his or her files. At the end, you can enforce some additional security settings on devices being used to access Work Folders. You can enforce Work Folders with encryption and an automatic lock screen with password requirements.

 **Note:** Enforcement of security settings related to Work Folders is not achieved by using Group Policy. These settings are enforced when a user establishes the Work Folders connection, and they are applied on computers that are domain-joined and on computers that are not domain-joined.

Configuring Clients to Use Work Folders

Windows 8.1 clients can be configured manually to use Work Folders or by using Group Policy. For domain-joined computers, it is easier to configure settings by using Group Policy, but non-domain clients must be configured manually.

If you are using Group Policy to configure Work Folders automatically, there are two places where you should look. Work Folders are user-based, so configuration is performed in the user part of the GPO. When you open the Group Policy Editor, you should navigate to the **User Configuration\Policies\Administrative Templates\Windows Components\Work Folders**. Then you should open the Specify Work Folders settings and enable the policy. Also, you have to configure the Work Folders URL. This URL is the location of your file server where you enabled Work Folders. It usually is https://*fileserverFQDN*. In

In this same GPO setting, you have the option to force automatic setup for each user. This option should be considered with caution. If you enable it, all users this GPO applies to will have their Work Folders configured on each device they sign on to, without being prompted to do so. In some scenarios, you might not want to have this outcome.

You also can manage some Work Folders settings in the computer part of the GPO. If you navigate to **Computer Configuration\Policies\Administrative Templates\Windows Components\Work Folders**, you will find the option to Force automatic setup of Work Folders for all users. Computers that have this GPO setting applied will configure Work Folders for every user that signs on.

After you apply these Group Policy settings to the users' and optionally the computers' domain, users will be able to start using Work Folders.

If you also want to enable Work Folders on a non-domain joined computer, for example, on the tablet that an employee is using, you have to make manual configurations by using the Work Folders item in the Control Panel of Windows 8.1. You will have to provide a valid user name and password for the domain account that is allowed to use Work Folders, and a file server URL.

Demonstration: Implementing Work Folders

This demonstration shows how to implement Work Folders.

Demonstration Steps

1. On LON-SVR2, in Server Manager, click File and Storage Services, and then select **Work Folders**.
2. Start the New Sync Share Wizard.
3. Select **WF-Share**.
4. Use **User alias** for the structure for user folders.
5. Grant access to the **WFSync** user group.
6. Switch to LON-DC1.
7. Open the Group Policy Management Console.
8. Create new GPO, and name it **Work Folders** GPO.
9. Open the Group Policy Management Editor for Work Folders GPO.
10. Expand **User Configuration**/**Policies**/**Administrative Templates**/**Windows Components**, and then click **Work Folders**.
11. Enable Work Folders support, and then type **https://lon-svr2.adatum.com** as the Work Folders URL.
12. Link the Work Folders GPO to the domain.

Lab: Implementing Secure Data Access

Scenario

You are working as an administrator at A. Datum Corporation. The company has a wide and complex file server infrastructure. It manages access control to folder shares by using NTFS file system ACLs, but in some cases, that approach does not provide the desired results.

Most of the files used by departments are stored in shared folders dedicated to specific departments, but confidential documents sometimes appear in other shared folders. Only members of the Research team should be able to access Research team folders, and only Executive department managers should be able to access highly confidential documents.

The Security department also is concerned that managers are accessing files by using their home computers, which might not be highly secure. Therefore, you must create a plan for securing documents regardless of where they are located, and you must ensure that documents can be accessed from authorized computers only. Authorized computers for managers are members of the security group ManagersWks.

The Support department reports that a high number of calls are generated by users who cannot access resources. You must implement a feature that helps users understand error messages better and will enable them to request access automatically.

Many users use personal devices such as tablets and laptops to work from home and while at work. You have to provide them with an efficient way to synchronize business data on all the devices that they use.

Objectives

After completing this lab, you will be able to:

- Prepare for DAC deployment.
- Implement DAC.
- Validate and remediate DAC.
- Implement Work Folders.

Lab Setup

Estimated Time: 110 minutes

Virtual machines: 20412C-LON-DC1,

20412C-LON-SVR1,

20412C-LON-SVR2,

20412C-LON-CL1,

20412C-LON-CL2

User name: **Adatum/Administrator**

Password: **Pa$$w0rd**

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

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, click **20412C-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 by using the following credentials:
   - User name: **Adatum\Administrator**
   - Password: **Pa$$w0rd**
5. Repeat steps 2 through 4 for **20412C-LON-SVR1** and **20412C-LON-SVR2**.
6. Do not start or sign in to **20412C-LON-CL1** and **20412C-LON-CL2** machines until instructed by lab steps**.**

Exercise 1: Preparing for DAC Deployment

Scenario

To address the requirements from the lab scenario, you decide to implement DAC technology. The first step in implementing DAC is to configure the claims for the users and devices that access the files. In this exercise, you will review the default claims and create new claims based on department and computer group attributes. Also, you will configure the Resource Property lists and the Resource Property definitions. You will do this and then use the resource properties to classify files.

The main tasks for this exercise are as follows:

1. Preparing AD DS for DAC deployment.
2. Configuring user and device claims.
3. Configuring resource properties and resource property lists.
4. Implement file classifications.
5. Assign property to the Research folder.

▶ Task 1: Preparing AD DS for DAC deployment

1. On LON-DC1, in Server Manager, open **Active Directory Domains and Trusts** console.
2. Raise the domain and forest functional level to Windows Server 2012.
3. On LON-DC1, in Server Manager, open **Active Directory Users and Computers**.
4. Create a new Organizational Unit named **DAC-Protected**.
5. Move the LON-SVR1 and LON-CL1 computer objects into the **DAC-Protected** OU.
6. On LON-DC1, from Server Manager, open the Group Policy Management Console.
7. Edit the **Default Domain Controllers Policy** GPO.
8. In the Group Policy Management Editor, under **Computer Configuration**, expand **Policies**, expand **Administrative Templates**, expand **System**, and then click **KDC**.
9. Enable the **KDC support for claims, compound authentication and Kerberos armoring** policy setting.
10. In the **Options** section, click **Always provide claims**.
11. On LON-DC1, refresh Group Policy.
12. Open Active Directory Users and Computers, and in the **Users** container, create a security group named **ManagersWKS**.

13. Add **LON-CL1** to the **ManagersWKS** group.
14. Verify that user Aidan Delaney is a member of **Managers** department, and that Allie Bellew is the member of the **Research** department. Department entries should be filled in for the appropriate Organization attribute in each user profile. After you verify these values, click Cancel and don't make any changes.

▶ Task 2: Configuring user and device claims

1. On LON-DC1, open the **Active Directory Administrative Center**.
2. In the Active Directory Administrative Center, in the navigation pane, click **Dynamic Access Control**.
3. Open the Claim Types container, and then create a new claim type for users and computers by using the following settings:
   - Source Attribute: **department**
   - Display name: **Company Department**
   - Suggested values**: Managers, Research**
4. In the Active Directory Administrative Center, in the Tasks pane, click **New**, and then click **Claim Type**.
5. Create a new claim type for computers by using the following settings:
   - Source Attribute: **description**
   - Display name: **description**

▶ Task 3: Configuring resource properties and resource property lists

1. In the Active Directory Administrative Center, click **Dynamic Access Control**, and then open the **Resource Properties** container.
2. Enable the **Department** and **Confidentiality** Resource properties.
3. Open Properties for the **Department** property.
4. Add **Research** as a suggested value.
5. Open the **Global Resource Property List**, ensure that **Department** and **Confidentiality** are included in the list, and then click **Cancel**.
6. Close the **Active Directory Administrative Center**.

▶ Task 4: Implement file classifications

1. On LON-SVR1, open the File Server Resource Manager.
2. Refresh **Classification Properties**, and then verify that **Confidentiality** and **Department** properties are listed.
3. Create a classification rule with following values:
   - Name: **Set Confidentiality**
   - Scope: **C:\Docs**
   - Classification method: **Content Classifier**
   - Property: **Confidentiality**
   - Value: **High**

- Classification Parameters: **String "secret"**
- Evaluation Type: **Re-evaluate existing property values**, and then click **Overwrite the existing value**

4. Run the classification rule.
5. Open a File Explorer window, browse to the **C:\Docs** folder, and then open the Properties window for files Doc1.txt, Doc2.txt, and Doc3.txt.
6. Verify values for Confidentiality. Doc1.txt and Doc2.txt should have confidentiality set to High.

▶ Task 5: Assign property to the Research folder

1. On LON-SVR1, open File Explorer..
2. Browse to **C:\Research**, and open its properties.
3. On the **Classification** tab, set the **Department** value to **Research**.

Results: After completing this exercise, you will have prepared Active Directory Domain Services (AD DS) for Dynamic Access Control (DAC) deployment, configured claims for users and devices, and configured resource properties to classify files.

Exercise 2: Implementing DAC

Scenario

The next step in implementing DAC is to configure the central access rules and policies that link claims and property definitions. You will configure rules for DAC to address the requirements from the lab scenario. After you configure DAC rules and policies, you will apply the policy to a file server.

The main tasks for this exercise are as follows:

- Configure central access rules.
- Configure central access policies.
- Apply central access policies to a file server.

▶ Task 1: Configure central access rules

1. On LON-DC1, in Server Manager, click **Tools**, and then click **Active Directory Administrative Center**.
2. Click **Dynamic Access Control**, and then open the **Central Access Rules** container.
3. Create a new central access rule with the following values:
   - Name: **Department Match**
   - Target Resource: **use condition Resource-Department-Equals-Value-Research**
   - Current Permissions:
     - Remove **Administrators**
     - Add **Authenticated Users**,
     - Modify, with condition **User-Company Department-Equals-Resource-Department**

4. Create another central access rule with the following values:
   - Name: **Access Confidential Docs**
   - Target Resource: **use condition Resource-Confidentiality-Equals-Value-High**
   - Current Permissions:
     - Remove **Administrators**
     - Add **Authenticated Users**
     - Modify, and set the first condition to: **User-Company Department-Equals-Value-Managers**
     - Permissions: Set the second condition to: **Device-Group-Member** of **each-Value-ManagersWKS**

▶ Task 2: Configure central access policies

1. On LON-DC1, in the Active Directory Administrative Center, create a new central access policy with following values:
   - Name: **Protect confidential docs**
   - Rules included : **Access Confidential Docs**
2. Create another central access policy with following values:
   - Name: **Department Match**
   - Rules included: **Department Match**
3. Close the Active Directory Administrative Center.

▶ Task 3: Apply central access policies to a file server

1. On LON-DC1, from the **Server Manager**, open the **Group Policy Management** console.
2. Create new GPO named **DAC Policy**, and in the Adatum.com domain, link it to the **DAC-Protected** OU.
3. Edit the **DAC Policy**, browse to **Computer Configuration /Policies/Windows Settings/Security Settings/File System**, and then right-click **Central Access Policy**.
4. Click **Manage Central Access Policies**, click both **Department Match** and **Protect confidential docs**, click **Add**, and then click **OK**.
5. Close the Group Policy Management Editor and the Group Policy Management Console.
6. On LON-SVR1, use Windows PowerShell to refresh Group Policy on LON-SVR1.
7. Open File Explorer, and then browse to the **C:\Docs** folder.
8. Apply the **Protect confidential docs** central policy to the **C:\Docs** folder.
9. Browse to the **C:\Research** folder.
10. Apply the **Department Match** central policy to the **C:\Research** folder.

Results: After completing this exercise, you will have implemented DAC.

Exercise 3: Validating and Remediating DAC

Scenario

To ensure that the DAC settings are configured correctly, you will test various scenarios for users to access files. You will try approved users and devices and unapproved users and devices. You also will validate the access-remediation configuration.

The main tasks for this exercise are as follows:

- Access file resources as an approved user.
- Access file resources as an unapproved user.
- Evaluate user access with DAC.
- Configure access-denied remediation.
- Request access remediation.

▶ Task 1: Access file resources as an approved user

1. Start LON-CL1 and LON-CL2 virtual machines.
2. Sign in to LON-CL1 as **Adatum\Allie** with the password **Pa$$w0rd**.
3. Try to open documents inside the **\\LON-SVR1\Research** folder.
4. Sign out of LON-CL1.
5. Sign in to LON-CL1 as **Adatum\Aidan** with the password **Pa$$w0rd**.
6. Try to open files inside the **\\LON-SVR1\Docs** folder.

Note: Both attempts should succeed.

7. Sign out of LON-CL1.

▶ Task 2: Access file resources as an unapproved user

1. Sign in to LON-CL2 as **Adatum\Aidan** with the password **Pa$$w0rd**.
2. Open the **\\LON-SVR1\Docs** folder. Try to open files Doc1.txt and Doc2.txt.
3. Sign out of LON-CL2.
4. Sign in to LON-CL2 as **Adatum\April** with the password **Pa$$word**.
5. Open **\\LON-SVR1\Docs** folder, and then try to open Doc3.txt file. You should be able to open that document.
6. While still signed in as April, try to open the **\\LON-SVR1\Research** folder. You should be unable to access the folder.
7. Sign out of LON-CL2.

▶ Task 3: Evaluate user access with DAC

1. On LON-SVR1, open the Properties for the **C:\Research** folder.
2. Open the Advanced options for **Security**, and then click **Effective Access**.
3. Click **select a user**, and in the Select User, Computer, Service Account, or Group window, type **April**, click **Check Names**, and then click **OK**.

4. Click **View effective access**, and then review the results. The user should not have access to this folder.
5. Click **Include a user claim**, and then in the drop-down list box, click **Company Department**.
6. In the **Value** text box, type **Research**, and then click **View Effective access**. The user should now have access.
7. Close all open windows.

▶ Task 4: Configure access-denied remediation

1. On LON-DC1, open the Group Policy Management Console, and then browse to Group Policy objects.
2. Edit the DAC Policy.
3. Under the Computer Configuration node, browse to **Policies\Administrative Templates\System**, and then click **Access-Denied Assistance**.
4. In the details pane, double-click **Customize Message for Access Denied errors**.
5. In the Customize Message for Access Denied errors window, click **Enabled**.
6. In the **Display the following message to users who are denied access** text box, type **You are denied access because of permission policy. Please request access.**
7. Select the **Enable users to request assistance** check box, and then click **OK**.
8. Double-click **Enable access-denied assistance on client for all file types**, enable it, and then click **OK**.
9. Close the Group Policy Management Editor and the Group Policy Management Console.
10. Switch to LON-SVR1, and then refresh Group Policy.

▶ Task 5: Request access remediation

1. Sign in to LON-CL1 as **Adatum\April** with the password **Pa$$w0rd**.
2. Try to access the **\\LON-SVR1\Research** folder
3. Request assistance when prompted. Review the options for sending a message, and then click **Close.**
4. Sign out of LON-CL1.

Results: After completing this exercise, you will have validated DAC functionality.

Exercise 4: Implementing Work Folders

Scenario

To address the requirements for allowing employees to use their own devices to access and synchronize company data, you decide to implement Work Folders for a limited number of users.

The main tasks for this exercise are as follows:

- Install Work Folders functionality, configure SSL certificate and create WFSync group.
- Provision a share for Work Folders.
- Configuring and implementing Work Folders.

- Validate Work Folders functionality.
- Prepare for the next module.

▶ Task 1: Install Work Folders functionality, configure SSL certificate and create WFSync group

1. Sign in to LON-SVR2 as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. Start Server Manager.
3. Add the **Work Folders** role service by using the Add Roles and Features Wizard.
4. Open **Internet Information Services (IIS) Manager** console.
5. Create a domain certificate for lon-svr2.adatum.com as follows:
    a. Common name: **lon-svr2.adatum.com**
    b. Organization: **Adatum**
    c. Organizational unit: **IT**
    d. City/locality : **Seattle**
    e. State/province : **WA**
    f. Country/region: **US**
6. Assign this certificate to the https protocol on the Default Web Site.
7. Open Active Directory Users and Computers console on LON-DC1.
8. Create security group WFSync in Users container.
9. Add Aidan Delaney from Managers OU as a member of the WFSync group.

▶ Task 2: Provision a share for Work Folders

1. On LON-SVR2, in Server Manager, expand **File and Storage Services**, and then click **Shares**.
2. Start the New Share Wizard.
3. Select the **SMB Share – Quick** profile.
4. Name the share **WF-Share**.
5. Enable access-based enumeration.
6. Leave all other options on default values.

▶ Task 3: Configuring and implementing Work Folders

1. On LON-SVR2, in Server Manager, expand **File and Storage Services**, and then select **Work Folders**.
2. Start the New Sync Share Wizard.
3. Select the share that you created in previous step: WF-Share.
4. Use user alias for the structure for user folders.
5. Grant access to the WFSync user group.
6. Switch to LON-DC1.
7. Open Group Policy Management.
8. Create a new GPO and name it **Work Folders GPO**.
9. Open the Group Policy Management Editor for Work Folders GPO.

10. Expand **User Configuration / Policies / Administrative Templates / Windows Components**, and then click **Work Folders**.
11. Enable the Work Folders support and type **https://lon-svr2.adatum.com** as the Work Folders URL.
12. Link the Work Folders GPO to the domain.

▶ Task 4: Validate Work Folders functionality

1. Sign in to LON-CL1 as **Adatum\Aidan** with the password **Pa$$w0rd**.
2. Open Windows PowerShell, and then refresh Group Policy.
3. Open File Explorer, click This PC and then make sure that Work Folders are created.
4. Open Work Folders applet from Control Panel and apply security policies.
5. Make a few text files in Work Folders.
6. Ensure that they are synchronized.
7. Sign in to LON-CL2 as **Adatum\Aidan** with the password **Pa$$w0rd**.
8. Open Windows PowerShell, and then refresh Group Policy.
9. Open File Explorer, click This PC and then make sure that Work Folders are created.
10. Open Work Folders applet from Control Panel and apply security policies.
11. Ensure that the files that you created on LON-CL1 are present.

▶ Task 5: Prepare for the next module

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20412C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps two and three for **20412C-LON-SVR1**, **20412C-LON-SVR2, 20412C-LON-CL1**, and **20412C-LON-CL2**.

Results: After completing this exercise, you will have configured Work Folders.

Question: How do file classifications enhance the usage of DAC?

Question: Can you implement DAC without central access policy?

Module Review and Takeaways

Best Practice

Use central access policies instead of configuring conditional expressions on resources.

- Enable Access Denied Assistance settings.
- Always test changes that you have made to central access rules and central access policies before you implement them.
- Use file classifications to assign properties to files.
- Use Work Folders to synchronize business data across devices.
- Use Workplace Join in Bring Your Own Device (BYOD) scenarios.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Claims are not populated with the appropriate values. A conditional expression does not allow access.	

Review Questions

Question: What is a claim?

Question: What is the purpose of Central Access Policy?

Question: What is the BYOD concept?

Question: How do Work Folders support BYOD concept?

Tools

Tool	Use	Location
Active Directory Administrative Center	Administering and creating claims, resource properties, rules, and policies	Administrative tools
Group Policy Management Console (GPMC)	Managing Group Policy	Administrative tools
Group Policy Management Editor	Editing GPOs	GPMC

Module 4

Implementing Distributed Active Directory® Domain Services Deployments

Contents:

Module Overview

For most organizations, the Active Directory® Domain Services (AD DS) deployment may be the single most important component in the IT infrastructure. When organizations deploy AD DS or any of the other Active Directory-linked services within the Windows Server® 2012 operating system, they are deploying a central authentication and authorization service that provides single sign-on (SSO) access to many other network services and applications in the organization. AD DS also enables policy-based management for user and computer accounts.

Most organizations deploy only a single AD DS domain. However, some organizations also have requirements that necessitate that they deploy a more complex AD DS deployment, which may include multiple domains or multiple forests.

This module describes the key components of a complex AD DS environment, and how to install and configure a complex AD DS deployment.

Objectives

After completing this module, you will be able to:

- Describe the components of distributed AD DS deployments.
- Explain how to deploy a distributed AD DS deployment.
- Explain how to configure AD DS trusts.

Lesson 1
Overview of Distributed AD DS Deployments

Before you start to configure a complex AD DS deployment, it is important to know the components that constitute the AD DS structure, and how they interact with each other to help provide a scalable and secure IT environment. The lesson starts by examining the various components of an AD DS environment, and then explores reasons why an organization may choose to deploy a complex AD DS environment.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the components of an AD DS environment.
- Explain how AD DS domains and forests form boundaries for security and administration.
- Describe reasons for having more than one domain in an AD DS environment.
- Explain reasons for having more than one forest in an AD DS environment.
- Explain the importance of Domain Name System (DNS) in a complex AD DS structure.

Discussion: AD DS Components Overview

An AD DS environment has various components, and it is important for you to understand the purpose of each component, and how they interact with each other.

Question: What is an AD DS domain?

Question: What is an AD DS domain tree?

Question: What is an AD DS forest?

Question: What are trust relationships?

Question: What is the global catalog?

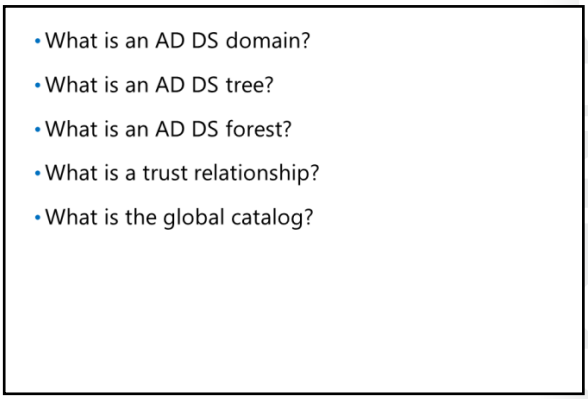

Overview of Domain and Forest Boundaries in an AD DS Structure

AD DS domains and forests provide different types of boundaries within an AD DS deployment. An understanding of the different types of boundaries is essential to managing a complex AD DS environment.

AD DS object	Boundary type
Domain	Domain partition replication
	Administrative permissions
	Group Policy application
	Auditing
	Password and account policies
	Domain DNS zone replication
Forest	Security boundary
	Schema partition replication
	Configuration partition replication
	Global catalog replication
	Forest DNS zone replication

AD DS Domain Boundaries

The AD DS domain provides the following boundaries:

- Replication boundary for the domain partition. All AD DS objects that exist in a single domain are stored in the domain partition in the AD DS database on each domain controller in the domain. The replication process ensures that all originating updates are replicated to all of the other domain controllers in the same domain. Data in the domain partition is not replicated to domain controllers in other forests.
- Administration boundary. By default, an AD DS domain includes several groups, such as the Domain Admins group, that have full administrative control over the domain. You can also assign administrative permissions to user accounts and groups within domains. With the exception of the Enterprise Admins group in the forest root domain, administrative accounts do not have any administrative rights in other domains in the forest or in other forests.
- Group Policy application boundary. Group Policies can be linked at the following levels: local, site, domain, and organizational unit (OU). Apart from site-level Group Policies, the scope of Group Policies is the AD DS domain. There is no inheritance of Group Policies from one AD DS domain to another, even if one AD DS domain is lower than another in a domain tree.
- Auditing boundary. Auditing is centrally managed by using Group Policy Objects (GPOs). The maximum scope of these settings is the AD DS domain. You can have the same audit settings in different AD DS domains, but then they must be managed separately in each domain.
- Password and account policy boundaries. By default, password and account policies are defined at the domain level and applied to all domain accounts. While it is possible to configure fine-grained password policies to configure different policies for specific users within a domain, you cannot apply the password and account policies beyond the scope of a single domain.
- Replication boundary for domain DNS zones. One of the options when you configure DNS zones in an AD DS environment is to configure Active Directory–integrated zones. This means that instead of the DNS records being stored locally on each DNS Server in text files, they are stored and replicated in the AD DS database. The administrator can then decide whether to replicate the DNS information to all domain controllers in the domain (regardless of whether they are DNS servers), to all domain controllers that are DNS servers in the domain, or to all domain controllers that are DNS servers in the forest. By default, when you deploy the first domain controller in an AD DS domain, and configure that server as a DNS server, two separate replication partitions called domainDnsZones and forestDnsZones are created. The domainDnsZones partition contains the domain-specific DNS records, and is replicated only to other DNS servers that are also AD DS domain controllers in the domain.

AD DS Forest Boundaries

The AD DS forest provides the following boundaries:

- Security boundary. The forest boundary is a security boundary because, by default, no account outside the forest has any administrative permissions inside the forest.
- Replication boundary for the schema partition. The schema partition contains the rules and syntax for the AD DS database. This is replicated to all the domain controllers in the AD DS forest.
- Replication boundary for the configuration partition. The configuration partition contains the details of the AD DS domain layout, including: domains, domain controllers, replication partners, site and subnet information, and Dynamic Host Configuration Protocol (DHCP) authorization or Dynamic Access Control configuration. The configuration partition also contains information about applications that are integrated with the AD DS database. An example of one application is Exchange Server® 2010. This partition is replicated to all domain controllers in the forest.
- Replication boundary for the global catalog. The global catalog is the read-only list containing every object in the entire AD DS forest. To keep it to a manageable size, the global catalog contains only some attributes for each object. The global catalog is replicated to all domain controllers in the entire forest that are also global catalog servers.
- Replication boundary for the Forest DNS zones. The forestDnsZones partition is replicated to all domain controllers in the entire forest that are also DNS servers. This zone contains records that are important to enable forest-wide DNS name resolution.

Why Implement Multiple Domains?

Many organizations can function adequately with a single AD DS domain. However, some organizations have requirements that necessitate that they deploy multiple domains. These requirements can include:

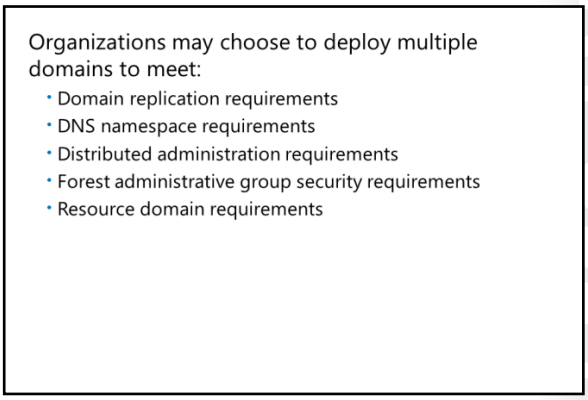

- Domain replication requirements. In some cases, organizations have several large offices that are connected by slow or unreliable wide area networks (WANs). The network connections may have enough bandwidth to support AD DS replication of the domain partition. In this case, it might be better to install a separate AD DS domain in each office.
- DNS namespace requirements. Some organizations have a requirement to have more than one DNS namespace in an AD DS forest. This is typically the case when one company acquires another company or merges with another organization, and there is need to preserve the domain names from the existing environment. It is possible to provide multiple user principal names (UPNs) for users in a single domain, but many organizations choose to deploy multiple domains in this scenario.
- Distributed administration requirements. Organizations may have corporate security or political requirements to have a distributed administration model. Organizations can achieve administrative autonomy by deploying a separate domain. With this deployment, domain administrators have complete control over their domains.

Note: Deploying separate domains provides administrative autonomy, but not administrative isolation. The only way to ensure administrative isolation is to deploy a separate forest.

- Forest administrative group security requirements. Some organizations may choose to deploy a dedicated or empty root domain. This is a domain that does not have any user accounts other than the default forest root domain accounts. The AD DS forest root domain has two groups—the Schema Admins group and the Enterprise Admins group—that do not exist in any other domain in the AD DS forest. Because these groups have far-reaching rights in the AD DS forest, you may want to restrict the use of these groups by only using the AD DS forest root domain to store them.
- Resource domain requirements. Some organizations deploy resource domains to deploy specific applications. With this deployment, all user accounts are located in one domain, whereas the application servers and application administration accounts are deployed in a separate domain. This enables the application administrators to have complete domain administrative permissions in the resource domain, without enabling any permissions in the domain that contains the regular user accounts.

Note: As a best practice, choose the simplest design that achieves the required goal, as it will be less costly to implement and more straightforward to administer.

Why Implement Multiple Forests?

Organizations may sometimes require that their AD DS design contains more than one forest. There are several reasons why one AD DS forest may not be sufficient:

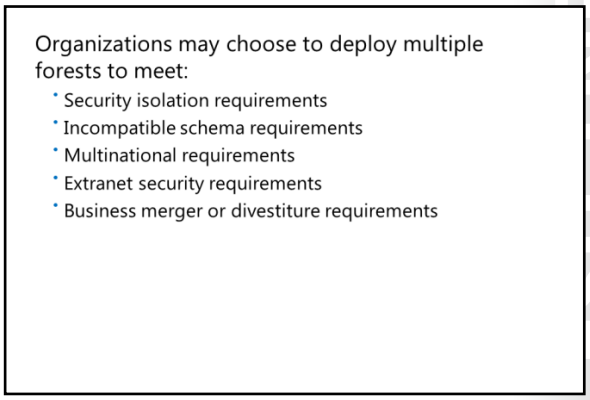

- Security isolation requirements. If an organization requires administrative isolation between two of its different parts, then the organization must deploy multiple AD DS forests. Separate AD DS forests are often deployed by government defense contractors and other organizations, for whom the isolation of security is a requirement.
- Incompatible schemas. Some organizations may require multiple forests because they require incompatible schemas or incompatible schema change processes. The schema is shared between all domains in a forest.
- Multinational requirements. Some countries have strict regulations regarding the ownership or management of enterprises within the country. Having a separate AD DS forest may provide the administrative isolation required by legislation.
- Extranet security requirements. Some organizations have several servers deployed in a perimeter network. These servers may need AD DS to authenticate user accounts, or may use AD DS to enforce policies on the servers in the perimeter network. To ensure that the extranet AD DS is as secure as possible, organizations often configure a separate AD DS forest in the perimeter network.

- Business merger or divestiture requirements. One of the most common reasons why organizations have multiple AD DS forests is because of business mergers. When organizations merge, or one organization purchases another, the organizations need to evaluate the requirement for merging the AD DS forests deployed in both organizations. Merging the AD DS forests provides benefits related to simplified collaboration and administration. However, if the two different groups in the organization will continue to be managed separately, and if there is little need for collaboration, it may not be worth the expense to merge the two forests. In particular, if there is a plan to sell one part of the company, it is preferable to retain the two organizations as separate forests.

Best Practice: As a best practice, choose the simplest design that achieves the required goal, as it will be less costly to implement and more straightforward to administer.

Integrating On-Premises AD DS with Cloud Services

There are two ways to extend your Active Directory Domain Services into the cloud. You can federate with the Windows Azure™ Active Directory, or you can install Windows Server 2012 R2 into Windows Azure virtual machine and promote the virtual machine to be a Domain Controller.

- Windows Azure AD:
  - Is a shared environment
  - Patching and upgrading is maintained by Microsoft
  - Can synchronize with on-premises AD DS
  - Does not support AD DS integrated applications
- AD in Azure:
  - Is a private Environment
  - Patching and upgrading is the responsibility of the customer
  - Can be part of on-premises AD DS
  - Supports AD DS aware applications

Windows Azure Active Directory

Windows Azure Active Directory (Windows Azure AD) is an Azure-based 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), and Active Directory Rights Management Service (AD RMS). Besides providing Windows Azure AD services, Windows Azure also currently supports Windows Azure Access Control Service. This service supports integration with third-party identity management as well as federation with your on-premises AD DS.

Installing Active Directory in 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 an on-premises infrastructure apply to deploying the same applications and servers to Windows Azure.

When you implement AD DS in Windows Azure, consider the following:

- Service healing. 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 UPN; when an AD DS system is rolled back, duplicate UPNs could be created. To prevent this, Windows Server 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 in Windows Azure.

When you deploy Windows Server 2012 AD DS on Windows Azure, the virtual machines is subject to the same guidelines as when you run AD DS on-premises on a virtual machine. These guidelines include the following:

- Time Synchronization. A Windows-based AD DS domain infrastructure relies on all communicating machines having the correct time.
- 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 or access resources, and applications and certain services may not run as well as other applications or services. It is very important that there is no single point of failure of your domain controller AD DS infrastructure when you virtualize domain controllers.

Special Considerations for Installing Active Directory in Windows Azure

Because you do not control several aspects of Windows Azure virtual machines, there are some special considerations for installing Active Directory in Windows Azure. These include:

- IP addressing. All Windows Azure virtual machines receive DHCP addresses. Your Windows Azure Virtual Network must be provisioned before the first Windows Azure-based domain controller is provisioned.
- DNS. Windows Azure built-in DNS does meet the requirements of Active Directory, such as Dynamic DNS and SRV records. You can install DNS with your domain controller; however, the domain controller cannot be configured with a static address. To alleviate potential issues, Windows Azure DHCP leases never expire.
- Disks. Windows Azure virtual machines use read-write host caching for operating system (OS) virtual hard disks. While this can improve the performance of the virtual machine, if Active Directory components are installed on the OS disk, data loss would be possible in the event of a disk failure. Additional Windows Azure hard disks attached to a VM have the caching turned off. When you install Active Directory in Windows Azure, the NTDS.DIT and SYSVOL folders should be located on an additional disk in the Windows Azure VM.

Implementing Windows Azure AD

Windows Azure AD is already built in and available for your use. When you sign up to use Windows Azure AD, you are using a portion of the existing Windows Azure AD. When you sign up for Windows Azure AD, you typically start with a trial subscription.

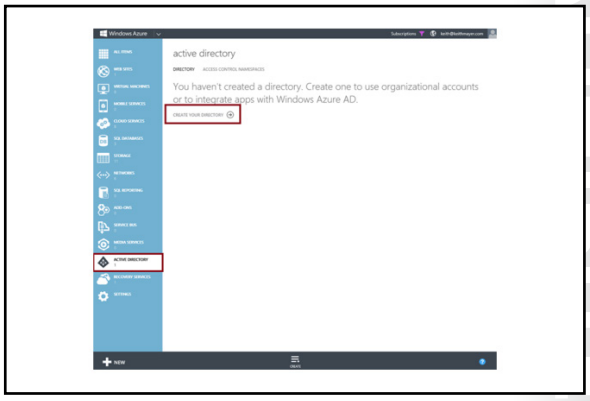

After you activate a trial, use the following steps to sign up for Windows Azure AD:

1. Sign in with your Windows Live account.
2. In the **Windows Azure Management Portal**, click **Active Directory** in the navigation tree to access the Active Directory page.
3. Click **CREATE YOUR DIRECTORY** to launch the **Create Directory** form and create your new Active Directory domain instance.
4. Complete the Create Directory form with the following information:
    - **Domain Name**. Enter a unique name for your new Active Directory domain instance. The domain you create will be provisioned as a subdomain inside the onmicrosoft.com public DNS domain. You can assign a custom DNS namespace to this domain after you complete initial provisioning.
    - **Country or region**. Select your closest country or region. This is used by Windows Azure to determine the Azure Datacenter Region in which your Active Directory domain instance is provisioned, and it cannot be changed after provisioning.
    - **Organization name**. Enter your organization's name.
5. Once the Domain is provisioned, you can continue to configure:
    - **Users**. Create and manage cloud-based users.
    - **Integrated Apps**. Integrate your cloud-based applications with Windows Azure AD.
    - **Domains**. Add a custom DNS domain name.
    - **Directory Integration**. Configure integration with an on-premises Windows Server Active Directory forest.

DNS Requirements for Complex AD DS Environments

AD DS requires DNS to function correctly, and implementing DNS in a multi-domain or multi-forest environment requires an extra level of planning.

When you deploy a DNS structure to support a complex AD DS environment, you will need to address several important configuration areas, including:

When implementing DNS in a complex AD DS environment, you should:
- Verify the DNS client configuration
- Verify and monitor DNS name resolution
- Optimize DNS name resolution between multiple namespaces
- Use AD DS integrated DNS zones
- Consider deploying a GlobalNames zone
- Design interoperability for DNS in Windows Azure and on-premise

- Verify the DNS client configuration. Configure all computers in the AD DS domain with at least two addresses of functional DNS servers. All computers must have good network connectivity with DNS servers.
- Verify and monitor DNS name resolution. Verify that all of your computers, including domain controllers, are able to perform successful DNS lookups for all domain controllers in the forest. Domain controllers need to be able to connect to other domain controllers to successfully replicate changes to AD DS. Client computers must be able to locate domain controllers by using service (SRV) resource records, and need to be able the resolve the domain controller names to IP addresses. In a multi-domain or multi-forest environment, client computers may need to locate a variety of cross-forest services, including Key Management Servers for Windows Activation, Terminal Services

Licensing servers, licensing servers for specific applications and domain controllers in any domain to validate trusts when accessing resources in another domain.

- Optimize DNS name resolution between multiple namespaces. When organizations deploy multiple trees in an AD DS forest, or when they deploy multiple forests, name resolution is more complicated because you need to manage multiple domain namespaces. Use DNS features such as conditional forwarding, stub zones, and delegation to optimize the process of resolving computer names across the namespaces.
- Use AD DS integrated DNS zones. When you configure a DNS zone as AD DS integrated, the DNS information is stored in AD DS and replicated through the normal AD DS replication process. This optimizes the process of replicating changes throughout the forest. You can also configure the scope of replication for the DNS zones. By default, domain-specific DNS records will be replicated to other domain controllers that are also DNS servers in the domain. DNS records that enable cross-domain lookups are stored in the _msdcs.forestrootdomainname zone and are replicated to domain controllers that are also DNS servers in the entire forest. This default configuration should not be changed.
- Deploying a GlobalNames zone. A GlobalNames zone allows you to configure single name resolution for DNS names in your forest. Previously, Windows Internet Name Service (WINS) was configured in a domain to support single-label name resolution. A GlobalNames zone can be used to replace WINS in your environment, especially if you deploy Internet Protocol version 6 (IPv6), as WINS does not support IPv6 addressing.
- When you extend your AD DS domain into Windows Azure, you must take a few extra steps. Windows Azure's built-in DNS does not support AD DS domains; to support your cloud-based domain components, you need to do the following:
  - Configure an AD DS site for your Windows Azure subnet.
  - Register your on-premises DNS with Windows Azure so that it is accessible from Windows Azure.
  - Register your cloud-based DNS with Windows Azure.

Lesson 2
Deploying a Distributed AD DS Environment

Some organizations need to deploy multiple domains or even multiple forests. Deploying AD DS domain controllers in this scenario is not much more complicated than deploying domain controllers in a single domain environment, but there are some special factors that you need to consider.

In this lesson, you will learn how to deploy a complex AD DS environment, and you will see how to upgrade from a previous version of AD DS.

Lesson Objectives

After completing this lesson, you will be able to:

- Explain how to install a domain controller in a new domain in a forest.
- Describe AD DS domain functional levels.
- Describe AD DS forest functional levels.
- Explain how to upgrade a previous version of AD DS to a Windows Server 2012 version.
- Explain how to migrate to Windows Server 2012 AD DS from a previous version.

Demonstration: Installing a Domain Controller in a New Domain in a Forest

In this demonstration, you will see how to:

- Configure an AD DS domain controller.
- Access the AD DS domain controller.

Demonstration Steps

Install the AD DS binaries on TOR-DC1

1. On TOR-DC1, in the Server Manager, use the **Add Roles and Features** Wizard to install the Active Directory Domain Services binaries.
2. Complete the AD DS **Add Roles and Features** Wizard using default settings.

Configure TOR-DC1 as an AD DS domain controller using the AD DS Installation Wizard

1. Use **Promote this server to a domain controller** to start the **Active Directory Domain Services Configuration Wizard.**
2. Use the **Active Directory Domain Services Configuration Wizard** to configure AD DS on TOR-DC1 with the following settings:
   - Deployment operation: **Add a new domain to an existing forest**
   - New domain name: **NA**
   - Directory Services Restore Mode (DSRM) password: **Pa$$w0rd**
3. Complete the **Active Directory Domain Services Configuration Wizard** with default settings.
4. Reboot and sign in as **NA\Administrator** with the password **Pa$$w0rd**, on the newly created AD DS domain controller **TOR-DC1**.

AD DS Domain Functional Levels

AD DS domains can run at different functional levels. Generally, upgrading the domain to a higher functional level will introduce additional features. Some of the domain functional levels are listed in the following table.

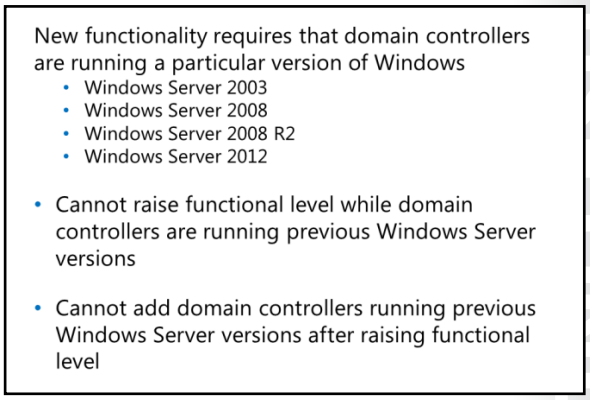

Domain functional level	Features
Microsoft Windows® 2000 Server native	• Universal groups • Group nesting • Group conversion, from security to distribution or vice versa • Security identifier (SID) history **Note:** Windows Server 2012 domain controllers cannot be installed in a domain running at the Windows 2000 Server native level.
Windows Server® 2003	• Netdom.exe. This domain management tool makes it possible to rename domain controllers. • **LastLogonTimestamp.** This attribute remembers the time of last domain logon for users, and replicates this to other AD DS domain controllers in the AD DS domain. • InetOrgPerson object support. The InetOrgPerson object is defined in Internet RFC 2798 and is used for federation with external directory services. • The ability to redirect the default location for user and computer objects. • Constrained delegation. This enables applications to take advantage of the secure delegation of user credentials by using Kerberos-based authentication. • Selective authentication. This allows you to specify the users and groups that are allowed to authenticate to specific resource servers in a trusting forest. • Application partitions, which are used to store information for AD integrated application. AD-integrated DNS uses an application partition, which allows the DNS partition to be replicated on domain controllers that are also DNS servers in the domain, or even across the forest.

Domain functional level	Features
Windows Server® 2008	• Distributed File System (DFS) Replication is available as a more efficient and robust file replication service for the SYSVOL folders. DFS Replication can replace the file replication service NT file replication service. • Additional interactive logon information is stored for each user, instead of just the last logon time. • Fine-grained password settings allow password and account lockout policies to be set for users and groups, which replaces the default domain settings for those users or group members. • Personal virtual desktops are available for users to connect to, by using RemoteApp and Remote Desktop. • Advanced Encryption Services (AES 128 and 256) support for Kerberos is available. • Read-only domain controllers (RODCs). These provide a secure and economical way to provide AD DS logon services in remote sites, without storing confidential information (such as passwords) in untrusted environments.
Windows Server® 2008 R2	• Authentication mechanism assurance, which packages information about a user's logon method, can be used in conjunction with application authentication—for example, with Active Directory Federation Services (AD FS). In another example, users who log on by using a smart card can be granted access to more resources than when they sign in with a username and password. • Automatic service principal name (SPN) management of managed services accounts. Managed service accounts allow account passwords to be managed by the Windows operating system.
Windows Server 2012	Windows Server 2012 domain functional level does not implement new features from Windows 2008 R2 functional level, with one exception: If the key distribution center (KDC) support for claims, compound authentication, and Kerberos armoring is configured for **Always provide claims** or **Fail unarmored authentication requests**, these functionalities will not be enabled until the domain is also set to Windows Server 2012 level.

Note: Generally, you cannot roll back AD DS domain functional levels. However, in Windows Server 2012 and Windows Server 2008 R2, you are able to roll back to a minimum of Windows Server 2008, as long as you do not have optional features (such as the Recycle Bin) enabled. If you have implemented a feature that is only available in a higher domain functional level, you cannot roll back to an earlier state.

Additional Reading: To learn more about the AD DS domain functional levels, see Understanding Active Directory Domain Services (AD DS) Functional Levels at http://go.microsoft.com/fwlink/?LinkId=270028.

AD DS Forest Functional Levels

The AD DS forest can run at different functional levels, and sometimes raising the AD DS forest functional level makes additional features available. The most noticeable additional features come with the upgrade to a Windows Server 2003 forest functional level. Additional features that are made available with Windows Server 2003 include:

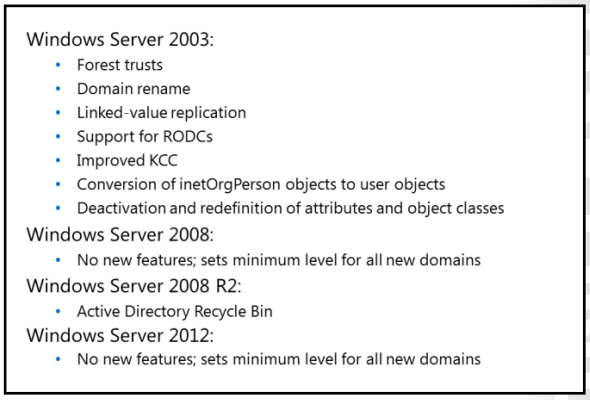

- Trusts. The basic feature of forests is that all domain trusts are transitive trusts, so that any user in any domain in the forest can access any resource in the forest, when given permission.
- Forest trusts. AD DS forests can have trusts set up between them, which enables resource sharing. There are full trusts and selective trusts.
- Linked-value replication. This feature improved Windows 2000 Server replication, and improved how group membership was handled. In previous versions of AD DS, the membership attribute of a group would be replicated as a single value. This meant that if two administrators changed the membership of the same group in two different instances of AD during the same replication period, the last write would win. The first changes made would be lost, because the new version of the group membership attribute would replace the previous one entirely. With linked-value replication, group membership is treated at the value level; therefore, all updates are merged together. This also greatly reduces the replication traffic that would occur. An additional benefit from this is the removal of the previous group membership restriction that limited the maximum number of members to 5,000.
- Improved AD DS replication calculation algorithms. Knowledge Consistency Checker (KCC) and intersite topology generator (ISTG) use improved algorithms to speed up the calculation of the AD DS replication infrastructure, and also provide much faster site link calculations.
- Support for RODCs. RODCs are supported at the Windows Server 2003 forest functional level. The RODC must be running Windows Server 2008 or newer, and requires at least one Windows Server 2008 or newer full domain controller as a replication partner.
- Conversion of inetOrgPerson objects to user objects. You can convert an instance of an inetOrgPerson object, used for compatibility with certain non-Microsoft directory services, into an instance of class user. You can also convert a user object to an inetOrgPerson object.
- Deactivation and redefinition of attributes and object classes. Although you cannot delete an attribute or object class in the schema at the Windows Server 2003 functional level, you can deactivate or redefine attributes or object classes.

The Windows Server 2008 forest functional level does not add new forest-wide features. The Windows Server 2008 R2 forest functional level adds the ability to activate AD features, such as the Active Directory Recycle Bin feature. This feature allows the ability to restore deleted Active Directory objects. The forest functional level cannot be rolled back if features requiring a certain forest level, such as the Active Directory Recycle Bin feature, have been enabled.

Although the Windows Server 2008 R2 AD DS forest functional level introduced AD DS Recycle Bin, the Recycle Bin had to be managed with Windows PowerShell®. However, the version of Remote Server Administration Tools (RSAT) that comes with Windows Server 2012 has the ability to manage the AD DS Recycle Bin by using graphical user interface (GUI) tools.

The Windows Server 2012 forest functional level does not provide any new forest-wide features. When you raise the forest functional level, you limit possible domain functional levels for domains that you add to the forest. For example, if you raise the forest functional level to Windows Server 2012, you cannot add a new domain running at Windows Server 2008 R2 domain functional level.

Upgrading a Previous Version of AD DS to Windows Server 2012 R2

To upgrade a previous version of AD DS to Windows Server 2012 AD DS, you can use either of the following two methods:

- Upgrade the operating system on the existing domain controllers to Windows Server 2012 R2.
- Introduce Windows Server 2012 R2 servers as domain controllers in the existing domain. You can then decommission AD DS domain controllers that are running earlier versions of AD DS.

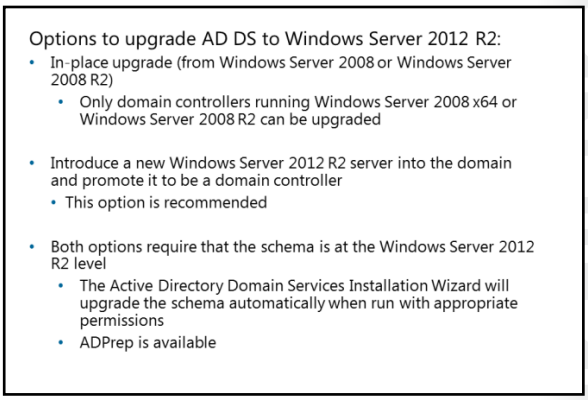

Of these two methods, the second is preferred, because upgrading operating systems—especially on servers that have been running for several years—is often difficult due to all the changes made through the years. By installing new domain controllers running Windows Server 2012 R2, you will have a clean installation of the Windows Server 2012 R2 operating system.

You can deploy Windows Server 2012 R2 servers as member servers in a domain with domain controllers running Windows Server 2003 or newer versions. However, before you can install the first domain controller that is running Windows Server 2012 R2, you must upgrade the schema. In versions of AD DS prior to Windows Server 2012 R2, you would run the adprep.exe tool to perform the schema upgrades. When you deploy new Windows Server 2012 R2 domain controllers in an existing domain, and if you are logged on with an account that is a member of the Schema Admins and Enterprise Admins groups, the Active Directory Domain Services Installation Wizard will automatically upgrade the AD DS forest schema.

Note: Windows Server 2012 R2 still provides a 64-bit version of ADPrep, so you can run Adprep.exe separately. For example, if the administrator who is installing the first Windows Server 2012 R2 domain controller is not a member of the Enterprise Admins group, you might need to run the command separately. You only have to run adprep.exe if you are planning an in-place upgrade for the first Windows Server 2012 R2 domain controller in the domain.

The Upgrade Process

To upgrade the operating system of a Windows Server 2008 domain controller to Windows Server 2012 R2, perform the following steps:

1. Insert the installation disk for Windows Server 2012 R2, and run **Setup**.
2. After the language selection page, click **Install now**.
3. After the operating system selection window and the license acceptance page, on the **Which type of installation do you want?** window, click **Upgrade: Install Windows and keep files, settings, and apps**.

With this type of upgrade, AD DS on the domain controller is upgraded to Windows Server 2012 AD DS. As a best practice, you should check for hardware and software compatibility before you perform an upgrade. Following the operating system upgrade, remember to update your drivers and other services (such as monitoring agents), and to check for updates for both Microsoft applications and non-Microsoft software.

Note: You can upgrade directly from Windows Server 2008 and Windows Server 2008 R2 to Windows Server 2012 R2. To upgrade servers that are running a version of Windows Server that is older than Windows Server 2008, you must either perform an interim upgrade to Windows Server 2008 or Windows Server 2008 R2, or perform a clean install. Note that Windows Server 2012 R2 AD DS domain controllers can coexist as domain controllers in the same domain as Windows Server 2003 domain controllers or newer.

The Clean Installation Process

To introduce a clean install of Windows Server 2012 R2 as a domain member, perform these steps:

1. Deploy and configure a new installation of Windows Server 2012, and then join it to the domain.
2. Promote the new server to be a domain controller in the domain by using Server Manager.

Migrating to Windows Server 2012 R2 AD DS from a Previous Version

As part of deploying AD DS, you might choose to restructure your environment for the following reasons:

- To optimize the logical AD DS structure. In some organizations, the business may have changed significantly since AD DS was first deployed. As a result, the AD DS domain or forest structure may no longer meet the business requirements.
- To assist in completing a business merger, acquisition, or divestiture.

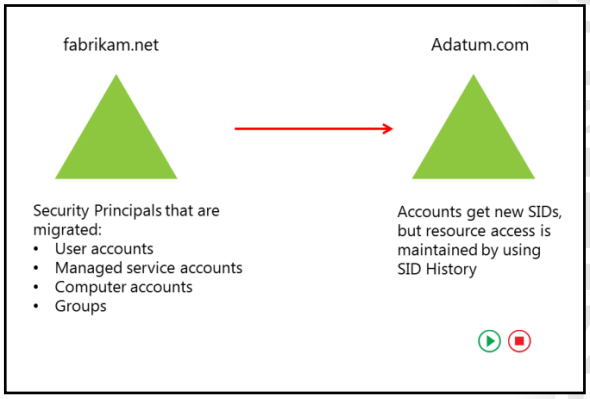

Restructuring involves the migration of resources between AD DS domains in either the same forest or in different forests. There is no option available in AD DS to detach a domain from one forest and then attach it to another forest. You can rename and rearrange domains within a forest under some circumstances, but there is no way to easily merge domains within or between forests. The only option for restructuring a domain in this way is to move all the accounts and resources from one domain to another.

You can use the Microsoft Active Directory Migration Tool (ADMT) to move user, group, and computer accounts from one domain to another, and to migrate server resources. If managed carefully, the migration can be completed without disrupting user access to the resources they need to do their work. ADMT provides both a GUI and a scripting interface, and supports the following tasks for completing the domain migration:

- User account migration
- Group account migration
- Computer account migration

- Service account migration
- Trust migration
- Exchange Server directory migration
- Security translation on migrated computer accounts
- Reporting features for viewing the migration's results
- Functionality to undo the last migration and retry the last migration

Pre-Migration Steps

Before performing the migration, you must perform several tasks to prepare the source and target domains. These tasks include:

- For domain member computers that are pre-Windows Vista® Service Pack 1 (SP1) or Windows Server 2008 R2, configure a registry on the target AD DS domain controller to allow cryptography algorithms that are compatible with the Microsoft Windows NT® Server 4.0 operating system.
- Enable firewall rules on source and target AD DS domain controllers, to allow file and printer sharing.
- Prepare the source and target AD DS domains to manage how the users, groups, and user profiles will be handled.
- Create a rollback plan.
- Establish the trust relationships that are required for the migration.
- Configure source and target AD DS domains to enable SID-History migration.
- Specify service accounts for the migration.
- Perform a test migration, and fix any errors that are reported.

Inter-forest Restructuring with ADMT

An inter-forest restructure involves moving resources from source domains that are in different forests than the target domain. To use ADMT to perform an inter-forest restructure, do the following:

1. Create a restructure plan. An adequate plan is critical to the success of the restructuring process. Complete the following steps to create your restructure plan:
    a. Determine the account-migration process.
    b. Assign object locations and location mapping.
    c. Develop a test plan.
    d. Create a rollback plan.
    e. Create a communication plan.
2. Prepare source and target domains. You must prepare both the source and target domains for the restructure process by performing the following tasks:
    a. Ensure 128-bit encryption on all domain controllers. Windows Server 2000 Service Pack 3 (SP3) and newer versions natively support 128-bit encryption. For older operating systems, you will need to download and install a separate encryption pack.
    b. Establish required trusts. You must configure at least a one-way trust between the source and target domains.
    c. Establish migration accounts. The ADMT uses migration accounts to migrate objects between source and target domains. Ensure that these accounts have permissions to move and modify objects on the source and target domains.

d. Determine whether ADMT will handle SID-History automatically, or if you must configure the target and source domains manually.

e. Ensure proper configuration of the target domain OU structure. Ensure that you configure the proper administrative rights and delegated administration in the target domain.

f. Install ADMT in the target domain.

g. Enable password migration.

h. Perform a test migration with a small test account group.

3. Migrate accounts. Perform the following steps to migrate accounts:

   a. Transition service accounts.

   b. Migrate global groups.

   c. Migrate accounts. Migrate user and computer accounts in batches to monitor the migration's progress. If you are migrating local profiles as part of the process, migrate the affected computers first, and then the associated user accounts.

4. Migrate resources. Migrate the remaining resources in the domain by performing the following steps:

   a. Migrate workstations and member servers.

   b. Migrate domain local groups.

   c. Migrate domain controllers.

5. Finalize migration. Finalize the migration and perform cleanup by performing the following steps:

   a. Transfer administration processes to the target domain.

   b. Ensure that at least two operable domain controllers exist in the target domain. Back up these domain controllers.

   c. Decommission the source domain.

The SID-History Attribute

During the migration, you may have moved user and group accounts to the new domain, but the resources that the users need to access may still be in the old domain. When you migrate a user account, AD DS assigns it a new SID. Because the resource in the source domain grants access based on the user SID from the source domain, the user cannot use the new SID to access the resource, until the resource is moved to the new domain.

To address this situation, you can configure the ADMT to migrate the SID from the source domain, and then store the SID in an attribute called **SID-History**. When the **SID**-**History** attribute is populated, the user's previous SID is used to grant access to resources in the source domain.

SID-History increases the size of the users' access token. After migrating the users to the new domain, the access control lists (ACLs) in your environment should be examined and ACLs migrated as well. Once a migration is complete, and the original domain has been removed, you should clean up your users' SID-History attribute. This task is best accomplished using the **Get-SIDHistroy** and **Remove-SIDHistory** PowerShell cmdlets. These activities should be carefully planned and executed as removing the SID-History before the environment is properly prepared could cause business interruptions.

To download the Active Directory Migration Tool version 3.2, go to:

http://go.microsoft.com/fwlink/?LinkId=270029

To download the Active Directory Migration Tool Guide, go to:

http://go.microsoft.com/fwlink/?LinkId=270045

Lesson 3
Configuring AD DS Trusts

AD DS trusts enable access to resources in a complex AD DS environment. When you deploy a single domain, you can easily grant access to resources within the domain to users and groups from the domain. When you implement multiple domains or forests, you need to ensure that the appropriate trusts are in place to enable the same access to resources. This lesson describes how trusts work in an AD DS environment, and how you can configure trusts to meet your business requirements.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the types of trusts that you can configure in a Windows Server 2012 environment.
- Explain how trusts work within an AD DS forest.
- Explain how trusts work between AD DS forests.
- Describe how to configure advanced trust settings.
- Describe how to configure a forest trust.

Overview of Different AD DS Trust Types

In a multi-domain AD DS forest, two-way transitive trust relationships are generated automatically between the AD DS domains, so that there is a path of trust between all of the AD DS domains. The trusts that are created automatically in the forest are all transitive trusts. That means that if domain A trusts domain B, and domain B trusts domain C, then domain A trusts domain C.

There are other types of trust that you can deploy. The following table describes the main trust types.

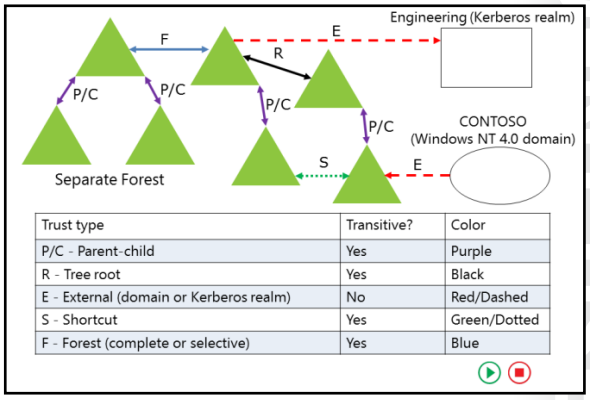

Trust type	Transitivity	Direction	Description
Parent and child	Transitive	Two-way	When a new AD DS domain is added to an existing AD DS tree, new parent and child trusts are created.
Tree-root	Transitive	Two-way	When a new AD DS tree is created in an existing AD DS forest, a new tree-root trust is created.
External	Non-transitive	One-way or two-way	External trusts enable resource access to be granted with a Windows NT 4.0 domain or an AD DS domain in another forest. These may also be set up to provide a framework for a migration.

Trust type	Transitivity	Direction	Description
Realm	Transitive or non-transitive	One-way or two-way	Realm trusts establish an authentication path between a Windows Server AD DS domain and a Kerberos v5 realm implemented using a directory service other than AD DS.
Forest (Complete or Selective)	Transitive	One-way or two-way	Trusts between AD DS forests allow two forests to share resources.
Shortcut	Non-transitive	One-way or two-way	Shortcut trusts can be configured to improve authentication times between AD DS domains that are in different parts of an AD DS forest. There are no shortcut trusts by default; they must be created by an administrator.

How Trusts Work Within a Forest

When you set up trusts between domains either within the same forest, across forests, or with an external realm, information about these trusts is stored in AD DS. A trusted domain object stores this information.

The trusted domain object stores information about the trust, such as the trust transitivity and type. Whenever you create a trust, a new trusted domain object is created and stored in the System container in AD DS.

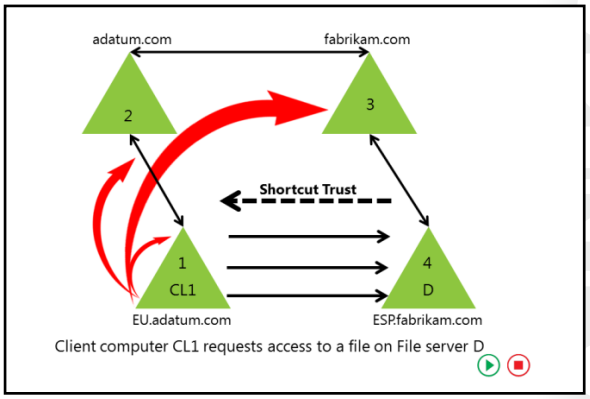

How Trusts Enable Users to Access Resources in a Forest

When the user in the domain attempts to access a shared resource in another domain in the forest, the user's computer first contacts a domain controller in its domain to request a session ticket to the resource. Because the resource is not in the user's domain, the domain controller needs to determine whether a trust exists with the target domain.

The domain controller can use the trust domain object to verify that the trust exists. However, to access the resource, the client computer must communicate with a domain controller in each domain along the trust path. The domain controller in the client computer's domain will refer the client computer to a domain controller in the next domain along the trust path. If that is not the domain where the resource is located, that domain controller will refer the client computer to a domain controller in the next domain. Eventually, the client computer will be referred to a domain controller in the domain where the resource is located, and the client will be issued a session ticket to access the resource.

The trust path is the shortest path through the trust hierarchy. In a forest with only the default trusts configured, the trust path will go up the domain tree to the forest root domain, and then down the domain tree to the target domain. If shortcut trusts are configured, the trust path may be a single hop from the client computer domain to the domain containing the resource.

How Trusts Work Between Forests

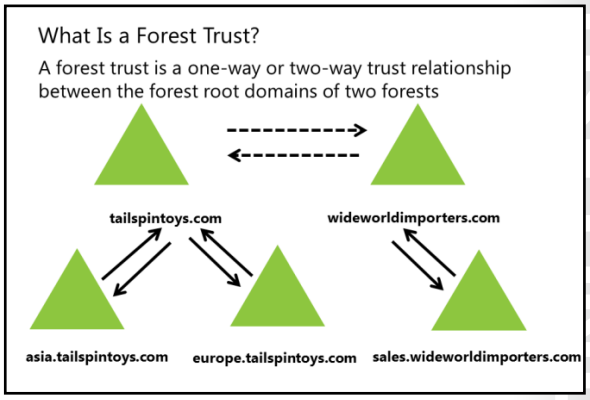

If the AD DS environment contains more than one forest, it is possible to set up trust relationships between the AD DS forest root domains. These forest trusts can be either forest-wide trusts or selective trusts. Forest trusts can be one-way or two-way. Forest trusts are also transitive for domains in each forest.

A forest trust relationship allows users who are authenticated by a domain in one forest to access resources that are in a domain in the other forest, provided they have been granted access rights. If the forest trust is one-way, domain controllers in the trusting forest can provide session tickets to users in any domain in the trusted forest. Forest trusts are significantly easier to establish, maintain, and administer than separate trust relationships between each of the domains in the forests.

Forest trusts are particularly useful in scenarios that involve cross-organization collaboration, or mergers and acquisitions, or within a single organization that has more than one forest in which to isolate Active Directory data and services. Forest trusts are also useful for application service providers, for collaborative business extranets, and for companies seeking a solution for administrative autonomy.

Forest trusts provide the following benefits:

- Simplified management of resources across two Windows Server 2008 (or newer version) forests, by reducing the number of external trusts necessary to share resources.
- Complete two-way trust relationships with every domain in each forest.
- Use of UPN authentication across two forests.
- Use of the Kerberos V5 protocol to improve the trustworthiness of authorization data that is transferred between forests.
- Flexibility of administration. Administrative tasks can be unique to each forest.

You can create a forest trust only between two AD DS forests, and you cannot extend the trust implicitly to a third forest. This means that if you create a forest trust between Forest 1 and Forest 2, and you create a forest trust between Forest 2 and Forest 3, Forest 1 does not have an implicit trust with Forest 3. Forest trusts are not transitive between multiple forests.

You must address several requirements before you can implement a forest trust, including ensuring that the forest functional level is Windows Server 2003 or newer, and that DNS name resolution exists between the forests.

Configuring Advanced AD DS Trust Settings

In some cases, trusts can present security issues. Additionally, if you do not configure a trust properly, users who belong to another domain can gain unwanted access to some resources. There are several technologies that you can use to help control and manage security in a trust.

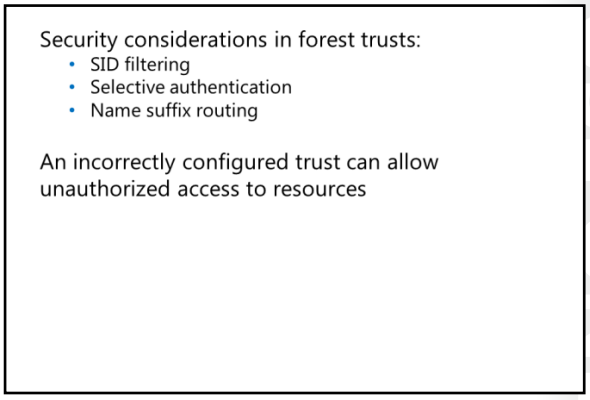

SID Filtering

By default, when you establish a forest or domain trust, you enable a domain quarantine, which is also known as SID filtering. When a user authenticates in a trusted domain, the user presents authorization data that includes the SIDs of all of the groups to which the user belongs. Additionally, the user's authorization data includes the SID-History of the user and the user's groups.

AD DS sets SID filtering by default to prevent users who have access at the domain or enterprise administrator level in a trusted forest or domain, from granting (to themselves or to other user accounts in their forest or domain) elevated user rights to a trusting forest or domain. SID filtering prevents misuse of the SID-History attribute, by only allowing reading the SID from the objectSID attribute and not the SID-History attribute.

In a trusted domain scenario, it is possible that an administrator could use administrative credentials in the trusted domain to load SIDs that are the same as SIDs of privileged accounts in your domain into the SID-History attribute of a user. That user would then have inappropriate levels of access to resources in your domain. SID filtering prevents this by enabling the trusting domain to filter out SIDs from the trusted domain that are not the primary SIDs of security principals. Each SID includes the SID of the originating domain, so that when a user from a trusted domain presents the list of the user's SIDs and the SIDs of the user's groups, SID filtering instructs the trusting domain to discard all SIDs without the domain SID of the trusted domain. SID filtering is enabled by default for all outgoing trusts to external domains and forests.

Selective Authentication

When you create an external trust or a forest trust, you can manage the scope of authentication of trusted security principals. There are two modes of authentication for an external or forest trust:

- Domain-wide authentication (for an external trust) or forest-wide authentication (for a forest trust)
- Selective authentication

If you choose domain-wide or forest-wide authentication, this enables all trusted users to authenticate for services and access on all computers in the trusting domain. Therefore, trusted users can be given permission to access resources anywhere in the trusting domain. If you use this authentication mode, all users from a trusted domain or forest are considered Authenticated Users in the trusting domain. Thus if you choose domain-wide or forest-wide authentication, any resource that has permissions granted to Authenticated Users is accessible immediately to trusted domain users.

If, however, you choose selective authentication, all users in the trusted domain are trusted identities. However, they are allowed to authenticate only for services on computers that you specify. When they use selective authentication, users will not become authenticated users in the target domain, however you can explicitly grant them the Allowed to Authenticate permission on specific computers.

For example, imagine that you have an external trust with a partner organization's domain. You want to ensure that only users from the partner organization's marketing group can access shared folders on only

one of your many file servers. You can configure selective authentication for the trust relationship, and then give the trusted users the right to authenticate only for that one file server.

Name Suffix Routing

Name suffix routing is a mechanism for managing how authentication requests are routed across forests running Windows Server 2003 or newer forests that are joined by forest trusts. To simplify the administration of authentication requests, when you create a forest trust, AD DS routes all unique name suffixes by default. A *unique name suffix* is a name suffix within a forest—such as a UPN suffix, SPN suffix, or DNS forest or domain tree name—that is not subordinate to any other name suffix. For example, the DNS forest name fabrikam.com is a unique name suffix within the fabrikam.com forest.

AD DS routes all names that are subordinate to unique name suffixes implicitly. For example, if your forest uses fabrikam.com as a unique name suffix, authentication requests for all child domains of fabrikam.com (childdomain.fabrikam.com) are routed, because the child domains are part of the fabrikam.com name suffix. Child names appear in the Active Directory Domains and Trusts snap-in. If you want to exclude members of a child domain from authenticating in the specified forest, you can disable name-suffix routing for that name. You also can disable routing for the forest name itself.

Demonstration: Configuring a Forest Trust

In this demonstration, you will see how to:

- Configure DNS name resolution by using a conditional forwarder.
- Configure a two-way selective forest trust.

Demonstration Steps

Configure DNS name resolution by using a conditional forwarder

- Configure DNS name resolution between adatum.com and treyresearch.net by creating a conditional forwarder so that LON-DC1 has a referral to TREY-DC1 as the DNS server for the DNS domain treyresearch.net.

Configure a two-way selective forest trust

- On LON-DC1, in Active Directory Domains and Trusts, create a two-way selective forest trust between adatum.com and treyresearch.net, by supplying the credentials of the treyresearch.net domain **Administrator** account.

Lab: Implementing Distributed AD DS Deployments

Scenario

A. Datum Corporation has deployed a single AD DS domain with all the domain controllers located in its London data center. As the company has grown and added branch offices with large numbers of users, it is becoming increasingly apparent that the current AD DS environment does not meet company requirements. The network team is concerned about the amount of AD DS–related network traffic that is crossing WAN links, which are becoming highly utilized.

The company has also become increasingly integrated with partner organizations, some of which need access to shared resources and applications that are located on the A. Datum internal network. The security department at A. Datum wants to ensure that the access for these external users is as secure as possible.

As one of the senior network administrators at A. Datum, you are responsible for implementing an AD DS infrastructure that will meet the company requirements. You are responsible for planning an AD DS domain and forest deployment that will provide optimal services for both internal and external users, while addressing the security requirements at A. Datum.

Objectives

After completing this lab, you will be able to:

- Implement child domains in AD DS.
- Implement forest trusts in AD DS.

Lab Setup

Estimated Time: 45 minutes

20412C-LON-DC1

20412C-TOR-DC1

20412C-LON-SVR2

20412C-TREY-DC1

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 the Hyper-V® Manager, click **20412C-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$$w0rd**
5. Repeat steps 2 through 4 for **20412C-LON-SVR2** and **20412C-TOR-DC1**.
6. Start **20412C-TREY-DC1**, and sign in as **Treyresearch\Administrator** with the password **Pa$$w0rd**.

Exercise 1: Implementing Child Domains in AD DS

Scenario

A. Datum has decided to deploy a new domain in the adatum.com forest for the North American region. The first domain controller will be deployed in Toronto, and the domain name will be na.adatum.com. You need to configure and install the new domain controller.

The main tasks for this exercise are as follows:

1. Install a domain controller in a child domain

2. Verify the default trust configuration

▶ Task 1: Install a domain controller in a child domain

1. On TOR-DC1, use the Server Manager to install the AD DS binaries.
2. When the AD DS binaries have installed, use the Active Directory Domain Services Configuration Wizard to install and configure TOR-DC1 as an AD DS domain controller for a new child domain named **na.adatum.com**.
3. When prompted, use **Pa$$w0rd** as the **Directory Services Restore Mode (DSRM)** password.

▶ Task 2: Verify the default trust configuration

1. Sign in to TOR-DC1 as **NA\Administrator** with the password **Pa$$w0rd**.
2. When the Server Manager opens, click **Local Server**. Verify that **Windows Firewall** shows **Domain: Off**. If it does not, then next to **Local Area Connection**, click **172.16.0.25, IPv6 enabled**. Right-click **Local Area Connection**, and then click **Disable**. Right-click **Local Area Connection**, and then click **Enable**. The Local Area Connection should now show **Adatum.com**.
3. From the Server Manager, launch the **Active Directory Domains and Trusts** management console, and verify the parent child trusts.

 **Note:** If you receive a message that the trust cannot be validated, or that the secure channel (SC) verification has failed, ensure that you have completed step 2, and then wait for at least 10 to 15 minutes. You can continue with the lab and come back later to verify this step.

Results: After completing this exercise, you will have implemented child domains in AD DS.

Exercise 2: Implementing Forest Trusts

Scenario

A. Datum is working on several high-priority projects with a partner organization named Trey Research. To simplify the process of enabling access to resources located in the two organizations, they have deployed a WAN between London and Munich, where Trey Research is located. You now need to implement and validate a forest trust between the two forests, and configure the trust to allow access to only selected servers in London.

The main tasks for this exercise are as follows:

- Configure stub zones for DNS name resolution.
- Configure a forest trust with selective authentication.
- Configure a server for selective authentication.
- To prepare for the next module.

▶ Task 1: Configure stub zones for DNS name resolution

1. On LON-DC1 using the DNS management console, configure a DNS stub zone for treyresearch.net.
2. Use **172.16.10.10** as the Master DNS server.
3. Close DNS Manager.
4. Sign in to TREY-DC1 as **TreyResearch\Administrator** with the password **Pa$$w0rd**.
5. Using the DNS management console, configure a DNS stub zone for adatum.com.
6. Use **172.16.0.10** as the Master DNS server.
7. Close DNS Manager.

▶ Task 2: Configure a forest trust with selective authentication

1. On LON-DC1, create a one-way outgoing trust between the treyresearch.net AD DS forest and the adatum.com forest. Configure the trust to use Selective authentication**.**
2. On LON-DC1, confirm and validate the trust from treyresearch.net.
3. Close Active Directory Domains and Trusts.

▶ Task 3: Configure a server for selective authentication

1. On LON-DC1, from the Server Manager, open **Active Directory Users and Computers**.
2. On LON-SVR2, configure the members of **TreyResearch\IT** group with the **Allowed to authenticate** permission. If you are prompted for credentials, type **TreyResearch\administrator** with the password **Pa$$w0rd**.
3. On LON-SVR2, create a shared folder named **IT-Data**, and grant **Read and Write** access to members of the **TreyResearch \IT** group. If you are prompted for credentials, type **TreyResearch\administrator** with the password **Pa$$w0rd**.
4. Sign out of TREY-DC1.
5. Sign in to TREY-DC1 as **TreyResearch\Alice** with the password **Pa$$w0rd**, and verify that you can access the shared folder on LON-SVR2.

▶ Task 4: To prepare for the next module

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

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20412C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20412C-TOR-DC1**, **20412C-TREY-DC1**, and **20412C-LON-SVR2**.

Results: After completing this exercise, you will have implemented forest trusts.

Question:

Why did you configure a delegated subdomain record in DNS on LON-DC1 before adding the child domain na.adatum.com?

Question: What are the alternatives to creating a delegated subdomain record in the previous question?

Question: When you are creating a forest trust, why would you create a selective trust instead of a complete trust?

Module Review and Takeaways

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
You receive error messages such as: DNS lookup failure, RPC server unavailable, domain does not exist, or domain controller could not be found.	
User cannot be authenticated to access resources on another AD DS domain or Kerberos realm.	

Module 5

Implementing Active Directory Domain Services Sites and Replication

Contents:

Module Overview

When you deploy Active Directory® Domain Services (AD DS), it is important that you provide an efficient logon infrastructure and a highly available directory service. Implementing multiple domain controllers throughout the infrastructure helps you meet both of these goals. However, you must ensure that AD DS replicates Active Directory information between each domain controller in the forest.

In this module, you will learn how AD DS replicates information between domain controllers within a single site and throughout multiple sites. You also will learn how to create multiple sites and monitor replication to help optimize AD DS replication and authentication traffic.

Objectives

After completing this module, you will be able to:

- Describe how AD DS replication works.
- Explain how to configure AD DS sites to help optimize authentication and replication traffic.
- Explain how to configure and monitor AD DS replication.

Lesson 1
AD DS Replication Overview

Within an AD DS infrastructure, standard domain controllers replicate Active Directory information by using a multimaster replication model. This means that if a change is made on one domain controller, that change then replicates to all other domain controllers in the domain, and potentially to all domain controllers throughout the entire forest. This lesson provides an overview of how AD DS replicates information between both standard and read-only domain controllers (RODCs).

Lesson Objectives

After completing this lesson, you will be able to:

- Describe AD DS partitions.
- Describe characteristics of AD DS replication.
- Explain how AD DS replication works within a site.
- Explain how to resolve replication conflicts.
- Explain how replication topology is generated.
- Explain how RODC replication works.
- Explain how SYSVOL replication works.

What Are AD DS Partitions?

The Active Directory data store contains information that AD DS distributes to all domain controllers throughout the forest infrastructure. Much of the information that the data store contains is distributed within a single domain. However, some information may be related to, and replicated throughout the entire forest regardless of the domain boundaries.

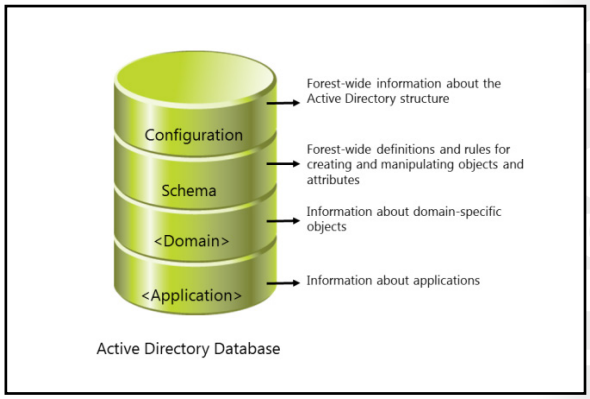

To help provide replication efficiency and scalability between domain controllers, the Active Directory data is separated logically into several partitions. Each partition is a unit of replication, and each partition has its own replication topology. The default partitions include the following:

- Configuration partition. The configuration partition is created automatically when you create the first domain controller in a forest. The configuration partition contains information about the forest-wide AD DS structure, including which domains and sites exist and which domain controllers exist in each domain. The configuration partition also stores information about forest-wide services such as Dynamic Host Configuration Protocol (DHCP) authorization and certificate templates. This partition replicates to all domain controllers in the forest. It is smaller than the other partitions, and its objects do not change frequently, therefore replication is also infrequent.
- Schema partition. The schema partition contains definitions of all the objects and attributes that you can create in the data store, and the rules for creating and manipulating them. Schema information replicates to all domain controllers in the forest. Therefore, all objects must comply with the schema

object and attribute definition rules. AD DS contains a default set of classes and attributes that you cannot modify. However, if you have Schema Admins group credentials, you can extend the schema by adding new attributes and classes to represent application-specific classes. Many applications such as Microsoft® Exchange Server and Microsoft® System Center 2012 Configuration Manager may extend the schema to provide application-specific configuration enhancements. These changes target the domain controller that contains the forest's schema master role. Only the schema master is permitted to make additions to classes and attributes. Similar to the configuration partition, the schema partition is small, and needs to replicate only when changes to the data stored there takes place, which does not happen frequently, except in those cases when the schema is extended.

- Domain partition. When you create a new domain, AD DS automatically creates and replicates an instance of the domain partition to all of the domain's domain controllers. The domain partition contains information about all domain-specific objects, including users, groups, computers, organizational units (OUs), and domain-related system settings. This is usually the largest of the AD DS partitions, as the objects stored here make up the bulk of the AD DS. Changes to this partition are fairly constant, as every time an object is created, deleted, or modified by changing an attribute's value, those changes must then be replicated. All objects in every domain partition in a forest are stored in the global catalog, with only a subset of their attribute values.
- Application partition. The application partition stores non-domain, application-related information that may have a tendency to be updated frequently or have a specified lifetime. An application is typically programed to determine how it stores, categorizes, and uses application-specific information that is stored in the Active Directory database. To prevent unnecessary replication of an application partition, you can designate which domain controllers in a forest will host the specific application's partition. Unlike a domain partition, an application partition does not store security principal objects, such as user accounts. Additionally, the global catalog does not store data that is contained in application partitions. The application partition's size and frequency of replication can vary widely, according to usage. Using Active Directory-integrated Domain Name System (DNS) with a large and robust DNS zone of many domain controllers, servers and client computers will result in the frequent replication of the partition.

Note: You can use the Active Directory Service Interfaces Editor (ADSI Edit) to connect to and view the partitions.

Characteristics of AD DS Replication

An effective AD DS replication design ensures that each partition on a domain controller is consistent with the replicas of that partition that are hosted on other domain controllers. Typically, not all domain controllers have exactly the same information in their replicas at any one moment because changes are occurring to the direction constantly. However, Active Directory replication ensures that all changes to a partition are transferred to all replicas of the partition. Active Directory replication balances accuracy, or integrity, and consistency (called *convergence*) with performance, this keeping replication traffic to a reasonable level.

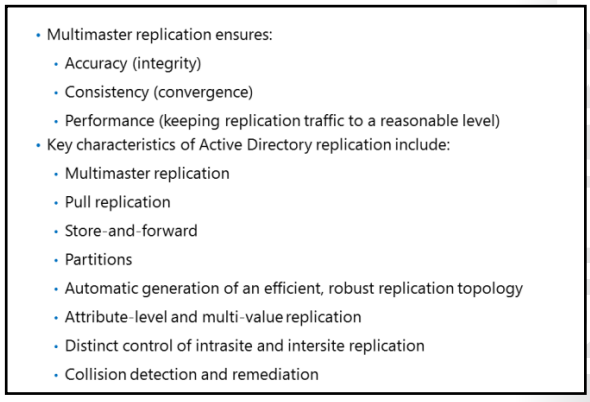

The key characteristics of Active Directory replication are:

- Multimaster replication. Any domain controller except an RODC can initiate and commit a change to AD DS. This provides fault tolerance, and eliminates dependency on a single domain controller to maintain the operations of the directory store.
- Pull replication. A domain controller requests, or *pulls* changes from other domain controllers. Even though a domain controller can notify its replication partners that it has changes to the directory, or poll its partners to see if they have changes to the directory, in the end, the target domain controller requests and pulls the changes themselves.
- Store-and-forward replication. A domain controller can pull changes from one partner, and then make those changes available to another partner. For example, domain controller B can pull changes initiated by domain controller A. Then, domain controller C can pull the changes from domain controller B. This helps balance the replication load for domains that contain several domain controllers.
- Data store partitioning. A domain's domain controllers host the domain-naming context for their domains, which helps minimize replication, particularly in multi-domain forests. The domain controllers also host copies of schema and configuration partitions, which are replicated forest wide. However, changes in configuration and schema partitions are much less frequent than in the domain partition. By default, other data, including application directory partitions and the partial attribute set (global catalog), do not replicate to every domain controller in the forest. You can enable replication to be universal by making all domain controllers in the forest global catalog.
- Automatic generation of an efficient and robust replication topology. By default, AD DS configures an effective, multi-way replication topology so that the loss of one domain controller does not impede replication. AD DS automatically updates this topology as domain controllers are added, removed, or moved between sites.
- Attribute-level replication. When an attribute of an object changes, only that attribute and minimal metadata describing that attribute replicates. The entire object does not replicate, except upon its initial creation. For multivalued attributes, such as account names in the *Member of* attribute of a group account, only changes to actual names are replicated, and not the entire list of names.
- Distinct control of intrasite replication and intersite replication. You can control replication within a single site and between sites.
- Collision detection and management. On rare occasions, you can modify an attribute on two different domain controllers during a single replication window. If this occurs, you must reconcile the two changes. AD DS has resolution algorithms that satisfy almost all scenarios.

How AD DS Replication Works Within a Site

AD DS replication within a single site, which takes place automatically, is called *intrasite replication*. However, you can also configure it to occur manually, as necessary. The following concepts are related to intrasite replication:

- Connection objects
- The knowledge consistency checker
- Notification
- Polling

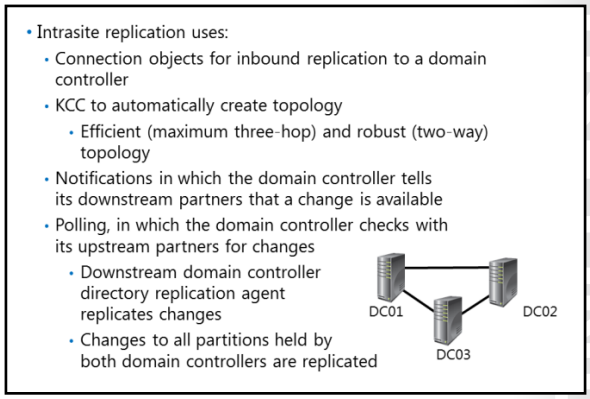

Connection Objects

A domain controller that replicates changes from another domain controller is called a replication partner. Replication partners are linked by connection objects. A connection object represents a replication path from one domain controller to another. Connection objects are one-way, representing inbound-only pull replication.

To view and configure connection objects, open Active Directory Sites and Services, and then select the NTDS Settings container of a domain controller's server object. You can force replication between two domain controllers by right-clicking the connection object, and then selecting Replicate Now. Note that replication is inbound-only, so if you want to replicate both domain controllers, you need to replicate the inbound connection object of each domain controller.

The Knowledge Consistency Checker

The replication paths that are built between domain controllers by connection objects create the forest's replication topology. You do not have to create the replication topology manually. By default, AD DS creates a topology that ensures effective replication. The topology is two-way, which means that if any one domain controller fails, replication continues uninterrupted. The topology also ensures that there are no more than three hops between any two domain controllers.

On each domain controller, a component of AD DS called the Knowledge Consistency Checker (KCC) helps generate and optimize the replication automatically between domain controllers within a site. The KCC evaluates the domain controllers in a site, and then creates connection objects to build the two-way, three-hop topology described earlier. If you add or remove a domain controller, or if a domain controller is not responsive, the KCC rearranges the topology dynamically, adding and deleting connection objects to rebuild an effective replication topology. The KCC runs at specified intervals (every 15 minutes by default) and designates replication routes between domain controllers that are the most favorable connections available at the time.

You can manually create connection objects to specify replication paths that should persist. However, creating a connection object manually is not typically required or recommended because the KCC does not verify or use the manual connection object for failover. The KCC will also not remove manual connection objects, which means that you must remember to delete connection objects that you create manually.

Notification

When a change is made to an Active Directory partition on a domain controller, the domain controller queues the change for replication to its partners. By default, the source server waits 15 seconds to notify its first replication partner of the change. Notification is the process by which an upstream partner informs its downstream partners that a change is available. By default, the source domain controller then waits three seconds between notifications to additional partners. These delays, called the *initial notification*

delay and the *subsequent notification delay*, are designed to stagger the network traffic that intrasite replication can cause.

Upon receiving the notification, the downstream partner requests the changes from the source domain controller, and the directory replication agent pulls the changes from the source domain controller. For example, suppose domain controller DC01 initializes a change to AD DS. When DC02 receives the change from DC01, it makes the change to its directory. DC02 then queues the change for replication to its own downstream partners.

Next, suppose DC03 is a downstream replication partner of DC02. After 15 seconds, DC02 notifies DC03 that it has a change. DC03 makes the replicated change to its directory, and then notifies its downstream partners. The change has made two hops, from DC01 to DC02, and then from DC02 to DC03. The replication topology ensures that no more than three hops occur before all domain controllers in the site receive the change. At approximately 15 seconds per hop, the change fully replicates in the site within one minute.

Polling

At times, a domain controller may not make any changes to its replicas for an extended time, particularly during off hours. Suppose this is the case with DC01. This means that DC02, its downstream replication partner, will not receive notifications from DC01. DC01 also might be offline, which would prevent it from sending notifications to DC02.

It is important for DC02 to know that its upstream partner is online and simply does not have any changes. This is achieved through a process called polling. During polling, the downstream replication partner contacts the upstream replication partner with queries as to whether any changes are queued for replication. By default, the polling interval for intrasite replication is once per hour. You can configure the polling frequency from a connection object's properties by clicking Change Schedule, although we do not recommend it.

If an upstream partner fails to respond to repeated polling queries, the downstream partner launches the KCC to check the replication topology. If the upstream server is indeed offline, the KCC rearranges the site's replication topology to accommodate the change.

Question: Describe the circumstances that result when you manually create a connection object between domain controllers within a site.

Resolving Replication Conflicts

Because AD DS supports a multimaster replication model, replication conflicts may occur. Typically, there are three types of replication conflicts that may occur in AD DS:

- Simultaneously modifying the same attribute value of the same object on two domain controllers.
- Adding or modifying the same object on one domain controller at the same time that the container object for the object is deleted on another domain controller.
- Adding objects with the same relative distinguished name into the same container on different domain controllers.

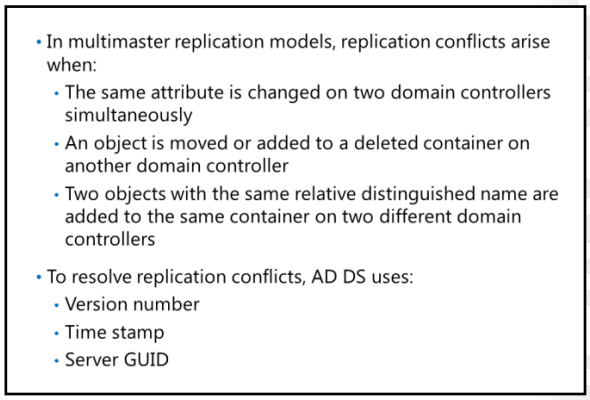

To help minimize conflicts, all domain controllers in the forest record and replicate object changes at the attribute or value level rather than at the object level. Therefore, changes to two different attributes of an object, such as the user's password and postal code, do not cause a conflict even if you change them at the same time from different locations.

When an originating update is applied to a domain controller, a stamp is created that travels with the update as it replicates to other domain controllers. The stamp contains the following components:

- Version number. The version number starts at one for each object attribute, and increases by one for each update. When performing an originating update, the version of the updated attribute is one number higher than the version of the attribute that is being overwritten.
- Timestamp. The timestamp is the update's originating time and date in the universal time zone, according to the system clock of the domain controller where the change is made.
- Server globally unique identifier (GUID). The server GUID identifies the domain controller that performed the originating update.

Common Replication Conflicts

The following table outlines several conflicts, and describes how AD DS resolves these issues:

Conflict	Resolution
Attribute value	If the version number value is the same, but the attribute value differs, then the timestamp is evaluated. The update operation that has the higher stamp value replaces the attribute value of the update operation with the lower stamp value. Certain multivalue attributes can be updated, such as a value in a group's Member of attribute, and will be processed as separate replicable events.
Add or move under a deleted container object, or the deletion of a container object	After resolution occurs at all replicas, AD DS deletes the container object, and the leaf object is made a child of the LostAndFound container. Stamps are not involved in this resolution.
Adding objects with the same relative distinguished name	The object with the later stamp keeps the relative distinguished name. AD DS assigns the sibling object a unique relative distinguished name by the domain controller. The name assignment is the relative distinguished name + CNF: + a reserved character (the asterisk,) + the object's GUID. This name assignment ensures that the generated name does not conflict with any other object's name.

How Replication Topology Is Generated

Replication topology is the route by which replication data travels through a network. To create a replication topology, AD DS must determine which domain controllers replicate data with other domain controllers. AD DS creates a replication topology based on the information that AD DS contains. Because each AD DS partition may be replicated to different domain controllers in a site, the replication topology can differ for schema, configuration, domain, and application partitions.

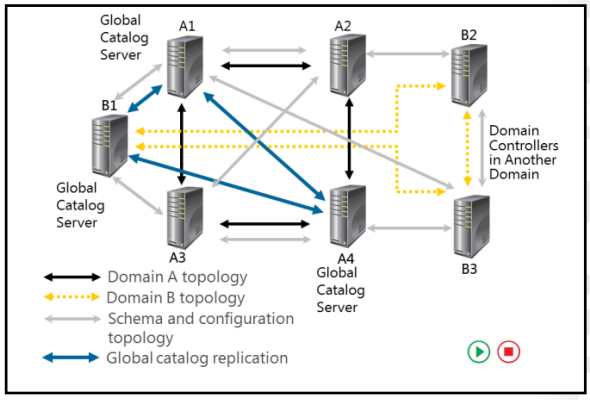

Because all domain controllers within a forest share schema and configuration partitions, AD DS replicates schema and configuration partitions to all domain controllers. Domain controllers in the same domain also replicate the domain partition. Additionally, domain controllers that host an application partition also replicate the application partition. To optimize replication traffic, a domain controller may have several replication partners for different partitions. In a single site, the replication topology will be fault tolerant and redundant. This means that if the site contains more than two domain controllers, each domain controller will have at least two replication partners for each AD DS partition.

How the Schema and Configuration Partitions Are Replicated

Replication of the schema and configuration partitions follows the same process as all other directory partitions. However, because these partitions are forest-wide rather than domain-wide, connection objects for these partitions may exist between any two domain controllers regardless of the domain controller's domain. Furthermore, the replication topology for these partitions includes all domain controllers in the forest.

How the Global Catalog Affects Replication

The configuration partition contains information about the site topology and other global data for all domains that are members of the forest. AD DS replicates the configuration partition to all domain controllers through normal forest-wide replication. Each global catalog server obtains domain information by contacting a domain controller for that domain and obtaining the partial replica information. Each global catalog server has full access to its own domain's domain partition, and therefore does not have to request a partial replication set of this information. The configuration partition also provides the domain controllers with a list of the forest's global catalog servers.

Global catalog servers register DNS service records in the DNS zone that corresponds to the forest root domain. These records, which are registered only in the forest root DNS zone, help clients and servers locate global catalog servers throughout the forest to provide client logon services.

How RODC Replication Works

As previously mentioned, domain controllers replicate data by pulling changes from other domain controllers. A RODC does not allow any non-replicated changes to be written to its database, and never replicates any information out to other domain controllers. Because changes are never written to an RODC directly, other domain controllers do not have to pull directory changes from an RODC. Restricting RODCs from originating changes prevents any changes or corruption that a malicious user or application might make from replicating to the rest of the forest.

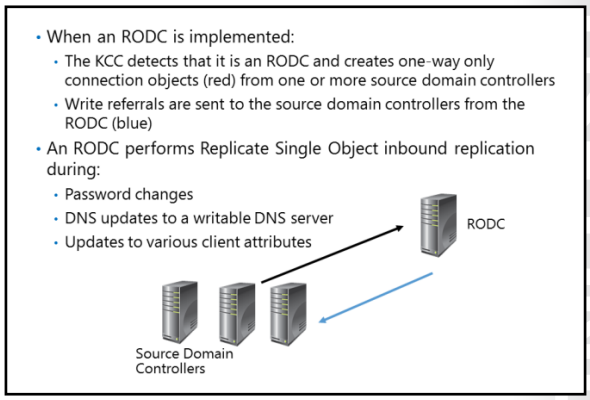

When a user or application attempts to perform a write request to a RODC, one of the following actions typically occurs:

- The RODC forwards the write request to a writable domain controller, which then replicates back to the RODC. Examples of this type of request include password changes, service principal name (SPN) updates, and computer\domain member attribute changes.
- The RODC responds to the client and provides a referral to a writable domain controller. The application can then communicate directly with a writable domain controller. Lightweight Directory Access Protocol (LDAP) is an example of acceptable RODC referrals.
- The write operation fails because it is not referred or forwarded to a writable domain controller. Remote procedure call (RPC) writes are an example of communication that may be prohibited from referrals or forwarding to another domain controller.

When you implement an RODC, the KCC detects that the domain controller is configured with a read-only replica of all applicable domain partitions. Because of this, the KCC creates one-way only connection objects from one or more source Windows Server® 2008 or newer Windows Server operating system domain controllers to the RODC.

For some tasks, an RODC performs inbound replication using a replicate-single-object operation. This is initiated on demand outside of the standard replication schedule. These tasks include:

- Password changes requests.
- DNS updates when a client is referred to a writable DNS server by the RODC. The RODC then attempts to pull the changes back using a replicate-single-object operation. This only occurs for Active Directory–integrated DNS zones.
- Updates for various client attributes including client name, DnsHostName, OsName, OsVersionInfo, supported encryption types, and the LastLogontimeStamp attribute.

How SYSVOL Replication Works

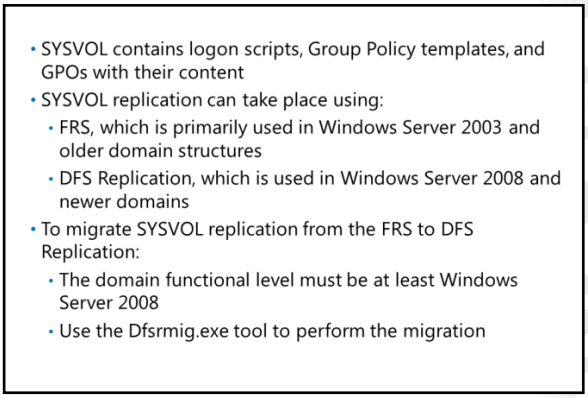

SYSVOL is a collection of files and folders on each domain controller that is linked to the %*SystemRoot*%\SYSVOL location. SYSVOL contains logon scripts and objects related to Group Policy such as Group Policy templates. The contents of the SYSVOL folder replicate to every domain controller in the domain using the connection object topology and schedule that the KCC creates.

Depending on the domain controller operating system version, the domain's functional level, and the migration status of SYSVOL, the File Replication Service (FRS) or Distributed File System (DFS) Replication replicates SYSVOL changes between domain controllers. The FRS was used primarily in Windows Server 2003 R2 and older domain structures. The FRS has limitations in both capacity and performance, which has led to the adoption of DFS Replication. The FRS is no longer available on domain controllers running Windows Server 2012 R2 when the domain is at the Windows Server® 2012 R2 domain functional level. If the forest functional level is Windows Server® 2008 R2 or newer, DFS Replication will be used. DFS database cloning is a new feature in Windows Server 2012 R2 DFS Replication but is not available for the AD DS SYSVOL. SYSVOL replication must use DFS initial replication.

In Windows Server 2008 and newer domains, you can use DFS Replication to replicate the contents of SYSVOL. DFS Replication supports replication scheduling and bandwidth throttling, and it uses a compression algorithm known as Remote Differential Compression (RDC). By using RDC, DFS Replication replicates only the differences or changes within files between the two servers, resulting in lower bandwidth use during replication. If any file that is stored in SYSVOL changes, DFS Replication will automatically replicate the file changes to the SYSVOL folders on the other domain controllers in the domain.

Note: You can use the dfsrmig.exe tool to migrate SYSVOL replication from the FRS to DFS Replication. For the migration to succeed, the domain functional level must be at least Windows Server 2008.

Lesson 2
Configuring AD DS Sites

Within a single site, AD DS replication occurs automatically without regard for network utilization. However, some organizations have multiple locations that are connected by wide area network (WAN) connections. If this is the case, you must ensure that AD DS replication does not impact network utilization negatively between locations. You also may need to localize network services to a specific location. For example, you may want users at a branch office to authenticate to a domain controller located in their local office, rather than over the WAN connection to a domain controller located in the main office. You can implement AD DS sites to help manage bandwidth over slow or unreliable network connections, and to assist in service localization for authentication and many other site-aware services on the network.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe AD DS sites.
- Explain why organizations might implement additional sites.
- Configure additional AD DS sites.
- Describe how AD DS replication works between sites.
- Describe the intersite topology generator.
- Describe optimizing domain controller coverage in multiple site scenarios.
- Describe how client computers locate domain controllers within sites.

What Are AD DS Sites?

To most administrators, a site is a physical location, such as an office or a city, typically separated by a WAN connection. These sites are physically connected by network links that might be as basic as dial-up connections or as sophisticated as fiber links. Together, the physical locations and links make up the physical network infrastructure.

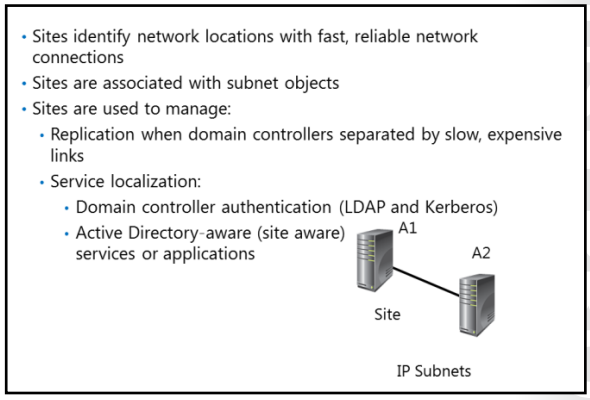

AD DS represents the physical network infrastructure with objects called *sites*. AD DS site objects are stored in the Configuration container (CN=Sites, CN=Configuration, DC=*forest root domain*) and are used to achieve two primary service management tasks:

- Manage replication traffic. Typically, there are two types of network connections within an enterprise environment: highly connected and less highly connected. Conceptually, a change made to AD DS should replicate immediately to other domain controllers within the highly connected network in which the change was made. However, you might not want the change to replicate to another site immediately if you have a slower, more expensive, or less reliable link. Instead, you might want to optimize performance, reduce costs, and manage bandwidth by managing replication over less highly connected segments of your enterprise. An Active Directory site represents a highly connected portion of your enterprise. When you define a site, the domain controllers within the site replicate

changes almost instantly. However, you can manage and schedule replication between sites as needed.

- Provide service localization. Active Directory sites help you localize services, including those provided by domain controllers. During logon, Windows clients are directed automatically to domain controllers in their sites. If domain controllers are not available in their sites, then they are directed to domain controllers in the nearest site that can authenticate the client efficiently. Many other services such as replicated DFS resources are also site-aware, to ensure that users are directed to a local copy of the resource.
- Group Policy Objects (GPOs) can be linked to a site. In that case, the site represents the top of the AD DS GPO hierarchy, and the AD DS GPO settings are applied here first.

What Are Subnet Objects?

Subnet objects identify the network addresses that map computers to AD DS sites. A subnet is a segment of a TCP/IP network to which a set of logical IP addresses are assigned. Because the subnet objects map to the physical network, so do the sites. A site can consist of one or more subnets. For example, if your network has three subnets in New York and two in London, you can create a site in New York and one in London, respectively, and then add the subnets to the respective sites.

Note: When you design your AD DS site configuration, it is critical that you correctly map IP subnets to sites. Likewise, if the underlying network configuration changes, you must ensure that these changes are updated to reflect the current IP subnet to site mapping. Domain controllers use the IP subnet information in AD DS to map client computers and servers to the correct AD DS site. If this mapping is not accurate, AD DS operations such as logon traffic and applying Group Policies are likely to happen across WAN links, and may get disrupted.

Default First Site

AD DS creates a default site when you install a forest's first domain controller. By default, this site is called *Default-First-Site-Name*. You can rename this site to a more descriptive name. When you install the forest's first domain controller, AD DS places it in the default site automatically. If you have a single site, it is not necessary to configure subnets or additional sites, because all machines will be covered by the Default-First-Site-Name default site. However, multiple sites need to have subnets associated with them as needed.

Why Implement Additional Sites?

Every Active Directory forest includes at least one site. You should create additional sites when:

- A slow link separates part of the network. As previously mentioned, a site is characterized by a location with fast, reliable, inexpensive connectivity. If two locations are connected by a slow link, you should configure each location as a separate AD DS site. A slow link typically is one that has a connection of less than 512 kilobits per second (Kbps).

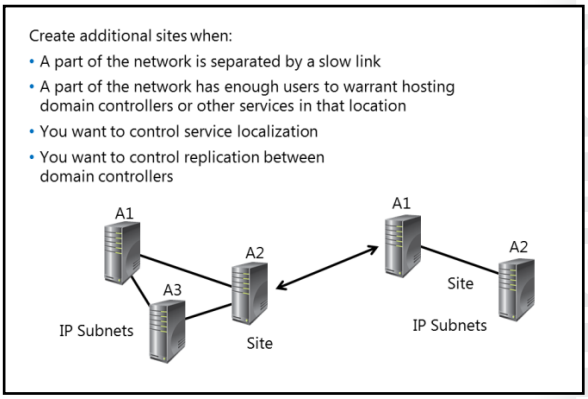

- A part of the network has enough users to warrant hosting domain controllers or other services in that location. Concentrations of users can also influence your site design. If a network location has a sufficient number of users, for whom the inability to authenticate would be problematic, place a domain controller in the location to support authentication within the location. After you place a domain controller or other distributed service in a location that will support those users, you might want to manage Active Directory replication to the location or localize service use by configuring an Active Directory site to represent the location.
- You want to control service localization. By establishing AD DS sites, you can ensure that clients use domain controllers that are nearest to them for authentication, which reduces authentication latency and traffic on WAN connections. In most scenarios, each site will contain a domain controller. However, you might configure sites to localize services other than authentication, such as DFS, Windows® BranchCache®, and Exchange Server services. In this case, some sites might be configured without a domain controller present in the site.
- You want to control replication between domain controllers. There may be scenarios in which two well-connected domain controllers are allowed to communicate only at certain times of the day. Creating sites allows you to control how and when replication takes place between domain controllers.

Demonstration: Configuring AD DS Sites

In this demonstration, you will see how to configure AD DS sites.

Demonstration Steps

1. From the Server Manager, open Active Directory Sites and Services.
2. Rename the **Default-First-Site-Name** site **LondonHQ** as needed.
3. Right-click the **Sites** node, and then click **New Site**. Specify the name **Toronto**, and then associate the new site with the default site link.
4. Create additional sites, as needed.
5. In the navigation pane, right-click **Subnets**, and then click **New Subnet**.
6. Provide the prefix **172.16.0.0/24**, and then associate the IP prefix to an available site object.
7. If required, move a domain controller to the new site.

How Replication Works Between Sites

The main characteristics of replication within a site are as follows:

- The network connections within a site are reliable, and have sufficient available bandwidth.

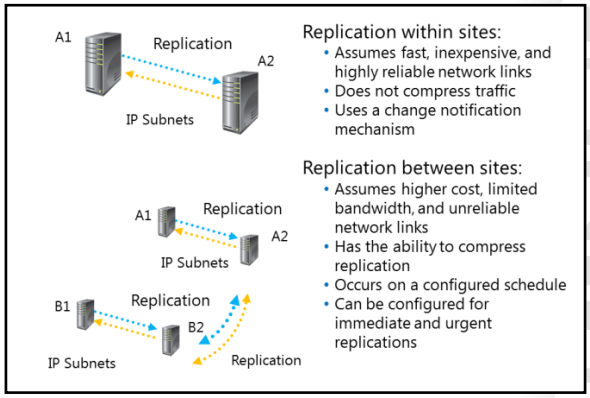

- Replication traffic within a site is not compressed, because a site assumes fast, highly reliable network connections. Not compressing replication traffic helps reduce the processing load on the domain controllers. However, uncompressed traffic may increase the network bandwidth.
- A change notification process initiates replication within a site.

The main characteristics of replication between sites are as follows:

- The network links between sites have limited available bandwidth, may have a higher cost, and may not be reliable.
- Replication traffic between sites can be designed to optimize bandwidth by compressing all replication traffic. Replication traffic is compressed to 10 to 15 percent of its original size before it is transmitted. Although compression optimizes network bandwidth, it imposes an additional processing load on domain controllers when it compresses and decompresses replication data.
- Replication between sites occurs automatically after you have defined configurable values, such as a schedule or a replication interval. You can schedule replication for inexpensive or off-peak hours. By default, changes are replicated between sites according to a schedule that you define, and not according to when changes occur. The schedule determines when replication can occur. The interval specifies how often domain controllers check for changes during the time that replication can occur.

Change Notifications Between AD DS Sites

By design, changes in AD DS replicate between domain controllers in different sites according to a defined replication schedule, and not according to when changes occur, such as with intrasite replication. Because of this, the replication latency in the forest can equal the sum of the greatest replication latencies along the longest replication path of any directory partition. In some scenarios, this can be inefficient.

To avoid latency in replication, you can configure change notifications on connections between sites. By modifying the site link object, you can enable change notification between sites for all connections that occur over that link. Because the replication partner across the site is notified of changes, the intersite replication interval is effectively ignored. The originating domain controller notifies the domain controller in the other site that it has a change, just as it does within a single site.

For changes such as account locks or similar security-related changes, immediate replication is crucial. In these situations, urgent replication is used. Urgent replication bypasses the notification delay and processes the change notifications immediately. This only affects change notifications. If you do not have change notifications enabled between sites, replication still honors the replication interval on the site link.

Note: When the user's password is changed, immediate replication is initiated to the primary domain controller (PDC) emulator operations master. This differs from urgent replication because it occurs immediately, without regard to the intersite replication interval.

What Is the Intersite Topology Generator?

When you configure multiple sites, the KCC on one domain controller in each site is designated as the site's intersite topology generator (ISTG). There is only one ISTG per site, regardless of how many domains or other directory partitions the site includes. ISTG is responsible for calculating the site's ideal replication topology across site links.

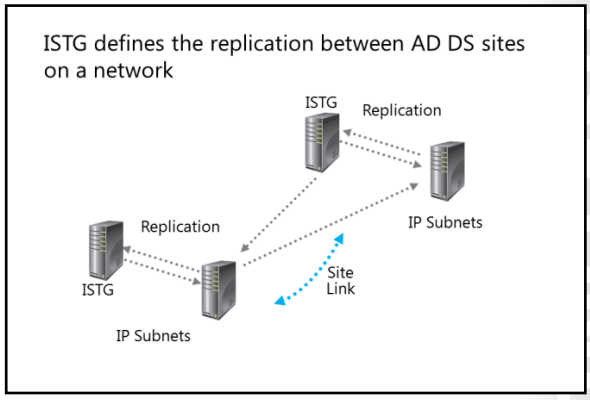

When you add a new site to the forest, each site's ISTG determines which directory partitions are present in the new site. The ISTG then calculates how many new connection objects are necessary to replicate the new site's required information.

In some networks, you might want to specify that only certain domain controllers are responsible for intersite replication. You can do this by specifying bridgehead servers. The bridgehead servers are responsible for all replication into, and out of, the site. ISTG creates the required connection agreement in its directory, and this information is then replicated to the bridgehead server. The bridgehead server then creates a replication connection with the bridgehead server in the remote site, and replication begins. If a replication partner becomes unavailable, the ITSG selects another domain controller automatically, if possible. If bridgehead servers have been assigned manually, and if they become unavailable, ISTG will not automatically select other servers.

The ISTG selects bridgehead servers automatically, and creates the intersite replication topology to ensure that changes replicate effectively between bridgeheads that share a site link. Bridgeheads are selected per partition, so it is possible that one domain controller in a site might be the bridgehead server for the schema, while another is for the configuration. However, you usually will find that one domain controller is the bridgehead server for all partitions in a site, unless there are domain controllers from other domains or application directory partitions. In this scenario, bridgeheads will be chosen for those partitions. Designated bridgehead servers are also useful when you have firewalls in between sites that only allow replication between specific domain controllers.

Optimizing Domain Controller Coverage in Multiple Site Scenarios

When you add a domain controller to a domain, the domain controller advertises its services by creating service (SRV) resource records (also known as *locator records*) in DNS. Unlike host A resource records, which map host names to IP addresses, SRV records map services to host names. For example, to publish its ability to provide authentication and directory access, a domain controller registers Kerberos v5 protocol and LDAP SRV records. These SRV records are added to several folders within the forest's DNS zones.

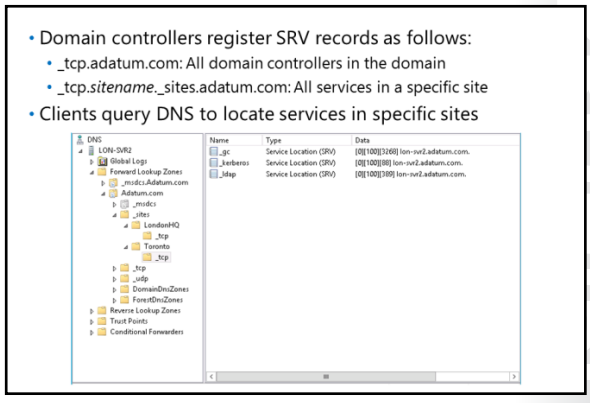

Within the domain zone, a folder called _tcp contains the SRV records for all domain controllers in the domain. Additionally, within the domain zone is a folder called _sites, which contains subfolders for each site configured in the domain. Each site-specific folder contains SRV records that represent services

available in the site. For example, if a domain controller is located in a site, a SRV record will be located at the path _sites*sitename*_tcp, where *sitename* is the name of the site.

A typical SRV record contains the following information:

- The service name and port. This portion of the SRV record indicates a service with a fixed port. It does not have to be a well-known port. SRV records in Windows Server 2012 include LDAP (port 389), Kerberos (port 88), Kerberos password protocol (KPASSWD, port 464), and global catalog services (port 3268).
- Protocol. The Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) is indicated as a transport protocol for the service. The same service can use both protocols in separate SRV records. Kerberos records, for example, are registered for both TCP and UDP. Microsoft clients use only TCP, but UNIX clients can use both UDP and TCP.
- Host name. The host name corresponds to the host A record for the server hosting the service. When a client queries for a service, the DNS server returns the SRV record and associated host A records, so the client does not need to submit a separate query to resolve the IP address of a service.

The service name in an SRV record follows the standard DNS hierarchy with components separated by dots. For example, a domain controller's Kerberos service is registered as: kerberos._tcp.*sitename*._sites.*domainName*, where:

- *domainName* is the domain or zone, for example contoso.com.
- _sites is all sites registered with DNS.
- *sitename* is the site of the domain controller registering the service.
- _tcp is any TCP-based services in the site.
- kerberos is a Kerberos Key Distribution Center that uses TCP as its transport protocol.

In certain situations, an organization might have computers in a location that does not have, nor would it be desirable to have, domain controllers. Sites can be created without domain controllers; however, as noted above, the site would not have a corresponding domain controller listing in the _sites\sitename_tcp path. In this case, there are several potential solutions. If, for example, maintenance of the domain controller and security of the AD DS database it contains are the main concerns, you could deploy RODCs. You can also use automatic site coverage. In the case of an empty site, a domain controller of the next closest site will automatically decide to take care of that site and also register its records for that site. This can also be adjusted or forced using group policy. Alternatively, if the site is well connected with only a few computers, you may want to avoid the costs of maintaining a server there. In this case, you could add the local subnet of the site to a central or a data center site location with multiple domain controllers. In the SRV record section example shown above, the client computers at the remote, domain controller-less location would be identified as belonging to the central site. The only time this would be a problem is if the central site's domain controllers were not available. In this case, the clients could use cached credentials to authenticate locally. Because automatic site link bridging, which will be discussed in the next lesson, is turned on by default, then domain authentication could still take place over the site link bridge where multiple sites exist.

How Client Computers Locate Domain Controllers Within Sites

When you join a Windows operating system client to a domain and then restart it, the client completes a domain controller location and registration process. The goal of this registration process is to locate the domain controller with the most efficient and closest location to the client's location based on IP subnet information.

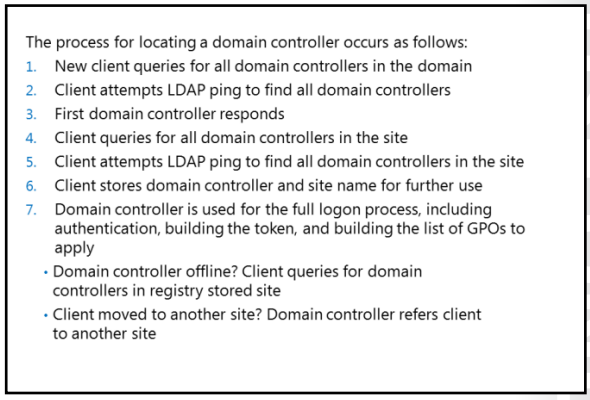

The process for locating a domain controller is as follows:

1. The new client queries for all domain controllers in the domain. As the new domain client restarts, it receives an IP address from a DHCP server, and is ready to authenticate to the domain. However, the client does not know where to find a domain controller. Therefore, the client queries for a domain controller by querying the _tcp folder, which contains the SRV records for all domain controllers in the domain.
2. The client attempts an LDAP ping to all domain controllers in a sequence. DNS returns a list of all matching domain controllers, and the client attempts to contact all of them on its first startup.
3. The first domain controller responds. The first domain controller that responds to the client examines the client's IP address, cross-references that address with subnet objects, and informs the client of the site to which the client belongs. The client stores the site name in its registry, and then queries for domain controllers in the site-specific _tcp folder.
4. The client queries for all domain controllers in the site. DNS returns a list of all domain controllers in the site.
5. The client attempts an LDAP ping sequentially to all domain controllers in the site. The domain controller that responds first authenticates the client.
6. The client forms an affinity. The client forms an affinity with the domain controller that responded first, and then attempts to authenticate with the same domain controller in the future. If the domain controller is unavailable, the client queries the site's _tcp folder again, and again attempts to bind with the first domain controller that responds in the site.

If the client moves to another site, which may be the case with a mobile computer, the client attempts to authenticate to its preferred domain controller. The domain controller notices that the client's IP address is associated with a different site, and then refers the client to the new site. The client then queries DNS for domain controllers in the local site

Automatic Site Coverage

As mentioned previously, you can configure sites to direct users to local copies of replicated resources, such as shared folders replicated within a DFS namespace. There may be scenarios in which you only require service localization with no need for a domain controller located within the site. In this case, a nearby domain controller will register its SRV records in the site by using a process called site coverage.

A site without a domain controller generally is covered by a domain controller in a site with the lowest site-link cost to the site that requires coverage. You also can configure site coverage and SRV record priority manually if you want to control authentication in sites without domain controllers.

For more information about how site coverage is evaluated, go to:
http://go.microsoft.com/fwlink/?LinkId=168550

Lesson 3
Configuring and Monitoring AD DS Replication

After you configure the sites that represent your network infrastructure, the next step is to determine if any additional site links are necessary to help manage AD DS replication. AD DS provides several options that you can configure to control how replication occurs over site links. You also need to understand the tools that you can use to monitor and manage replication in an AD DS network environment.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe AD DS site links.
- Explain the concept of site link bridging.
- Describe universal group membership caching.
- Describe how to manage intersite replication.
- Configure AD DS intersite replication.
- Describe best practices for deploying RODCs to support remote sites.
- Configure password replication policies.
- Describe tools used for monitoring and managing replication.

What Are AD DS Site Links?

For two sites to exchange replication data, a site link must connect them. A site link is a logical path that the KCC\ISTG uses to establish replication between sites. When you create additional sites, you must select at least one site link that will connect the new site to an existing site. Unless a site link is in place, the KCC cannot make connections between computers at different sites, nor can replication occur between sites.

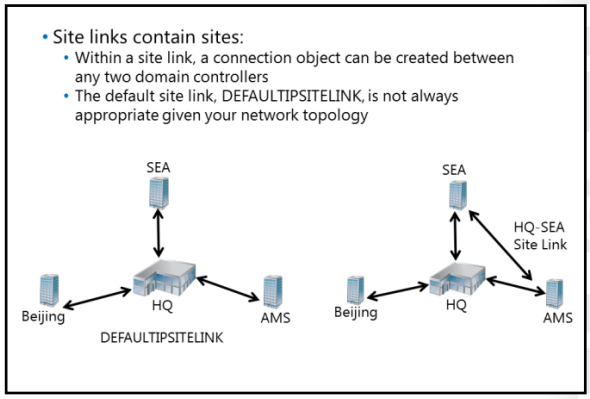

The important thing to remember about a site link is that it represents an available path for replication. A single site link does not control the network routes that are used. When you create a site link and add sites to it, you are telling AD DS that it can replicate between any of the sites associated with the site link. The ISTG creates connection objects, and those objects will determine the actual replication path. Although the replication topology that the ISTG builds does replicate AD DS effectively, it might not be efficient, given your network topology.

To understand this concept better, consider the following example. When you create a forest, one site link object is created: DEFAULTIPSITELINK. By default, each new site that you add is associated with the DEFAULTIPSITELINK. The DEFAULTIPSITELINK and any site links created have a default cost of 100 and a default replication period of 180 minutes. Consider an organization with a data center at the Headquarters and three branch offices. The three branch offices are each connected to the data center with a dedicated link. You create sites for each branch office: Seattle (SEA), Amsterdam (AMS), and Beijing. Each of the sites, including headquarters, is associated with the DEFAULTIPSITELINK site link object.

Because all four sites are on the same site link, you are instructing AD DS that all four sites can replicate with each other. That means that Seattle may replicate changes from Amsterdam; Amsterdam may replicate changes from Beijing; and Beijing may replicate changes from the headquarters, which in turn replicates changes from Seattle. In several of these replication paths, the replication traffic on the network flows from one branch through the headquarters on its way to another branch. With a single site link, you do not create a hub-and-spoke replication topology even though your network topology is hub-and-spoke.

To align your network topology with Active Directory replication, you must create specific site links. That is, you can manually create site links that reflect your intended replication topology. Continuing the preceding example, you would create three site links as follows:

- HQ-AMS includes the Headquarters and Amsterdam sites.
- HQ-SEA includes the Headquarters and Seattle sites.
- HQ-Beijing includes the Headquarters and Beijing sites.

After you create site links, the ISTG will use the topology to build an intersite replication topology that connects each site, and then creates connection objects automatically to configure the replication paths. As a best practice, you should set up your site topology correctly and avoid creating connection objects manually.

What Is Site Link Bridging?

After you have created site links and the ISTG generates connection objects to replicate partitions between domain controllers that share a site link, your work might be complete. In many environments, particularly those with straightforward network topologies, site links might be sufficient to manage intersite replication. In more complex networks, however, you can configure additional components and replication properties.

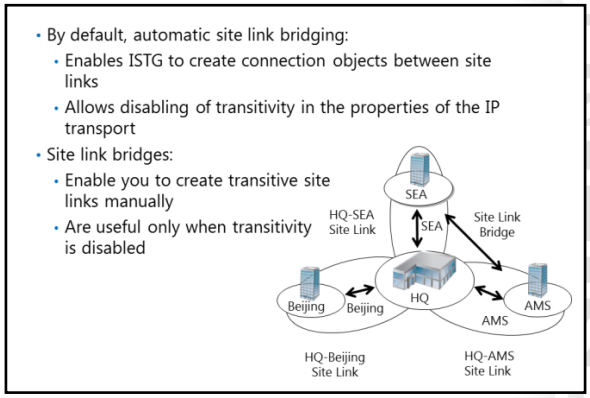

Automatic Site Link Bridging

By default, all site links are bridged. For example, if the Amsterdam and Headquarters sites are linked, and the Headquarters and Seattle sites are linked, then Amsterdam and Seattle are linked with a combined cost. This means, theoretically, that the ISTG could create a connection object directly between a domain controller in Seattle and a domain controller in Amsterdam, if a domain controller was not available at the headquarters for replication. This is accomplished by working around the hub-and-spoke network topology.

You can disable automatic site-link bridging by opening the properties of the IP transport in the Intersite Transports container, and then clearing the Bridge All Site Links check box. Before you do this in a production environment, read the technical resources about replication in the Windows Server technical libraries on the Microsoft TechNet website.

Site Link Bridges

A site link bridge connects two or more site links in a way that creates a transitive link. Site link bridges are necessary only when you have cleared the Bridge All Site Links check box for the transport protocol. Remember that automatic site-link bridging is enabled by default, which means that site link bridges are not required. However, you can keep automatic site link bridging enabled for the majority of the sites, but

also configure a site link bridge with in-between-costs. For example, suppose you have many sites, but branches A and B are both directly connected to the corporate headquarters with the default cost of 100. The corporate headquarters has a backup datacenter, HQ-HA, which is also connected with the cost of 100 between the corporate headquarters location and the locations of site A and B. In the event that all domain controllers are not available in HQ-HA, you want to ensure that site A can go to B. This is so you can keep site link bridging on but configure a site link bridge with the cost of 150 for A to B, this being greater than the cost of 100 for either to HQ-HA. Note that this is less than the cost without the site link bridge which would be 200, that is 100 from Site A to HQ-HA, plus the cost of 100 to HQ-HA to Site B. This makes the site link bridge cost of 150 an in-between-cost.

The figure on the previous slide illustrates how you can use a site link bridge in a forest in which automatic site-link bridging is disabled. By creating the site link bridge AMS-HQ-SEA, which includes the HQ-AMS and HQ-SEA site links, those two site links become transitive—or bridged. Therefore, a replication connection can be made between a domain controller in Amsterdam and a domain controller in Seattle.

What Is Universal Group Membership Caching?

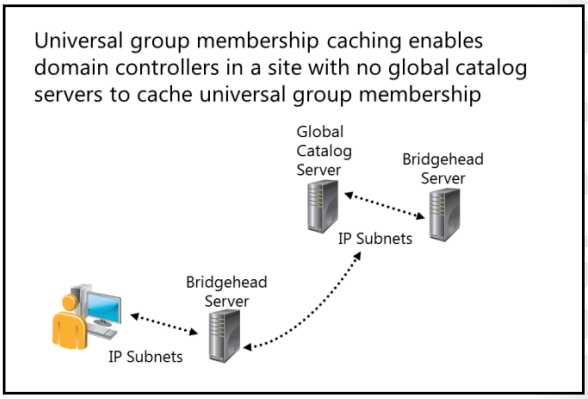

One of the issues that you may need to address when configuring AD DS replication is whether to deploy global catalog servers in each site. Because global catalog servers are required when users sign in to the domain, deploying a global catalog server in each site optimizes the user experience. However, if you deploy a global catalog server in a site, additional replication traffic might occur. That could be an issue if the network connection between AD DS sites has limited bandwidth and there are other domains with a large number of objects in the forest. In these scenarios, you can deploy domain controllers that are running Windows Server 2008 or newer, and then enable universal group membership caching for the site.

How Universal Group Membership Caching Works

A domain controller in a site that has enabled universal group membership caching stores the universal group information locally after a user attempts to log on for the first time. The domain controller obtains the user's universal group membership information from a global catalog server in another site. It then caches the information indefinitely, and periodically refreshes it. The next time that the user tries to sign in, the domain controller obtains the universal group membership information from its local cache without contacting a global catalog server.

By default, the universal group membership information contained in each domain controller's cache refreshes every eight hours. To refresh the cache, domain controllers send a universal group membership confirmation request to a designated global catalog server.

You can configure universal group membership caching from the properties of the NTDS Site Settings node.

Managing Intersite Replication

When you create a site link, you can use several configuration options to help manage intersite replication. These options include:

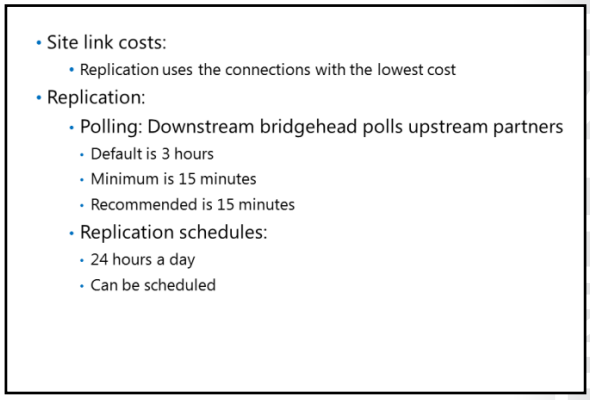

- Site link costs. Site link costs manage the flow of replication traffic when there is more than one route for replication traffic. You can configure site link costs to indicate that a particular link meets one or more requirements/conditions; for example, it might be faster or more reliable, or it might be preferred. Higher costs are used for slow links, and lower costs are used for fast links. AD DS replicates by using the connection with the lowest cost. By default, all site links are configured with a cost of 100.
- Replication frequency. Intersite replication is based only on polling. By default, every three hours a replication partner polls its upstream replication partners to determine whether changes are available. This replication interval may be too long for organizations that want directory changes to replicate more quickly. You can change the polling interval by accessing the properties of the site link object. The minimum polling interval is 15 minutes.
- Replication schedules. By default, replication occurs 24 hours a day. However, you can restrict intersite replication to specific times by changing the schedule attributes of a site link.

Demonstration: Configuring AD DS Intersite Replication

In this demonstration, you will see how to configure AD DS intersite replication.

Demonstration Steps

1. From the Server Manager, open **Active Directory Sites and Services**.
2. Rename the DEFAULTIPSITELINK as **LON-TOR**.
3. Right-click the site link, and then click **Properties**.
4. Modify the Cost, Replication interval, and Schedule as needed.
5. If necessary, open the properties of the IP node, and then modify the **Bridge all site links** option.

Best Practices When Deploying RODCs to Support Remote Sites

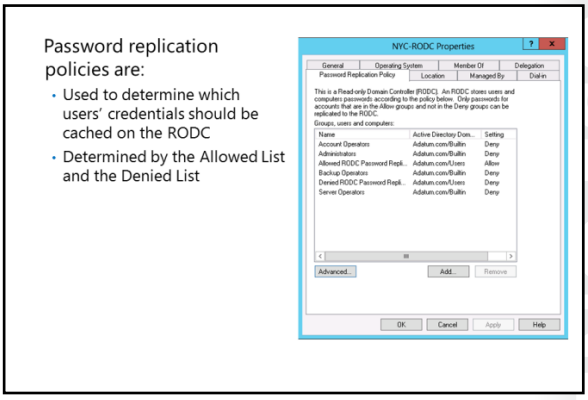

RODCs have unique AD DS replication requirements related to cached user and computer credentials. They use password replication policies to determine which users' or computers' credentials might be cached on the server. If a password replication policy allows an RODC to cache a user's or computer's credentials, the RODC can process the authentication and service-ticket activities. To address this issue, you can create one allowed list that contains both user and computer credentials. If a credential is not allowed to be cached on the RODC, then the RODC refers the authentication and service-ticket activities to a writable domain controller. If the computer accounts in that site are not allowed to authenticate and the WAN link is offline, no user can authenticate since the user can only authenticate to the domain after the client computer has been authenticated.

To access the password replication policy, open the properties of the RODC in the domain controllers organizational unit (OU), and then click the Password Replication Policy tab. An RODC's password replication policy is determined by two multivalued attributes of the RODC's computer account. These attributes are known commonly as the *Allowed List* and the *Denied List*. If a user, computer, or service account is on the Allowed List, the credentials are cached after the user logs on via the RODC. You can include groups on the Allowed List, in which case all users or computers who belong to the group can have their credentials cached on the RODC. If the account is on both the Allowed List and the Denied List, the RODC does not cache the credentials. The Denied List takes precedence.

To facilitate the management of password replication policy, two domain local security groups are created in the AD DS Users container. The first security group, the Allowed RODC Password Replication Group, is added to the Allowed List for each new RODC. By default, this group has no members. Therefore, by default, a new RODC will not cache any credentials. If there are accounts with credentials that you want to be cached by all domain RODCs, you add those users to the Allowed RODC Password Replication Group. As a best practice, you can create one Allow List per site, and configure only the accounts assigned to that site in the Allow List.

The second group, the Denied RODC Password Replication Group, is added to the Denied List for each new RODC. If you want to ensure that domain RODCs never cache certain credentials, you can add those accounts to the Denied RODC Password Replication Group. By default, this group contains security-sensitive accounts that are members of groups, including Domain Admins, Enterprise Admins, Schema Admins, Cert Publishers, and Group Policy Creator Owners.

Demonstration: Configuring Password Replication Policies

In this demonstration, you will see how to configure password replication policies.

Demonstration Steps

1. Run **Active Directory Users and Computers**.
2. Pre-create an RODC computer object named **LON-RODC1**.
3. In the **Domain Controllers** organizational unit (OU), open the properties of **LON-RODC1**.

4. Click the **Password Replication Policy** tab, and view the default policy.
5. Close the **LON-RODC1 Properties** dialog box.
6. In the **Active Directory Users and Computers** console, click the **Users** container.
7. Double-click **Allowed RODC Password Replication Group**, then click the **Members** tab and examine the default membership of **Allowed RODC Password Replication Group**. There should be no members by default.
8. Click **OK**.
9. Double-click **Denied RODC Password Replication Group**, and then click the **Members** tab.
10. Click **Cancel** to close the **Denied RODC Password Replication Group Properties** dialog box.

Tools for Monitoring and Managing Replication

After you have implemented your replication configuration, you must be able to monitor replication for ongoing support, optimization, and troubleshooting. Two tools are particularly useful for reporting and analyzing replication: the Replication Diagnostics tool, Repadmin.exe, and the Directory Server Diagnosis, Dcdiag.exe, tool.

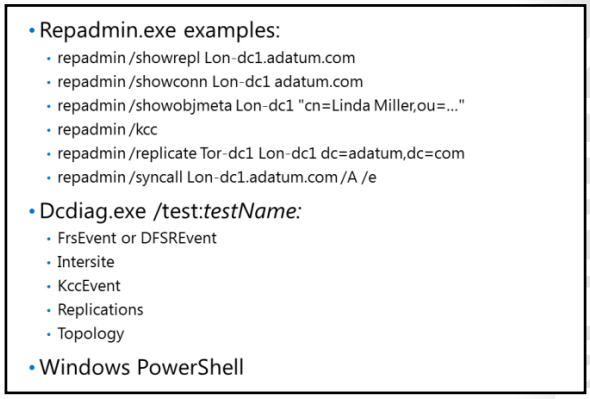

The Replication Diagnostics Tool

The Replication Diagnostics tool, Repadmin.exe, is a command-line tool that enables you to report the status of replication on each domain controller. The information that Repadmin.exe produces can help you spot a potential problem with replication in the forest. You can view levels of detail down to the replication metadata for specific objects and attributes, enabling you to identify where and when a problematic change was made to AD DS. You can even use Repadmin.exe to create the replication topology and force replication between domain controllers.

Repadmin.exe supports a number of commands that perform specific tasks. You can learn about each command by typing **repadmin /?:command** at a command line. Most commands require arguments. Many commands take a *DC_LIST* parameter, which is simply a network label (DNS, NetBIOS name, or IP address) of a domain controller. Some of the replication monitoring tasks you can perform by using Repadmin.exe are:

- Display the replication partners for a domain controller. To display the replication connections of a domain controller, type **repadmin /showrepl *DC_LIST***. By default, Repadmin.exe shows only intersite connections. Add the **/repsto** argument to see intersite connections, as well.
- Display connection objects for a domain controller. Type **repadmin /showconn *DC_LIST*** to show the connection objects for a domain controller.
- Display metadata about an object, its attributes, and replication. You can learn much about replication by examining an object on two different domain controllers to find out which attributes have or have not replicated. Type **repadmin /showobjmeta *DC_LIST Object***, where *DC_LIST* indicates the domain controller(s) to query. You can use an asterisk to indicate all domain controllers. *Object* is a unique identifier for the object, its distinguished name or GUID, for example.

You can also make changes to your replication infrastructure by using the Repadmin.exe tool. Some of the management tasks you can perform are:

- Launching the KCC. Type **repadmin /kcc** to force the KCC to recalculate the inbound replication topology for the server.
- Forcing replication between two partners. You can use Repadmin.exe to force replication of a partition between a source and a target domain controller. Type **repadmin /replicate *Destination_DC_LIST Source_DC_Name Naming_Context***.
- Synchronizing a domain controller with all replication partners. Type **repadmin /syncall *DC*/A /e** to synchronize a domain controller with all its partners, including those in other sites.

The Directory Server Diagnosis Tool

The Directory Service Diagnosis tool, Dcdiag.exe, performs a number of tests and reports on the overall health of replication and security for AD DS. Run by itself, dcdiag.exe performs summary tests and then reports the results. At the other extreme, **dcdiag.exe /c** performs almost every test. The output of tests can be redirected to files of various types, including XML. Type **dcdiag /?** for full usage information.

You can also specify one or more tests to perform using the **/test:*Test Name*** parameter. Tests that are directly related to replication include:

- FrsEvent. Reports any operation errors in the FRS.
- DFSREvent. Reports any operation errors in the DFS Replication system.
- Intersite. Checks for failures that would prevent or delay intersite replication.
- KccEvent. Identifies errors in the KCC.
- Replications. Checks for timely replication between domain controllers.
- Topology. Checks that the replication topology is connected fully for all domain controllers.
- VerifyReplicas. Verifies that all application directory partitions are instantiated fully on all domain controllers that are hosting replicas.

Monitoring Replication with Microsoft System Center 2012 Operations Manager

You can install the Active Directory Domain Services Management Pack for Operations Manager on the domain controller. This management pack contains many alerts, views, tasks, and reports for a variety of AD DS functions, including replication.

The section on replication monitoring collects replication performance data to include AD DS replication alerts, intersite replication, replication latency, and both inbound and outbound replication traffic bytes per second. The management pack also contains several replication topology diagrams that cover site links, connection objects and broken connection objects. It also contains reports on replication connection objects, replication site links, replication bandwidth, and replication latency.

There are primarily four replication areas that Operations Manager monitors as part of the management pack:

- Operations master consistency check. This critical part of replication allows replication partners to be in agreement on which domain controllers are in an operations master role.
- Replication latency monitoring. This ensures that changes to the AD DS are replicated in a timely manner, and can periodically send replication events of its own to ensure that all replication partners are functioning properly.
- Replication partner count. This keeps track of how many replication partners a domain controller has. If the number is either below or above a particular threshold, it will trigger an alert.

- Replication provider. This monitors and reports on all replication links for each domain controller. You use Windows Management Instrumentation (WMI) to find link status.

New Windows PowerShell cmdlets for AD DS Replication

Windows Server 2012 has added several new Windows PowerShell® cmdlets to create, configure, and monitor AD DS replication. The following table lists these cmdlets:

Cmdlet	Data returned
Get-ADReplicationConnection	Specific Active Directory replication connection or set of Active Directory replication connection objects based on a specified filter
Get-ADReplicationFailure	Description of an Active Directory replication failure
Get-ADReplicationPartnerMetadata	Replication metadata for a set of one or more replication partners
Get-ADReplicationSite	Specific Active Directory replication site or a set of replication site objects based on a specified filter
Get-ADReplicationSiteLink	Specific Active Directory site link or a set of site links based on a specified filter
Get-ADReplicationSiteLinkBridge	Specific Active Directory site link bridge or a set of site link bridge objects based on a specified filter
Get-ADReplicationSubnet	Specific Active Directory subnet or a set of Active Directory subnets based on a specified filter

For more Windows PowerShell AD DS replication cmdlets, go to:
http://go.microsoft.com/fwlink/?LinkID=386640

Lab: Implementing AD DS Sites and Replication

Scenario

A. Datum Corporation has deployed a single AD DS domain with all the domain controllers located in the London data center. As the company has grown and added branch offices with large numbers of users, it has become apparent that the current AD DS environment does not meet the company requirements. Users in some of the branch offices report that it can take a long time for them to sign in on their computers. Access to network resources such as the company's Microsoft Exchange® 2013 servers and the Microsoft SharePoint® servers can be slow, and they fail sporadically.

As one of the senior network administrators, you are responsible for planning and implementing an AD DS infrastructure that will help address the business requirements for the organization. You are responsible for configuring AD DS sites and replication to optimize the user experience and network utilization within the organization.

Objectives

After completing this lab, you will be able to:

- Modify the default site created in AD DS.
- Create and configure additional sites and subnets.
- Configure AD DS replication.
- Monitor and troubleshoot AD DS replication.

Lab Setup

Estimated Time: 30 minutes

Virtual machines	20412C-LON-DC1 20412C-TOR-DC1
User Name	**Adatum\Administrator**
Password	**Pa$$w0rd**

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

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V® Manager, click **20412C-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 two through four for 20412C-TOR-DC1.

Exercise 1: Modifying the Default Site

Scenario

A. Datum Corporation has decided to implement additional AD DS sites to optimize the network utilization for AD DS network traffic. The first step in implementing the new environment is to install a new domain controller for the Toronto site. You then will reconfigure the default site and assign appropriate IP address subnets to the site.

Finally, you have been asked to change the name of the default site to LondonHQ and associate it with the IP subnet 172.16.0.0/24, which is the subnet range used for the London head office.

The main tasks for this exercise are as follows:

1. Install the Toronto domain controller
2. Rename the default site
3. Configure IP subnets associated with the default site

▶ Task 1: Install the Toronto domain controller

1. On TOR-DC1, use Server Manager to install **Active Directory Domain Services**.
2. When the AD DS binaries have installed, use the Active Directory Domain Services Configuration Wizard to install and configure TOR-DC1 as an additional domain controller for Adatum.com.
3. After the server restarts, sign in as **Adatum\Administrator** with the password of **Pa$$w0rd**.

▶ Task 2: Rename the default site

1. If necessary, on LON-DC1, open the Server Manager console.
2. Open Active Directory Sites and Services, and then rename the **Default-First-Site-Name** site to **LondonHQ**.
3. Verify that both LON-DC1 and TOR-DC1 are members of the **LondonHQ** site.

▶ Task 3: Configure IP subnets associated with the default site

1. If necessary, on LON-DC1, open the Server Manager console, and then open Active Directory Sites and Services.
2. Create a new subnet with the following configuration:
   - Prefix: **172.16.0.0/24**
   - Site object: **LondonHQ**

Results: After completing this exercise, you will have reconfigured the default site and assigned IP address subnets to the site.

Exercise 2: Creating Additional Sites and Subnets

Scenario

The next step you take to implement the AD DS site design is to configure the new AD DS site. The first site that you need to implement is the Toronto site for the North American data center. The network team in Toronto would also like to dedicate a site called TestSite in the Toronto data center. You have been instructed that the Toronto IP subnet address is 172.16.1.0/24, and the test network IP subnet address is 172.16.100.0/24.

The main tasks for this exercise are as follows:

1. Create the AD DS sites for Toronto
2. Create IP subnets associated with the Toronto sites

▶ Task 1: Create the AD DS sites for Toronto

1. If necessary, on LON-DC1, open the Server Manager console, and then open Active Directory Sites and Services.
2. Create a new site with the following configuration:
   - Name: **Toronto**
   - Site link object: **DEFAULTIPSITELINK**
3. Create another new site with the following configuration:
   - Name: **TestSite**
   - Site link object: **DEFAULTIPSITELINK**

▶ Task 2: Create IP subnets associated with the Toronto sites

1. If necessary, on LON-DC1, open Active Directory Sites and Services.
2. Create a new subnet with the following configuration:
   - Prefix: **172.16.1.0/24**
   - Site object: **Toronto**
3. Create another new subnet with the following configuration:
   - Prefix: **172.16.100.0/24**
   - Site object: **TestSite**
4. In the navigation pane, click the **Subnets** folder. Verify in the details pane that the three subnets are created and associated with their appropriate site.

Results: After this exercise, you will have created two additional sites representing the IP subnet addresses located in Toronto.

Exercise 3: Configuring AD DS Replication

Scenario

Now that the AD DS sites have been configured for Toronto, your next step is to configure the site links to manage replication between the sites, and then to move the TOR-DC1 domain controller to the Toronto site. Currently, all sites belong to DEFAULTIPSITELINK.

You need to modify site linking so that LondonHQ and Toronto belong to one common site link called LON-TOR. You should configure this link to replicate every hour. Additionally, you should link the TestSite site only to the Toronto site using a site link named TOR-TEST. Replication should not be available from the Toronto site to the TestSite during the working hours of 9 A.M. to 3 P.M. You then will use tools to monitor replication between the sites.

The main tasks for this exercise are as follows:

1. Configure site links between AD DS sites
2. Move TOR-DC1 to the Toronto site
3. Monitor AD DS site replication

▶ Task 1: Configure site links between AD DS sites

1. If necessary, on LON-DC1, open Active Directory Sites and Services.
2. Create a new IP-based site link with the following configuration:
   - Name: **TOR-TEST**
   - Sites: **Toronto, TestSite**
   - Modify the schedule to only allow replication from **Monday 9 AM to Friday 3 PM**
3. Rename DEFAULTIPSITELINK, and configure it with the following settings:
   - Name: **LON-TOR**
   - Sites: **LondonHQ, Toronto**
   - Replication: Every **60** minutes

▶ Task 2: Move TOR-DC1 to the Toronto site

1. If necessary, on LON-DC1, open Active Directory Sites and Services.
2. Move **TOR-DC1** from the **LondonHQ** site to the **Toronto** site.
3. Verify that TOR-DC1 is located under the Servers node in the Toronto site.

▶ Task 3: Monitor AD DS site replication

1. On LON-DC1, on the taskbar, click the Windows PowerShell icon.
2. Use the following commands to monitor site replication:

```
Repadmin /kcc
```

This command recalculates the inbound replication topology for the server.

```
Repadmin /showrepl
```

Verify that the last replication with TOR-DC1 was successful.

```
Repadmin /bridgeheads
```

This command displays the bridgehead servers for the site topology.

```
Repadmin /replsummary
```

This command displays a summary of replication tasks. Verify that no errors appear.

```
DCDiag /test:replications
```

Verify that all connectivity and replication tests pass successfully.

3. Switch to TOR-DC1, and then repeat the commands to view information from the TOR-DC1 perspective.

Results: After this exercise, you will have configured site links and monitored replication.

Exercise 4: Monitoring and Troubleshooting AD DS Replication

Scenario

After AD DS sites and replication are established, A. Datum experiences replication issues. You have to use monitoring and troubleshooting tools to diagnose the issue and resolve it.

The main tasks for this exercise are as follows:

1. Produce an error
2. Monitor AD DS site replication
3. Troubleshoot AD DS replication
4. To prepare for the next module

▶ Task 1: Produce an error

1. On LON-DC1, in Active Directory Sites and Services, replicate TOR-DC1 with LON-DC1 from the LondonHQ site.
2. In Windows PowerShell, run:

```
Get-ADReplicationUpToDatenessVectorTable –Target “adatum.com”
```

3. Observe the results and note the date/time of the most recent replication event.
4. Go to TOR-DC1 and execute the following Windows PowerShell script:

```
\\LON-DC1\E$\Mod05\Mod05Ex4.ps1
```

▶ Task 2: Monitor AD DS site replication

1. On TOR-DC1, in Active Directory Sites and Services, replicate **LON-DC1** with **TOR-DC1** from the **Toronto** site. Acknowledge the error.
2. In Windows PowerShell, run the following cmdlets, and observe the results:

```
Get-ADReplicationUpToDatenessVectorTable –Target “adatum.com”
Get-AdReplicationSubnet –filter *
Get-AdReplicationSiteLink–filter *
```

▶ Task 3: Troubleshoot AD DS replication

1. On TOR-DC1, in Windows PowerShell, determine the IP address settings for the computer, then run the following cmdlet:

```
Get-DnsClient | Set-DnsClientServerAddress -ServerAddresses
("172.16.0.10","172.16.0.25")
```

Ensure that the IP address settings are correct.

2. Go to **Active Directory Sites and Services** and replicate **LON-DC1** with **TOR-DC1** from the **Toronto** site. Acknowledge the error.

3. Attempt to go to the **DNS Console**. Close any error windows.
4. In Windows PowerShell, investigate the DNS Server service with the Get-Service cmdlet. If it is not running, start the service with the Start-Service cmdlet.
5. Go to **Active Directory Sites and Services** and replicate **LON-DC1** with **TOR-DC1** from the **Toronto** site. You should not get an error. Review the objects to determine if any are missing.
6. From **TOR-DC1**, open the following script in **Windows PowerShell ISI**:

```
\\LON-DC1\E$\Mod05\Mod05Ex4Fix.ps1
```

7. Run the **recreate Site Links** and **recreate subnets** sections of the script.
8. Return to **Active Directory Sites and Services** and determine if anything is still missing.
9. Close all open windows, and sign off **LON-DC1** and **TOR-DC1**.

▶ Task 4: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. On the **Virtual Machines** list, right-click **20412C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps two and three for 20412C-TOR-DC1.

Question: You decide to add a new domain controller to the LondonHQ site named LON-DC2. How can you ensure that LON-DC2 is used to pass all replication traffic to the Toronto site?

Question: You have added the new domain controller named LON-DC2 to the LondonHQ site. Which AD DS partitions will be modified as a result?

Question: In the lab, you created a separate site link for the Toronto and TestSite sites. What might you also have to do to ensure that LondonHQ does not automatically create a connection object directly with the TestSite site?

Module Review and Takeaways

Review Questions

Question: Why is it important that all subnets are identified and associated with a site in a multisite enterprise?

Question: What are the advantages and disadvantages of reducing the intersite replication interval?

Question: What is the purpose of a bridgehead server?

Best Practice:

- Implement the following best practices when you manage Active Directory sites and replication in your environment: Always provide at least one or more global catalog servers per site.
- Ensure that all sites have appropriate subnets associated.
- Do not set up long intervals without replication when you configure replication schedules for intersite replication.
- Avoid using SMTP as a protocol for replication.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Client cannot locate domain controller in its site.	• Verify whether all SRV records for the domain controller are present in DNS. • Verify whether the domain controller has an IP address from the subnet that is associated with that site. • Verify that the client is a domain member and has the correct time.
Replication between sites does not work.	• Verify whether site links are configured correctly. • Verify the replication schedule. • Verify whether the firewall between the sites permits traffic for Active Directory replication. Use repadmin /bind.
Replication between two domain controllers in the same site does not work.	• Verify whether both domain controllers appear in same site. • Verify whether AD DS is operating correctly on the domain controllers. • Verify network communication, and that the time on each server is valid.

Module 6

Implementing AD CS

Contents:

Module Overview

Public key infrastructure (PKI) consists of several components that help you secure corporate communications and transactions. One component is the certification authority (CA). You can use CAs to manage, distribute, and validate digital certificates that are used to secure information. You can install Active Directory® Certificate Services (AD CS) as a root CA or a subordinate CA in your organization. In this module, you will learn about implementing AD CS server role and certificates.

Objectives

After completing this module, you will be able to:

- Describe PKI.
- Deploy CAs.
- Explain how to deploy and configure a CA hierarchy.
- Explain how to deploy and manage certificate templates.
- Explain how to implement certificate distribution and revocation.
- Describe how to manage certificate recovery.

Lesson 1
Using Certificates in a Business Environment

Certificates are often used in today's electronic communications. Each time you open an HTTPS URL, a certificate is used to enable encryption. Also, certificates are used to sign content digitally, encrypt it, or authenticate a user or a device. In this lesson, you will learn some of the most common ways to use certificates in business environments.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe how to use certificates for Secure Sockets Layer (SSL).
- Describe how to use certificates for digital signatures.
- Describe how to sign a document digitally.
- Describe how to use certificates for content encryption.
- Describe how to use certificates for authentication.

Using Certificates for SSL

Most websites that deal with critical security data are protected with SSL security technology. SSL establishes a secure, encrypted link between a server and a client. Most commonly, the connection is between a web server and a browser or email client on a client computer. SSL is commonly referred to as a security protocol, because it specifies algorithms for encryption and necessary variables for connection encryption. The purpose of securing a connection with SSL is to protect data, such as credit card numbers, logon credentials, and other critical data, while the data transfers between a client and a server.

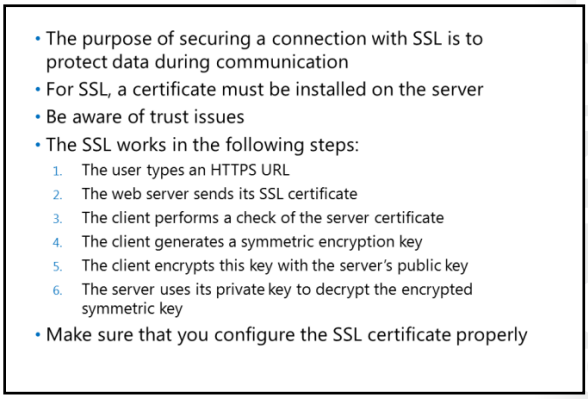

To establish a connection protected by SSL, the certificate must be installed on the server. Your internal CA or a public CA can issue a certificate for SSL. For websites available on the Internet, it is common to have a certificate issued by a public CA because of the trust issues. However, you can also use a certificate issued by your local CA. A connection can be secured with both types of certificates, but an internally issued certificate cannot be trusted by most clients that connect to the website where the certificate is installed. The fact that the certificate is untrusted will not prevent it from securing a connection, but it will present clients with a warning message when they connect to your website. Most companies want to avoid this, so public certificates are used for the most secure websites on the Internet. Internet browsers come with a preinstalled list of trusted CAs, and they store this list in the trusted root CA store.

Note: Buying a public SSL certificate does not automatically guarantee that the certificate will be trusted by all clients. Make sure that you choose a certificate vendor that is trusted globally and has its CA certificates present in the client's preinstalled trusted root CA store.

Securing a Connection with an SSL Certificate

Each certificate has a key pair associated with it after it is issued. The key pair consists of a public key and a private key. These keys work together in an encryption process. Data that is encrypted with a public key can be decrypted only with a corresponding private key, and vice versa. Each key pair is unique.

Besides having a key pair, each certificate also has its subject name that specifies the identity of the server or website where the certificate is installed.

When a web browser connects to a secure website, the client and server establish an SSL connection. The SSL connection establishes during the SSL *handshake*. This handshake process occurs through the following steps:

1. The user types or clicks an HTTPS URL in a web browser.
2. The web browser software connects to a website and requests for the server to identify itself.
3. The web server sends its SSL certificate. With the certificate, the server also distributes its public key to the client.
4. The client performs a check of the server certificate. It checks the subject name and compares it with the URL that it used to access the server. It also checks if the certificate is issued by any of the CAs in the trusted root CA store, and it checks the certificate revocation list distribution point (CDP) locations to verify if the certificate is revoked.
5. If all checks pass, the client generates a symmetric encryption key. The client and server use a symmetric key for decrypting data because the public and private key pairs are not very efficient in encrypting and decrypting large amounts of data. The client generates a symmetric key and then encrypts this key with the server's public key. After that, the client sends the encrypted symmetric key to the server.
6. The server uses its private key to decrypt the encrypted symmetric key. Now both the server and the client have a symmetric key, and the secure data transfer can begin.

During this process, several very important checks are performed. First, the server proves its identity to the client by presenting its SSL certificate. If the server name in the certificate matches the URL that the client requested, and if the certificate is issued by a trusted CA, the client trusts that the server has valid identity. Also, the client has checked the validity of the certificate by checking its lifetime and CDP location for the certificate revocation lists (CRLs). This means that establishing an SSL session does more than manage encryption; it also provides authentication from a server to a client.

 **Note:** Client authentication is not part of the classic SSL handshake. This means that a client does not have to provide its identity to the server. However, you also can configure your website to require client authentication. The client also can use a certificate to authenticate itself.

Configuring an SSL Certificate on a Server

To use SSL to protect communication between a server and a client, you must install the certificate on the server. You can install it in several ways. However, before you install the certificate on the server, you must define the name or names that the certificate supports. For example, if you want to protect your website on the URL www.adatum.com, you need to issue the certificate with the common name www.adatum.com.

 **Note:** A certificate can be issued only for a server name or alias, not for a full URL. For example, a certificate with the common name www.adatum.com will also protect URL www.adatum.com/sales or similar.

In some scenarios, you need to have more than one server name on the same server. A typical example for this is Microsoft Exchange Server® Client Access server. A certificate installed on the Client Access server must support its public name—for example, mail.adatum.com and autodiscover.adatum.com. Since both names are associated to the same website, and you cannot assign more than one certificate to a single website, you must use a certificate that supports multiple names, also known as *subject alternative names*. This means that you have one certificate with more than one name. These certificates can be issued from both and internal CA on Windows Server® 2012 and from public CAs.

Note: Instead of having one certificate with multiple names on the same domain, you also can issue a wild card certificate with a common name—for example, *.adatum.com. This certificate will be valid for all names with domain suffix adatum.com. However, we do not recommend using these certificates for security reasons.

To issue an SSL certificate from an internal CA, you can use following approaches:

- Use the CA console on the server to make the certificate request to the CA. By using this approach, you can specify any additional attributes for the certificate, such as the certificate template or the subject alternative name. However, after the certificate is installed, you must assign it to the appropriate website manually.
- Use the Internet Information Services (IIS) console. In the IIS console, you make a certificate request directly to the CA. However, when you use this approach, you are not able to choose a certificate template—it looks for a web server template by default—and you cannot specify a subject alternative name. This is, however, the simplest way to install a certificate on a website.
- Use CA Web enrollment. This approach is appropriate if you want to issue a certificate to a server that is not a member of your domain. For this type of enrollment, you must first make a certificate request (.req) file and then submit that request on the CA Web enrollment page. There, you also can specify the certificate template and add subject alternative names, if needed.

If you buy a publicly trusted SSL certificate, the procedure is somewhat different. After you choose a certificate vendor, you will first have to go through an administrative procedure to prove the identity of your company and domain name ownership. After you have completed that, you have to create a Certificate Signing Request (CSR) on your server. This CSR creates the private key and a CSR data file, which is basically a certificate request. You then send the CSR to the certificate issuer. The CA uses the CSR data file to create a public key to match your private key without compromising the key itself. The CA never sees the private key in this or any previous scenario for certificate issuing, except when key archival is configured—but even then, the key is encrypted.

Using Certificates for Digital Signatures

Besides protecting communications, certificates also can be used for protecting content and verifying the identity of the content author. When you receive a message with confidential content, it is important that you are certain about two things: first, that the message was not modified in transit, and second, that the identity of the author is verifiable.

You can use certificates to protect and verify content and verify the identity of an author. It is a common scenario for a user to sign a document digitally.

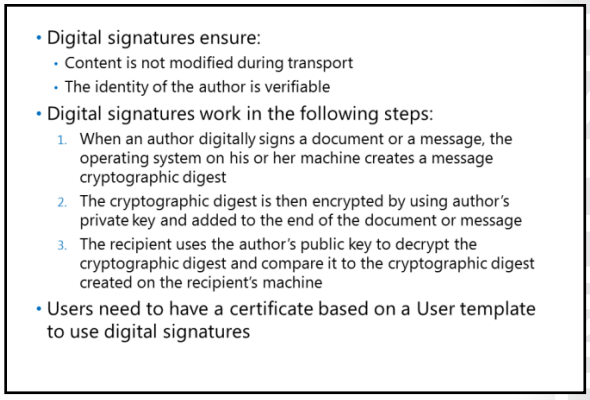

Digital Signatures

When a person digitally signs a document in an application—such as in email, a Microsoft Word document, or similar—he or she confirms that the document is *authentic*. In this context, authentic means that it is known who created a document, and that the document has not been altered in any way since the person created and signed it.

PKI can achieve this level of security. Compared to the web server from the previous topic, a user also can have certificate with a public and private key pair. This certificate is used in the process of digital signing.

When an author digitally signs a document or a message, the operating system on his or her machine creates a message cryptographic digest, which ranges from a 128-bit to a 256-bit number. It is generated by running the entire message through a hash algorithm. This number is then encrypted by using the author's private key, and it is added to the end of the document or message.

When the document or message reaches the recipient, it will go through the same hash algorithm as when it was digitally signed. Also, the recipient uses the author's public key to decrypt the digest that is added to the message. After it is decrypted, it is compared to the digest that the recipient has generated. If they are the same, the document or the message was not altered during transport. Also, if the recipient is able to decrypt the digest by using the author's public key, this means that the digest was encrypted by using the author's private key, and that confirms the author's identity. At the end, the recipient also verifies the certificate that was used to prove author's identity. During this check, the validity period, CRL, subject name, and certificate chain trust also are verified.

Implementing Digital Signatures

To implement digital signatures in internal communications, you need to issue certificates based on the user template. You must issue certificates to all users who use digital signatures. You can issue the certificate without any user intervention if you use autoenrollment. Also, users must use an application that supports content signing. For example, you can use digital signatures by default in Microsoft Word or Microsoft Outlook.

Digital signatures are ready to use after the certificate has been issued and configured in the application.

However, if you want to send digitally signed content outside of your organization, you might experience CA trust issues. In this scenario, a recipient is not in the same Active Directory® Domain Services (AD DS) domain as the author, so it does not trust the CA that issued the certificate for the digital signature. Although this kind of digital signature will still be valid from the content protection perspective, an application being used will probably generate a warning on the recipient's side.

If you need to send digitally signed content to recipients outside of your organization, we recommend that you buy certificates from a public, globally trusted CA.

Demonstration: Signing a Document Digitally

In this demonstration, your instructor will show you how to digitally sign a document in Microsoft Word.

Demonstration Steps

1. On LON-CL1, open the Windows® PowerShell command-line interface, and then run **mmc.exe**.
2. Add the Certificates snap-in, and then choose **My user account**.
3. Start the Request New Certificate Wizard, and then enroll for a User certificate.
4. Open Microsoft® Word 2013, type some text in the blank document, and then save the document.
5. Click **Insert** on the ribbon, and then insert the signature line.
6. Fill the signature fields with your data.
7. Right-click the signature line, and then select to sign the document.
8. Select the certificate.
9. Sign the document.
10. Make sure that the document cannot be edited anymore.

Using Certificates for Content Encryption

While digital signatures can verify an author's identity and ensure content consistency, they cannot protect the content itself. For example, if someone intercepts a digitally signed message, the person can still read its content; however, the attempt to alter the content is detected because the digital signature check will fail.

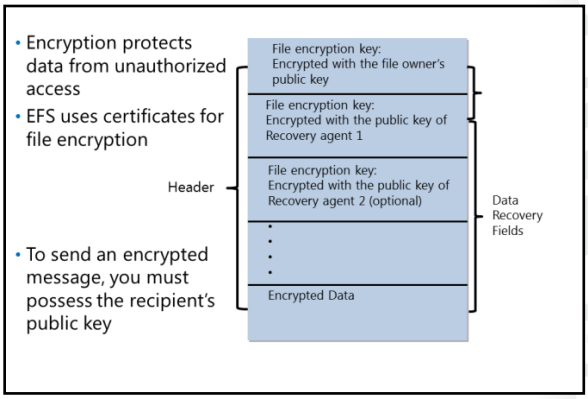

If you want to protect the content of the document so that it cannot be read, you must use encryption.

The Windows operating system supports file-based encryption called Encrypting File System (EFS). Outlook also supports the encryption of email messages.

EFS

To encrypt a file by using EFS, you must have an EFS certificate issued. Like other certificates, this certificate also provides a private and public key pair. However, these keys are not used directly to encrypt or decrypt content. The reason for this is because, the algorithms that use *asymmetric encryption*, where one key is used for encryption and another for decryption, are inefficient. These algorithms are 100 to 1,000 times slower than algorithms that use the same key for both encryption and decryption, which is called *symmetric encryption*. To overcome this problem, EFS uses a somewhat hybrid approach.

When a user selects the option to encrypt a file, the local computer generates a symmetric key, which is also known as a file encryption key, and uses that key to encrypt the file. After the file is encrypted, the system uses the user's public key to encrypt the symmetric key and then store it in the file header.

When the user who originally encrypted the file wants to decrypt the file and access its content, the local computer accesses the user's private key and first decrypts the symmetric key from file header, which also is called the Data Decryption Field. After that, the symmetric key is used to decrypt the content.

This works well if the file's owner is the only person who accesses the encrypted file. However, there are scenarios where you want to share encrypted files with other users, and it might be inconvenient or unacceptable to decrypt the file before sharing it with other people. Also, if the user who originally encrypted the file loses their private key, the file might be inaccessible to anyone.

To resolve this, Data Recovery Field is defined for each file encrypted with EFS. When you configure EFS to use locally or in an AD DS domain, the Data Recovery Agent role is defined by default and assigned to local or Domain Admin. The Data Recovery Agent is actually a certificate with a key pair that can be used to decrypt files in case the private key of the originating user is not accessible for any reason.

When a user encrypts the file with EFS, his or her public key is used to encrypt the symmetric key, and that encrypted key then is stored to the Data Decryption Field in the file header. At the same time, the public key of the Data Recovery Agent is used to encrypt the symmetric key once more. The symmetric key is encrypted with a public key of the Data Recovery Agent and then is stored to the Data Recovery Field in the file header. If more than one Data Recovery Agent is defined, the symmetric key is encrypted with each Data Recovery Agent's public key. Then, if the user who originally encrypted the file does not have a private key available for any reason, the Data Recovery Agent can use its private key to decrypt the symmetric key from Data Recovery Field, and then decrypt the file.

 **Note:** As an alternative to the Data Recovery Agent, you also can use the Key Recovery Agent (KRA) to retrieve a user's private key from a CA database, if key archival is enabled for the EFS certificate template and on the CA.

When a user wants to share an encrypted file with other users, an approach is used that is similar to the one taken when Data Recovery Agent is used. When EFS sharing is selected, the file's owner must select a certificate from each user who shares the file. These certificates can be published to AD DS and taken from there. When the certificate is selected, the public key of destination user is taken, and the symmetric key is encrypted and added to the file header. Then the other user also can access the EFS encrypted content, as he or she can use their private keys to decrypt the symmetric key.

 **Note:** The Data Recovery Agent also can be defined for BitLocker® Drive Encryption. Because the BitLocker Data Recovery Agent certificate template is not predefined, you can copy the KRA template and then add the BitLocker encryption and the BitLocker Drive Recovery Agent from the application policies. After you enroll a user for this certificate, you can add it to be the BitLocker Data Recovery Agent at the domain level if you use Group Policy settings, in following path: Computer Configuration\Windows Settings\Security\Public Key Policies\BitLocker Drive Encryption.

Email Encryption

Besides the use of EFS to encrypt files and BitLocker® to encrypt drives, you also can use certificates to encrypt emails. However, email encryption is more complicated than a digital signature. While you can send a digitally signed email to anyone, you cannot do the same with an encrypted email. To send an encrypted email to someone with a PKI, you must possess the recipient's public key from his or her key pair. In the AD DS environment, where Exchange Server is used as an email system, you can publish the public keys of all mailbox users to a global address list (GAL). When you do that, an application such as Outlook can extract a recipient's public key easily from the GAL, if you are sending encrypted email. When you send an encrypted email to an internal user, your email application takes the recipient public key from the GAL, encrypts the email with it, and then sends the email. After the email is received, the recipient uses his or her private key from the certificate to decrypt the content of the email.

Sending an encrypted email to external users is more complicated. While public keys of internal users can be published to the AD DS or the GAL, you cannot do the same with external users. To send an encrypted

email to an external user, you first must get his or her public key. You can get the key if the external user sends it to you in a .cer file, which you can import to your local address book. Also, if an external user sends you one digitally signed email, you will get his or her public key, which you also can import to your local address book. After the public key is imported into your address book, you can use it to send encrypted emails to external users.

Note: If you want to provide authenticity, content consistency, and protection, you can send a message that is both digitally signed and encrypted.

Using Certificates for Authentication

Besides the use of certificates for digital signing and encryption, they often are used for user and device authentication. Also, certificates commonly are used for network access authentication because they provide strong security for authenticating users and computers, and they eliminate the need for less secure password-based authentication methods.

You can use certificates for user and device authentication, and in network and application access scenarios such as:

- L2TP/IPsec VPN
- EAP-TLS
- PEAP
- NAP with IPsec
- Outlook Web App
- Mobile device authentication

For example, you can use certificates on computers that are allowed to access your network by using virtual private network (VPN) connections. This enables you to authenticate devices and users. A user can authenticate with a user name and password, while a device authenticates with a certificate. Devices that do not have your company's certificate will not be allowed to connect, even if a user is authorized. This approach improves security.

Two authentication methods for network access use of certificates are: Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) and Protected Extensible Authentication Protocol (PEAP). Both methods always use certificates for server authentication. Depending on the authentication type configured with the authentication method, certificates might be used for user and client device authentication.

You must use certificate-based authentication for VPN connections based on Layer Two Tunneling Protocol over Internet Protocol security (IPsec).

Certificates also are used to authenticate clients when Network Access Protection (NAP) is implemented with IPsec. In this scenario, the Health Registration Authority issues a certificate to a computer that satisfies the health policy for establishing an IPsec connection.

IIS in Windows Server 2012 also supports certificate authentication for users. For example, you can configure Exchange Outlook® Web App to use certificate-based authentication.

Finally, you also can use certificates for mobile device authentication. Some types of mobile devices can install certificates and use them to authenticate a user or device to the network resource.

Lesson 2
PKI Overview

PKI helps you verify and authenticate the identity of each party involved in an electronic transaction. It also helps you establish trust between computers and the corresponding applications that are hosted on application servers. A common example includes the use of PKI technology to secure websites. Digital certificates are key PKI components that contain electronic credentials, which are used to authenticate users or computers. Moreover, certificates can be validated using certificate discovery, path validation, and revocation checking processes. Windows Server 2012 supports building a certificate services infrastructure in your organization by using AD CS components.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe PKI.
- Describe components of a PKI solution.
- Describe CAs.
- Describe the AD CS server role in Windows Server 2012.
- Describe new features in AD CS in Windows Server 2012.
- Explain the difference between public and private CAs.
- Describe cross-certification hierarchy.

What Is PKI?

PKI is a combination of software, encryption technologies, processes, and services that assist an organization with securing its communications and business transactions. It is a system of digital certificates, CAs, and other registration authorities. When an electronic transaction takes place, PKI verifies and authenticates the validity of each party involved. PKI standards are still evolving, but they are widely implemented as an essential component of electronic commerce.

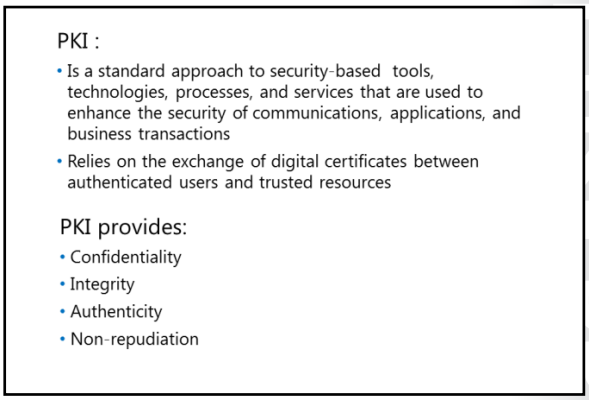

General Concepts of PKI

In general, a PKI solution relies on several technologies and components. When you plan to implement PKI, you should consider and understand the following:

- Infrastructure. The meaning in this context is the same as in any other context, such as electricity, transportation, or water supply. Each of these elements has a specific job, and requirements that must be met for it to function efficiently. The sum of these elements allows for the efficient and safe use of PKI. The elements that make up a PKI include the following:
    - A CA
    - A certificate repository
    - A registration authority

    - An ability to revoke certificates
    - An ability to back up, recover, and update keys
    - An ability to regulate and track time
    - Client-side processing

Most of these components will be discussed in later topics and lessons throughout this module.

- Public/Private Keys. In general, there are two methods for encrypting and decrypting data:
    - Symmetric encryption: The methods to encrypt and decrypt data are identical, or mirrors of each other. Data is encrypted by using a particular method or key. To decrypt the data, you must have the same identical method or key. Therefore, anyone who has the key can decrypt the data. The key must remain private to maintain the integrity of the encryption.
    - Asymmetric encryption: In this case, the methods to encrypt and decrypt data are neither identical nor mirrors of each other. Data is encrypted by using a particular method or key. However, a different key is used to decrypt data. This is achieved by using a pair of keys. Each person gets a key pair, which consists of a public key and a private key. These keys are unique, and data that the public key encrypts can be decrypted by using the private key, and vice versa. In this situation, the keys are sufficiently different, and knowing or possessing one does not allow you to determine the other. Therefore, one of the keys (public) can be made publicly available without reducing the security of the data, as long as the other key (private) remains private—hence the name Public Key Infrastructure.

Algorithms that use symmetric encryption are fast and efficient for large amounts of data. However, because they use a symmetric key, they are not considered secure enough, because you always must transport the key to the other party. Alternatively, algorithms that use asymmetric encryption are secure, but are very slow. Because of this, it is common to use a hybrid approach, which means that data is encrypted by using symmetric encryption, while the symmetric encryption key is protected with asymmetric encryption.

When you implement a PKI solution, your entire system, and especially the security aspect, can benefit. The benefits of using PKI include:

- Confidentiality. A PKI solution enables you to encrypt both stored and transmitted data.
- Integrity. You can use PKI to sign data digitally. A digital signature identifies whether any data was modified while information was in transit.
- Authenticity and non-repudiation. Authentication data passes through hash algorithms such as Secure Hash Algorithm 1 to produce a message digest. The message digest is then digitally signed using the sender's private key to prove that the message digest was produced by the sender. Non-repudiation is digitally signed data in which the digital signature provides both proof of the integrity of signed data, and proof of the origin of data.
- Standards-based approach. PKI is standards-based, which means that multiple technology vendors are compelled to support PKI-based security infrastructures. It is based on industry standards defined in RFC 2527, "Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework."

Components of a PKI Solution

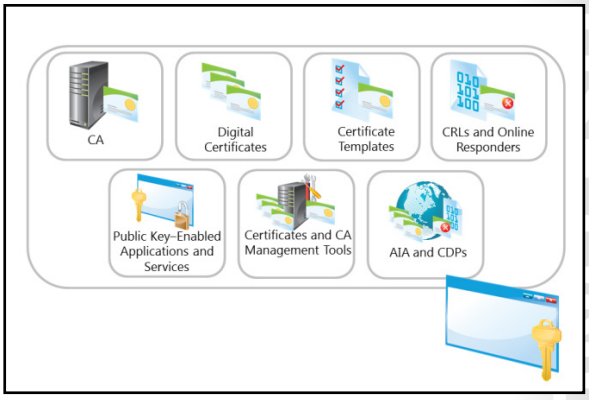

Many components must work together to provide a complete PKI solution. The PKI components in Windows Server 2012 are as follows:

- CA. CA issues and manages digital certificates for users, services, and computers. By deploying CA, you establish the PKI in your organization.
- Digital certificates. Digital certificates are similar in function to an electronic passport. A digital certificate is used to prove the identity of the user (or other entity). Digital certificates contain the electronic credentials that are associated with a public key and a private key, which are used to authenticate users and other devices such as web servers and mail servers. Digital certificates also ensure that software or code is run from a trusted source. Digital certificates contain various fields, such as Subject**,** Issuer**,** and Common Name. These fields are used to determine the specific use of the certificate. For example, a web server certificate might contain the Common Name field of web01.contoso.com, which would make that certificate valid only for that web server. If an attempt were made to use that certificate on a web server named web02.contoso.com, the user of that server would receive a warning.
- Certificate templates. This component describes the content and purpose of a digital certificate. When requesting a certificate from an AD CS enterprise CA, the certificate requestor will be able, depending on his or her access rights, to select from a variety of certificate types based on certificate templates, such as User and Code Signing. The certificate template saves users from low-level, technical decisions about the type of certificate they need. In addition, they allow administrators to distinguish who might request which certificates.
- CRLs and Online Responders.
  - CRLs are complete, digitally signed lists of certificates that have been revoked. These lists are published periodically and can be retrieved and cached by clients (based on the configured lifetime of the CRL). The lists are used to verify a certificate's revocation status.
  - Online Responders are part of the Online Certificate Status Protocol (OCSP) role service in Windows Server 2008 and Windows Server 2012. An Online Responder can receive a request to check for revocation of a certificate without requiring the client to download the entire CRL. This speeds up certificate revocation checking, and reduces the network bandwidth. It also increases scalability and fault tolerance, by allowing for array configuration of Online Responders.
- Public key–based applications and services. This relates to applications or services that support public key encryption. In other words, the application or services must be able to support public key implementations to gain the benefits of it.
- Certificate and CA management tools. Management tools provide command-line and graphical user interface (GUI)-based tools to:
  - Configure CAs
  - Recover archived private keys
  - Import and export keys and certificates
  - Publish CA certificates and CRLs
  - Manage issued certificates

- Authority information access (AIA) and CDPs. AIA points determine the location where CA certificates can be found and validated, and CDP locations determine the points where CRLs can be found during certificate validation process. Because CRLs can become large, (depending on the number of certificates issued and revoked by a CA), you can also publish smaller, interim CRLs called *delta CRLs*. Delta CRLs contain only the certificates revoked since the last regular CRL was published. This allows clients to retrieve the smaller delta CRLs and quickly build a complete list of revoked certificates. The use of delta CRLs also allows revocation data to be published more frequently, because the size of a delta CRL means that it usually does not require as much time to transfer as a full CRL.
- Hardware security module (HSM). A HSM is an optional secure cryptographic hardware device that accelerates cryptographic processing for managing digital keys. It is a high-security, specialized storage device that is connected to the CA for managing the certificates. An HSM is typically physically attached to a computer. This is an optional add-on in your PKI, and is the most widely used in high security environments where there would be a significant impact if a key were compromised.

Note: The most important component of any security infrastructure is physical security. A security infrastructure is not just the PKI implementation. Other elements—such as physical security and adequate security policies—are also important parts of a holistic security infrastructure.

What Are CAs?

A CA is a well-designed and highly trusted service in an enterprise, that provides users and computers with certificates, manages and publishes the CRLs, and optionally responds to OCSP requests. You can install a CA in your environment by deploying the AD CS role on Windows Server 2012. When you install the first CA, it establishes the PKI in the network, and it provides the highest point in the entire structure. You can have one or more CAs in a network, but only one CA can be at the highest point on the CA hierarchy (that CA is called the *root CA*, which will be discussed later in this module).

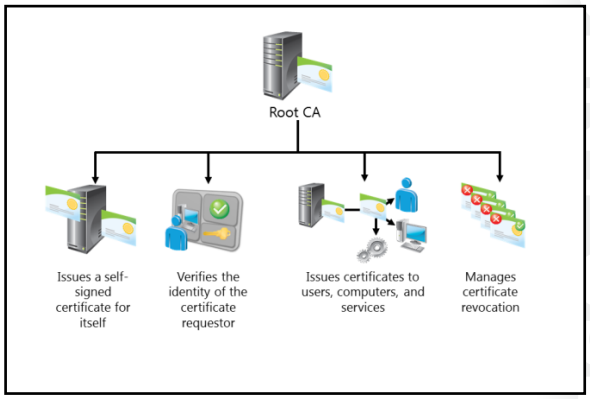

The main purposes of the CA are to issue certificates, revoke certificates, and publish AIA and CRL information. By doing this, the CA ensures that users, services, and computers are issued certificates that can be validated.

A CA performs multiple functions or roles in a PKI. In a large PKI, separation of CA roles among multiple servers is common. A CA provides several management tasks, including:

- Verifying the identity of the certificate requestor.
- Issuing certificates to requesting users, computers, and services.
- Managing certificate revocation.

When you deploy a first CA (root CA) in your network, it issues a certificate for itself. After that, other CAs receive certificates from the root CA. You can also choose to issue a certificate for your CA by using one of the public CAs.

Overview of the AD CS Server Role in Windows Server 2012

All PKI-related components are deployed as role services of the AD CS server role. The AC CS server role is made up of several components that are known as *role services*. While each role service is responsible for a specific portion of the certificate infrastructure, all services work together to form a complete solution.

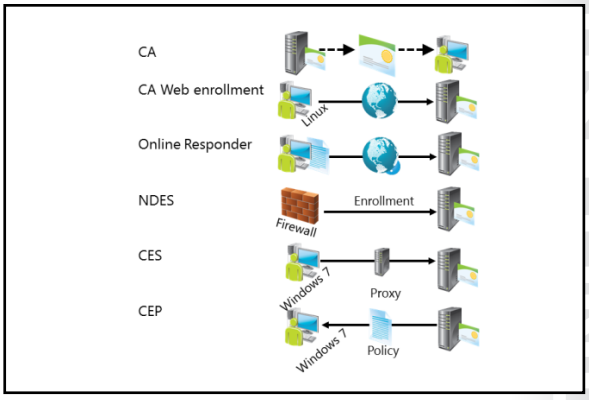

Role services of the AD CS role are:

- CA. This component issues certificates to users, computers, and services. It also manages certificate validity. Multiple CAs can be chained to form a PKI hierarchy.
- CA Web enrollment. This component provides a method to issue and renew certificates for users, computers, and devices that are not joined to the domain, are not connected directly to the network, or are for users of operating systems other than Windows.
- Online Responder. You can use this component to configure and manage OCSP validation and revocation checking. Online Responder decodes revocation status requests for specific certificates, evaluates the status of those certificates, and returns a signed response containing the requested certificate status information. Unlike in Windows Server 2008 R2, you can install Online Responder on any version of Windows Server 2012. When using Online Responder, the certificate revocation data can come from a CA on a computer that is running Windows Server 2003, Windows Server 2008, or a CA other than Microsoft.
- Network Device Enrollment Service (NDES). With this component, routers, switches, and other network devices can obtain certificates from AD CS. On Windows Server 2008 R2, this component is only available on the Enterprise and Datacenter editions, but with Windows Server 2012, you can install this role service on any version.
- Certificate Enrollment Web Service (CES). This component works as a proxy between Windows 7 and Windows 8 client computers and the CA. This component is new to Windows Server 2008 R2 and Windows Server 2012, and requires that the Active Directory forest be at least at the Windows Server 2008 R2 level. It enables users to connect to a CA by means of a web browser to perform the following:
    - Request, renew, and install issued certificates.
    - Retrieve CRLs.
    - Download a root certificate.
    - Enroll over the internet or across forests (new to Windows Server 2008 R2).
- Certificate Enrollment Policy Web Service (CEP). This component is new to Windows Server 2008 R2 and Windows Server 2012. It enables users to obtain certificate enrollment policy information. Combined with the CES, it enables policy-based certificate enrollment when the client computer is not a member of a domain, or when a domain member is not connected to the domain.

New Features of AD CS in Windows Server 2012

Like many other Windows Server roles, AD CS is improved and enhanced in Windows Server 2012. The AD CS role in Windows Server 2012 still has the same six role services as described in the previous topic. In addition, it now provides multiple new features and capabilities when compared with previous versions.

- All AD CS role services run on all versions of Windows Server
- Full integration with Server Manager
- Manageable through Windows PowerShell
- New certificate template version (v4)
- Support for automatic renewal of certificates for non-domain joined computers
- Enforcement of certificate renewal with the same key
- Additional security for certificate requests
- Support for Virtual Smart Cards

In Windows Server 2008 R2, some of the AD CS role services require a specific version of Windows Server. For example, NDES does not work on the Windows Server 2008 Standard edition, only on the Windows Server 2008 Enterprise edition. In Windows Server 2012, all role services are available on all versions of Windows Server.

The AD CS Server role, in addition to all related role services, can run on Windows Server 2012 with full GUI, Minimal Server Interface, or on a Server Core installation. You can deploy the AD CS role services in Windows Server 2012 using Server Manager, or using Windows® PowerShell cmdlets. Additionally, you can deploy the role services while working locally at the computer, or remotely over the network.

From a management perspective, AD CS and its events, and the Best Practices Analyzer tool are now fully integrated into the Server Manager console, which means that you can access all of its options directly from Server Manager. AD CS is also fully manageable by using the Windows PowerShell command-line interface.

The Windows Server 2012 version of AD CS also introduces a new certificate template version—version 4 (v4)—which provides some new capabilities. This will be discussed later in this module.

CES is also enhanced in Windows Server 2012. This feature, introduced in Windows 7 and Windows Server 2008 R2, allows online certificate requests to come from untrusted AD DS domains or even from computers or devices that are not joined to a domain. AD CS in Windows Server 2012 adds the ability to renew certificates automatically for computers that are part of untrusted AD DS domains, or are not joined to a domain.

From a security perspective, AD CS in Windows Server 2012 provides the ability to require the renewal of a certificate with the same key. Windows Server 2012 also supports generating trusted platform module (TPM)–protected keys using TPM-based key storage providers. The benefit of using a TPM-based key storage provider is true non-exportability of keys that are backed up by the TPM mechanism that blocks users if they enter a wrong PIN too many times. To enhance security even further, you can now force encryption of all certificate requests that come to AD CS in Windows Server 2012.

Virtual Smart Cards

Smart cards, as an option for multi-factor authentication, have been used since the Microsoft Windows 2000 Server operating system. Smart cards provide enhanced security over passwords, because it is much more difficult for an unauthorized user to gain and maintain access to a system. In addition, access to a smart card–protected system requires that a user both have a valid card and know the PIN that provides access to that card. By default, only one copy of the smart card exists, so only one individual can use their logon credentials at a time. In addition, a user will quickly notice if their card has been lost or stolen, especially when their card is combined with physical access to doors or other functions. This greatly reduces the risk window of credential theft in comparison to passwords.

However, implementation of the smart card infrastructure has proved historically too expensive in some situations. To implement smart cards, companies had to buy hardware, including smart card readers and smart cards. This cost, in some cases, prevented the deployment of Multi-factor Authentication.

To address these issues, Windows Server 2012 AD CS introduces a technology that provides the security of smart cards while reducing material and support costs. This is done by providing Virtual Smart Cards. Virtual Smart Cards emulate the functionality of traditional smart cards, but instead of requiring the purchase of additional hardware, they utilize technology that users already own and are more likely to have with them at all times.

Virtual Smart Cards in Windows Server 2012 leverage the capabilities of the TPM chip that is present on most of the computer motherboards produced in the past two years. Because the chip is already in the computer, there is no cost for buying smart cards and smart card readers. However, unlike traditional smart cards, which required that the user be in physical possession of the card, in the Virtual Smart Card scenario, a computer (or to be more specific, TPM chip on its motherboard) acts like a smart card. When using a TPM chip, you also achieve two-factor authentication, similar to when you use a smart card with a PIN. A user must have his or her computer (which has been set up with the Virtual Smart Card), and also know the PIN required to use his or her Virtual Smart Card.

It is important to understand how Virtual Smart Cards protect private keys. Traditional smart cards have their own storage and cryptographic mechanism for protecting the private keys. In the Virtual Smart Card scenario, private keys are protected not by isolation of physical memory, but rather by the cryptographic capabilities of the TPM: All sensitive information that is stored on a smart card is encrypted using the TPM and then stored on the hard drive in its encrypted form. Although private keys are stored on a hard drive (in encrypted form), all cryptographic operations occur in the secure, isolated environment of the TPM. Private keys never leave this environment in unencrypted form. If the hard drive of the machine is compromised in any way, private keys cannot be accessed, because they are protected and encrypted by TPM. To provide more security, you can also encrypt the drive with BitLocker Drive Encryption. To deploy Virtual Smart Cards, you need Windows Server 2012 AD CS and a Windows 8 client machine with a TPM chip on the motherboard.

Public vs. Private CAs

When you plan PKI implementation for your organization, one of the first choices you should make is whether to use private or public CAs. It is possible for you to establish PKI by using either of these approaches. If you decide to use a private CA, then you deploy the AD CS server role, and then establish an internal PKI. If you decide to use an external PKI, you do not have to deploy any service internally.

Both approaches have advantages and disadvantages, as specified in the following table:

Internal private CAs:
- Require greater administration than external public CAs
- Cost less than external public CAs, and provide greater control over certificate management
- Are not trusted by external clients by default
- Offer advantages such as customized templates and autoenrollment

External public CAs:
- Are trusted by many external clients
- Have slower certificate procurement

CA type	Advantages	Disadvantages
External public CA	• Trusted by many external clients (web browsers and operating systems) • Requires minimal administration	• Higher cost as compared to an internal CA • Cost is based per certificate • Certificate procurement is slower

CA type	Advantages	Disadvantages
Internal private CA	• Provides greater control over certificate management • Lower cost as compared to a public CA • Customized templates • Autoenrollment	• By default, not trusted by external clients (web browsers and operating systems) • Requires greater administration

Some organizations have started using a hybrid approach to their PKI architecture. A hybrid approach uses an external public CA for the root CA, and a hierarchy of internal CAs for distribution of certificates. This gives organizations the advantage of having their internally issued certificates trusted by external clients, while still providing the advantages of an internal CA. The only disadvantage to this method is cost. A hybrid approach is typically the most expensive approach, because public certificates for CAs are very expensive.

You can also choose to deploy an internal PKI for internal purposes such as EFS and digital signatures. For external purposes, such as protecting web or mail servers with SSL, you must buy a public certificate. This approach is not very expensive, and it is probably the most cost-effective solution.

Some organizations that require a higher security level might also choose to define their own list of trusted root CAs, both public and internal.

What Is a Cross-Certification Hierarchy?

As the term implies, a cross-certification hierarchy is one in which the root CA in each CA hierarchy provides a cross-certification certificate to the root CA in the other CA hierarchy. The other hierarchy root CA then installs the supplied certificate. By doing this, the trust flows down to all the subordinate CAs below the level where the cross-certification certificate was installed.

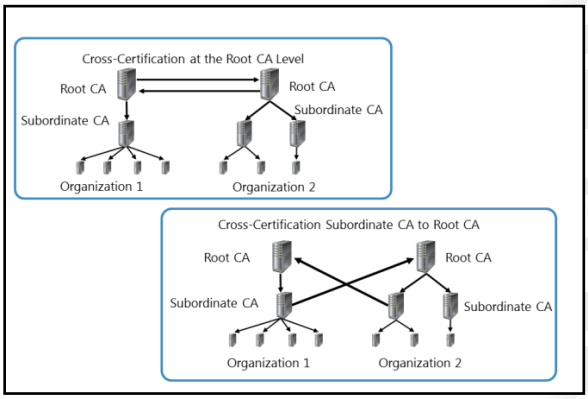

Cross-Certification Benefits

A cross-certification hierarchy provides the following benefits:

- Provides interoperability between businesses and between PKI products.
- Joins disparate PKIs.
- Assumes complete trust of a foreign CA hierarchy.

Companies usually deploy cross-certifications to establish a mutual trust on PKI level, and also to implement some other applications that rely on PKI, such as Active Directory® Rights Management Services (AD RMS).

Question: Your company is currently acquiring another company. Both companies run their own PKI. What could you do to minimize disruption and continue to provide PKI services seamlessly?

Lesson 3
Deploying CAs

The first CA that you install will be a root CA. After you install the root CA, you can optionally install a subordinate CA to apply policy restrictions and distribute certificates. You can also use a CAPolicy.inf file to automate additional CA installations and provide additional configuration settings that are not available with the standard GUI–based installation. Also, you can use Policy and Exit modules in the CA to integrate your CA with other services, such as Microsoft® Forefront Identity Manager (FIM).In this lesson, you will learn about deploying and managing CAs in the Windows Server 2012 environment.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe options for implementing CA hierarchies.
- Explain differences between stand-alone and enterprise CAs.
- Describe considerations for deploying a root CA.
- Deploy a root CA.
- Describe considerations for deploying a subordinate CA.
- Describe how to use CAPolicy.inf file for installing the CA.
- Explain how to configure CA administration and security.
- Explain how to configure CA Policy and Exit modules.
- Configure CA properties.
- Describe how to perform CA backup and recovery.

Options for Implementing CA Hierarchies

When you decide to implement a PKI in your organization, one of the first decisions you must make is how to design your CA hierarchy. CA hierarchy determines the core design of your internal PKI, and also it determines the purpose of each CA in the hierarchy. Each CA hierarchy usually includes two or more CAs. Usually, the second CA and all subsequent CAs are deployed with a specific purpose. Only the root CA is mandatory.

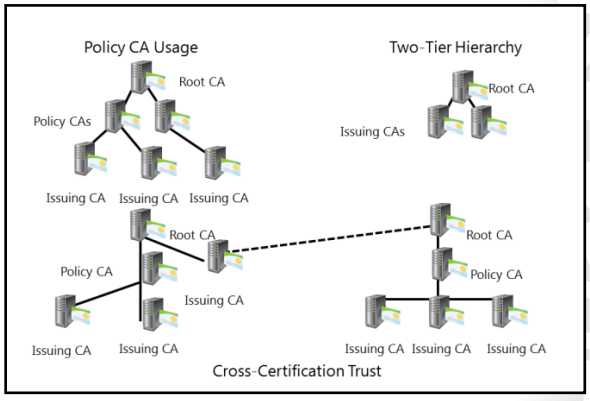

Note: It is not mandatory to have a CA hierarchy deployed to use PKI and certificates. For smaller and simpler environments, you can have a CA hierarchy with just one CA deployed. This CA usually is deployed as an enterprise root CA.

If you decide to implement a CA hierarchy, and you have already deployed a root CA, you must decide which roles to assign CAs on the second and third tiers. In general, we do not recommended that you build a CA hierarchy deeper than three levels, unless it is in a complex and distributed environment.

Most commonly, CA hierarchies have two levels, with the root CA at the top level and the subordinate issuing CA on the second level. The root CA usually is taken offline while the subordinate CA issues and manages certificates for all clients. However, in some more complex scenarios, you also can deploy other types of CA hierarchies.

In general, CA hierarchies fall into one of following categories:

- CA hierarchies with a policy CA. Policy CAs are types of subordinate CAs that are located directly below the root CA in a CA hierarchy. You use policy CAs to issue CA certificates to subordinate CAs that are located directly below the policy CA in the hierarchy. The role of a policy CA is to describe the policies and procedures that an organization implements to secure its PKI, the processes that validate the identity of certificate holders, and the processes that enforce the procedures that manage certificates. A policy CA issues certificates only to other CAs. The CAs that receive these certificates must uphold and enforce the policies that the policy CA defined. It is not mandatory to use policy CAs unless different divisions, sectors, or locations of your organization require different issuance policies and procedures. However, if your organization requires different issuance policies and procedures, you must add policy CAs to the hierarchy to define each unique policy. For example, an organization can implement one policy CA for all certificates that it issues internally to employees, and another policy CA for all certificates that it issues to users who are not employees.
- CA hierarchies with cross-certification trust. In this scenario, two independent CA hierarchies interoperate when a CA in one hierarchy issues a cross-certified CA certificate to a CA in another hierarchy. When you do this, you establish mutual trust between different CA hierarchies.
- CAs with a two-tier hierarchy. In a two-tier hierarchy, there is a root CA and at least one subordinate CA. In this scenario, the subordinate CA is responsible for policies and for issuing certificates to the requestors.

Stand-alone vs. Enterprise CAs

In Windows Server 2012, you can deploy two types of CAs: a stand-alone CA and an enterprise CA. These CA types are not structured around functionality and configuration storage rather than hierarchy. The most important difference between these two CA types is Active Directory integration and dependency. A stand-alone CA can work without AD DS, and does not depend on it in any way. An enterprise CA requires AD DS, but it also provides several benefits, such as autoenrollment.

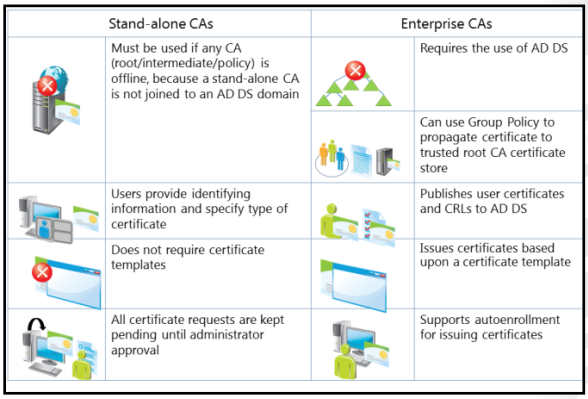

The following table details the most significant differences between stand-alone and enterprise CAs.

Characteristic	Stand-alone CA	Enterprise CA
Typical usage	A stand-alone CA is typically used for offline CAs, but it can be used for a CA that is consistently available on the network.	An enterprise CA is typically used to issue certificates to users, computers, and services, and is not typically used as an offline CA.

Characteristic	Stand-alone CA	Enterprise CA
Active Directory dependencies	A stand-alone CA does not depend on AD DS and can be deployed in non–Active Directory environments.	An enterprise CA requires AD DS, which can be used as a configuration and registration database. An enterprise CA also provides a publication point for certificates issued to users and computers.
Certificate request methods	Users can only request certificates from a stand-alone CA by using a manual procedure or CA Web enrollment.	• Users can request certificates from an enterprise CA using the following methods: • Manual enrollment • Web Enrollment • Autoenrollment • Enrollment agent
Certificate issuance methods	All requests must be manually approved by a certificate administrator.	Requests can be automatically issued or denied, based on the template's discretionary access control list (DACL).

Most commonly, the root CA is deployed as stand-alone CA, and it is taken offline after it issues a certificate for itself and for a subordinate CA. Alternatively, a subordinate CA is usually deployed as an enterprise CA, and is configured in one of the scenarios described in the previous topic.

Considerations for Deploying a Root CA

There are several decisions that you should make before you deploy a root CA. First, you should decide whether you need to deploy an offline root CA. Based on that decision, you also need to decide if you are going to deploy a stand-alone root CA or an enterprise root CA.

Usually, if you deploy a single-layer CA hierarchy—which means that you deploy only a single CA—it is most common to choose an enterprise root CA. However, if you deploy a two-layer hierarchy, the most common scenario is to deploy a stand-alone root CA and an enterprise subordinate CA.

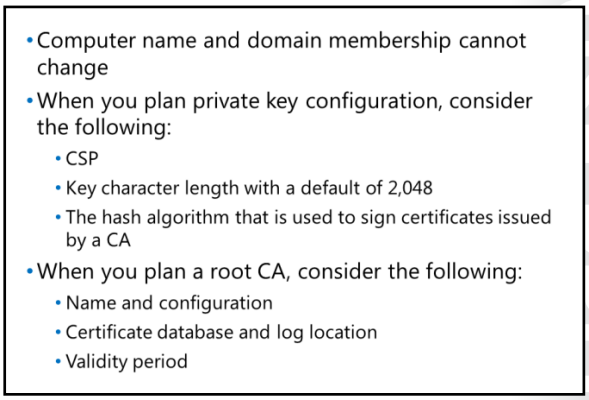

The next factor to consider is the operating system installation type. AD CS is supported in both the full installation and the Server Core installation scenarios. Server Core installation provides a smaller attack surface and less administrative overhead, and therefore should be considered for AD CS in an enterprise environment. You also can install AD CS in the Minimal Server Interface environment.

In Windows Server 2012, you also can use Windows PowerShell to deploy and manage the AD CS role.

You should be aware that you cannot change computer names or computer domain memberships after you deploy a CA of any type on that computer, nor can you change the domain name. Therefore, it is important to determine these attributes before installing a CA.

The following table details additional considerations:

Consideration	Description
A cryptographic service provider (CSP) that is used to generate a new key	• The default CSP is the Microsoft Strong Cryptographic Provider. • Any provider whose name contains a number sign (#) is a Cryptography Next Generation (CNG) provider.
The key character length	The default key length for the Microsoft Strong Cryptographic Provider is 2,048 characters. This is the minimum recommended value for a root CA, although it is a best practice to select a 4096-bit key.
The hash algorithm that is used to sign certificates issued by a CA	The default hash algorithm is Secure Hash Algorithm 1. If you do not have older clients such as Windows XP, you can choose newer hash algorithms such as SHA256.
The validity period for certificates issued by a CA	The default value for certificates is defined by templates. You can choose various validity periods on various certificate templates.
The status of the root server (online or offline)	The root server should be deployed as an offline CA, if possible. This enhances security and safeguards the root certificate because it is not available to attack over the network.

Specifically, if you decide to deploy an offline, stand-alone root CA, there are some specific considerations that you should keep in mind:

- Before you issue a subordinate certificate from the root CA, make sure that you provide at least one CDP and AIA location that will be available to all clients. This is because, by default, a stand-alone root CA has the CDP and AIA located on itself. Therefore, when you take the root CA off the network, the revocation check will fail because the CDP and AIA locations will be inaccessible. When you define these locations, you should copy the CRL and AIA information manually to that location.
- Set a validity period, for example one year, for CRLs that the root CA publishes to. This means that you will have to turn on root CA once per year to publish a new CRL, and then copy it to a location that is available to the clients. If you fail to do this, after the CRL on the root CA expires, the revocation check for all certificates also will fail.
- Use Group Policy to publish the root CA certificate to a trusted root CA store on all server and client machines. You must do this manually because a stand-alone CA cannot do it automatically, unlike an enterprise CA. You also can publish the root CA certificate to AD DS by using the certutil command-line tool.

Demonstration: Deploying a Root CA

In this demonstration, you will see how to deploy an enterprise root CA.

Demonstration Steps

Deploy a root CA

1. In the Server Manager, add the **Active Directory Certificate Services** role.
2. Select the **Certification Authority** role service.
3. After the installation completes successfully, click the text **Configure Active Directory Certificate Services on the destination server**.
4. Select to install **Enterprise root CA**.
5. Set the Key length to **4096**.
6. Name the CA **AdatumRootCA**.

Considerations for Deploying a Subordinate CA

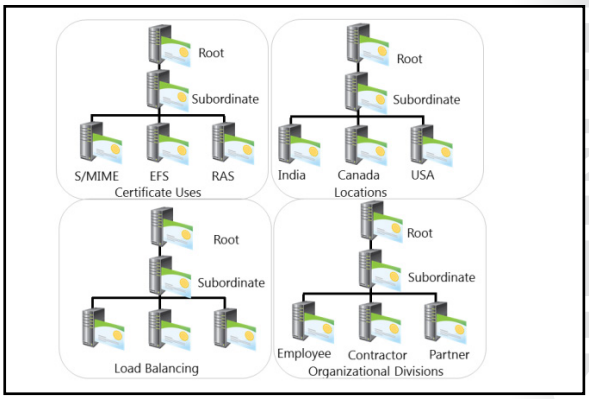

You can use a subordinate CA to implement policy restrictions for PKI, and to distribute certificates to clients. After installing a root CA for the organization, you can install one or more subordinate CAs.

When you use a subordinate CA to distribute certificates to users or computers that have an account in an AD DS environment, you can install the subordinate CA as an enterprise CA. Then, you can use the data from the client accounts in AD DS to distribute and manage certificates, and to publish certificates to AD DS. To complete this procedure, however, you must be a member of the local Administrators group or have equivalent permissions. If the subordinate CA will be an enterprise CA, you also need to be a member of the Domain Admins group or have equivalent permissions. From a security perspective, we recommended that you have an offline root stand-alone CA and an enterprise subordinate CA.

A subordinate CA is usually deployed to achieve some of the following functionalities:

- Usage. You may issue certificates for a number of purposes, such as Secure Multipurpose Internet Mail Extensions (S/MIME), EFS, or Remote Access Service (RAS). The issuing policy for these uses may be distinct, and separation provides a basis for administering these polices.
- Organizational divisions. You may have different policies for issuing certificates, depending upon an entity's role in the organization. You can create subordinate CAs to separate and administer these policies.
- Geographic divisions. Organizations often have entities at multiple physical sites. Limited network connectivity between these sites may require individual subordinate CAs for many or all sites.

- Load balancing. If you will be using your PKI to issue and manage a large number of certificates, having only one CA can result in a considerable network load for that single CA. Using multiple subordinate CAs to issue the same kind of certificates divides the network load among CAs.
- Backup and fault tolerance. Multiple CAs increase the possibility that your network will always have operational CAs available to respond to user requests.

How to Use the CAPolicy.inf File for Installation

If you want to deploy a root or subordinate CA, you want to predefine certain values for use during installation and define certain additional parameters. You can use the CAPolicy.inf file to complete these steps. The CAPolicy.inf file is a plaintext file that contains various settings that are used when installing the AD CS role, or when renewing the CA certificate. The CAPolicy.inf file is not required to install AD CS, but without it, the default settings will be applied, and in many cases, the default settings are insufficient. You can use the CAPolicy.inf file to configure CAs in more complicated deployments.

The CAPolicy.inf file is stored in the %*Windir*% folder of the root or subordinate CA, and defines the following:

- CPS
- Object Identifier
- CRL publication intervals
- CA renewal settings
- Key size
- Certificate validity period
- CDP and AIA paths

Each CAPolicy.inf file is divided into sections, and has a simple structure, which can be described as follows:

- A *section* is an area in the .inf file that contains a logical group of keys. A section always appears in brackets in the .inf file.
- A *key* is the parameter that is to the left of the equal (=) sign.
- A *value* is the parameter that is to the right of the equal (=) sign.

For example, if you want to specify an Authority Information Access point in the CAPolicy.inf file, you use following syntax:

```
[AuthorityInformationAccess]
URL=http://pki.adatum.com/CertData/adatumCA.crt
```

In this example, AuthorityInformationAccess is a section, URL is the key, and http://pki.adatum.com/CertData/adatumCA.crt is the value.

You can also specify some CA server settings in the CAPolicy.inf file. One example of the section that specifies these settings is:

```
[certsrv_server]
RenewalKeyLength=2048
RenewalValidityPeriod=Years
RenewalValidityPeriodUnits=5
CRLPeriod=Days
CRLPeriodUnits=2
CRLDeltaPeriod=Hours
CRLDeltaPeriodUnits=4
ClockSkewMinutes=20
LoadDefaultTemplates=True
AlternateSignatureAlgorithm=0
ForceUTF8=0
EnableKeyCounting=0
```

Note: All parameters from the previous examples are optional.

You can also use the CAPolicy.inf file when installing AD CS to define the following:

- Certification practice statement: Describes the practices that the CA uses to issue certificates. This includes the types of certificates issued, information for issuing, renewing, and recovering certificates, and other details about the CA's configuration.
- Object identifier (OID): Identifies a specific object or attribute.
- CRL publication intervals: Defines the interval between publications for the base CRL.
- CA renewal settings: Defines renewal settings as follows:
  - Key size: Defines the length of the key pair used during the root CA renewal.
  - Certificate validity period: Defines the validity period for a root CA certificate.
  - CDP and AIA paths: Provides the path used for root CA installations and renewals.

Once you have created your CAPolicy.inf file, you must copy it into the %*SystemRoot*% folder of your server (for example, C:\Windows) before you install the AD CS role, or before you renew the CA certificate.

Note: The CAPolicy.inf file is processed for both the root and subordinate CA installations and renewals.

Configuring CA Administration and Security

Role-based administration in AD CS provides the ability to delegate predefined permissions to users or groups based on built-in CA roles. Each role can perform a predetermined task or series of tasks. The following table shows the details of roles and groups involved in role-based administration.

- You can establish role-based administration for CA hierarchy by defining the following roles:
  - CA administrator
  - Certificate manager
  - Backup operator
  - Auditor
  - Enrollees
- You can assign the following permissions on the CA level:
  - Read
  - Issue and Manage Certificates
  - Manage CA
  - Request Certificates
- Certificate managers can be restricted to a template

Role/group	Purpose	Information
CA administrator	Manage the CA	Assigned by using the CA console
Certificate manager	Issue and manage certificates	Assigned by using the CA console
Backup operator	Backup and restore files and directories	Operating system role

Role/group	Purpose	Information
Auditor	Manage auditing and Security Event log	Operating system role
Enrollees	Read and enroll	Can request certificates

Role-based administration combines operating system roles and AD CS roles to provide a complete, segmented management solution for your CAs. Instead of assigning local administrative privileges to the various IT personnel involved in managing the CA, you can assign roles, which ensure that administrators have the minimum permissions necessary to perform their jobs.

Role-based administration also reduces the administrative overhead of granting rights to administrators because the process involves adding a user to a group or role.

Managing CA Security

To manage and configure role-based administration of a CA, and to manage security on the CA, you can use the Security tab of the Certification Authority when you open Properties in the certsrv admin console. The following are security permissions that you can set on a CA object level:

- Read. Security principals assigned with this permission can locate the CA in AD DS or access it by using the web console or services if the stand-alone is deployed.
- Issue and Manage Certificates. Security principals assigned with this permission are able to approve or deny certificate requests that are in a pending state. Also, they are able to revoke an issued certificate, to specify a revoke reason, and to perform an unrevoke. They also are able to read all issued certificates and export them into files.
- Manage CA. Security principals assigned with this permission are able to manage and configure all options on the CA level. They are not able to manage certificates, and can manage only a CA.
- Request Certificates. Security principals assigned with this permission are able to perform certificate requests against this CA. However, this does not mean that they are able to enroll certificates. This is specified on the certificate template level.

Together with defining security permissions on the access control list (ACL) of the CA object, you also can use the Certificate Managers tab on the CA Properties. You can then narrow these security principals to specific certificate templates when you configure security principals that can issue and manage certificates on the Security ACL. For example, if you want to assign user Bob permission to issue and manage only user's certificates, you will put Bob on the ACL and assign it Issue and Manage Certificates permission. However, you will use the Certificate Managers tab to restrict Bob to the User certificate template, because you do not want Bob to be able to issue and manage all certificates.

Configuring CA Policy and Exit Modules

More advanced deployments of CA hierarchies, or scenarios where a CA is integrated with another PKI-related service, require that you configure and manage *policy* and *exit modules* on your CA. Policy and exit modules exist on every CA, stand-alone or enterprise. Each CA is configured with default policy and exit modules, and in most scenarios, you will not have to configure these modules. You can manage both policy and exit modules if you use the CA administrator console. For more complex configuration, however, you must use the certutil command-line tool.

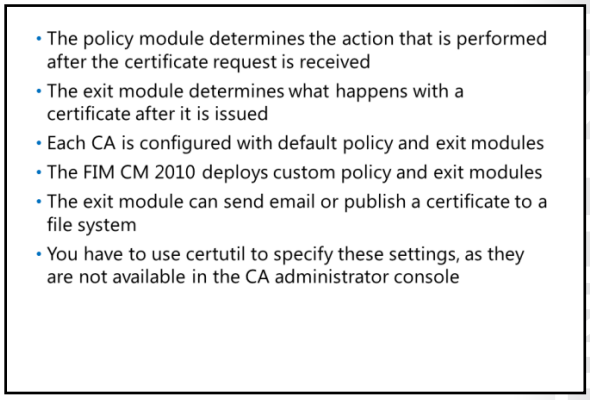

What Is a Policy Module?

A policy module determines the action that is performed after the CA receives the certificate request. You can configure a default policy module to put every certificate request in a pending state until the administrator approves or denies it. The behavior of the default policy module is to issue a certificate if the settings in the certificate template allow it. However, you can install a custom policy module to do other tasks when the CA receives the certificate request. For example, if you install Microsoft Forefront Identity Manager 2010 Certificate Management (FIM CM 2010) in your internal PKI, you will have to deploy the FIM CM 2010 policy module on your CA that issues certificates. FIM CM 2010 manages certification issuing through workflows. Each request for a certificate managed by FIM CM 2010 is forwarded to FIM CM 2010 by the FIM CM 2010 policy module when a CA receives a request. After the request is processed in the workflow on the FIM, the certificate is issued or the request is denied. The FIM CM 2010 policy module also specifies the signature certificate thumbprint for an agent that passed certificate requests from users to a CA. Each request that is signed with a thumbprint specified in the FIM CM 2010 policy module is passed to the FIM workflow before it is issued. This is one example of using the custom policy module, but there also are other third-party applications that might use custom policy modules.

What Is an Exit Module?

Unlike the policy module, the exit module determines what happens with a certificate after it is issued. The most common actions are to send an email or publish a certificate to a file system. These actions are possible even with a default exit module on each CA.

However, you also can deploy a custom policy Exit module. To use the same example as the policy module, if you deploy a FIM CM 2010 in your environment, you also will have to deploy a custom exit module to your CA. The exit module forwards data about each issued certificate to a Microsoft SQL Server specified in the exit module. If you write information about issued certificates to a computer that is running SQL Server, FIM CM is able to view and monitor issued certificates without direct interaction with the CA database.

A CA can use multiple exit modules simultaneously, unlike the policy module, where you can have only one active policy module at a time.

For example, if you want to send an email to a specific address each time the certificate is issued, you have to use certutil to specify these settings, as they are not available in the CA administrator console.

First, you should specify the Simple Mail Transfer Protocol (SMTP) server that is used to send emails, which you can do by typing following certutil command:

```
certutil -setreg exit\smtp\<smtpServerName>
```

You have to enter the fully qualified domain name (FQDN) of your email server instead of the <*smtpServerName*> variable. After this, you have to specify the event and email address to which the notification is sent by typing the following command:

```
certutil -setreg exit\smtp\CRLIssued\To<E-mailString>
```

Note: The exit module on the CA that is configured to send emails on an event does not use SMTP authentication. If your SMTP server requires authentication, you have to configure it on the CA side by typing the following command:

```
certutil -setreg exit\smtp\SMTPAuthenticate 1
certutil -setsmtpinfo<UserName>
```

The <*UserName*> specifies the user name of a valid account on the SMTP server. You will be prompted to provide the password for this user name.

Besides sending notification emails when the certificate is issued, you also can configure an exit module to send notifications of following events:

- Certificate request in pending state.
- Certificate request denied.
- Certificate revoked.
- CRL is issued.
- CA service startup.
- CA service shutdown.

If you want to configure an exit module to publish certificates to the file system, you can use the CA admin console to open the properties of the exit module. After you enable the Allow certificates to be published to the file system option and restart the CA, certificates issued from that CA are copied into the .cer file in the C:\Windows\System32\CertEnroll folder on the CA. However, for this to happen, the certificate requestors must include a **certfile:true** attribute in their request.

If you deploy custom exit modules, their configuration might be possible through the CA admin console or with some other utility.

Demonstration: Configuring CA Properties

In this demonstration, you instructor will show you how to configure CA properties.

Demonstration Steps

1. On LON-SVR1, open the Certification Authority console, and then open the Properties for **AdatumRootCA**.
2. View the Certification Authority certificate.
3. Review the settings for the active policy module.
4. Review the settings for the exit module.
5. Review the values provided on the Extension tab.

6. Review the Security settings of the CA.
7. Review the Certificate Managers settings.

CA Backup and Recovery

As discussed earlier, CAs are designed and configured to work for many years, during which time you might want to upgrade the hardware and operating system that supports the CA. However, because the CA is a critical service in your network infrastructure, it is necessary to have backup and restore procedures defined. Also, scenarios when you need to move a CA role from one computer to another also require that you have a valid and up-to-date backup.

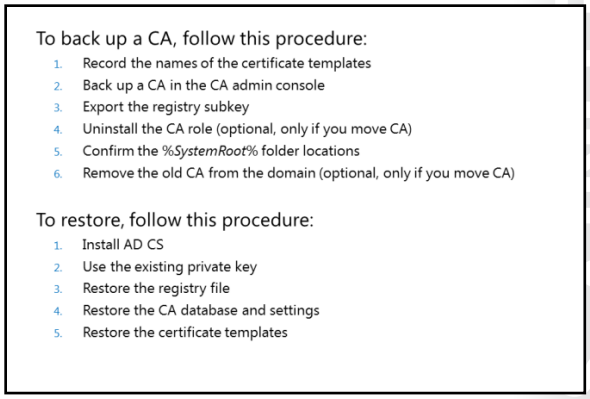

Note: A CA is unlike some other services that you simply can install on a new computer and continue to work. When you move a CA from one computer to another, it is very important that you keep the identity of the CA during this process so that it can continue to work on the new hardware or operating system with the same identity.

Performing a CA Backup

You should have a CA backup even if you are not moving a CA to another computer. A CA backup is different from ordinary backup scenarios. To perform a CA backup to move a CA to another computer, you should perform the following procedure:

1. If you are backing up an enterprise CA, click the **Certificate Templates** item in the CA console, and then record the names of the certificate templates that are listed. These templates are in AD DS, so you do not have to back them up. You must note which templates are published on the CA that you are moving because you will have to add them manually after the CA is moved.
2. In the CA snap-in, right-click the **CA name**, click **All Tasks**, and then click **Back up CA** to start the Certification Authority Backup Wizard. In the backup wizard, you have to select the option to make the backup of the CA's private key, CA certificate, certificate database and certificate database log. You also have to provide an appropriate location for the backup content. CA's private key should be protected with a password for security reasons.
3. After the backup is done, you should open Registry Editor.

   Locate and export the following registry subkey, which is located at: **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\CertSvc\Configuration**.

Note: We recommend that you save this registry key to file in the same folder with the CA backup from the previous step.

4. After this is done, in case you want to move the CA to another computer, you should uninstall the CA from the old server, and then rename the old server or permanently disconnect it from the network.

Before you begin the restore procedure, confirm that the *%SystemRoot%* folder of the target server matches the *%SystemRoot%* folder of the server from which the backup is taken.

In addition, the location of the CA restore must match the location of the CA backup. For example, if you back up the CA from the D:\Windows\System32\Certlog folder, you must restore the backup to the D:\Windows\System32\Certlog folder. After you restore the backup, you can move the CA database files to a different location.

Performing a CA Restore

Restore procedure for CA is initiated when you have to repair your current CA or when you want to move the CA role to another computer.

To restore the CA, perform the following procedure:

1. Install AD CS on the target computer. Select to install either a **Stand-alone** or an **Enterprise** CA, depending on the type of CA that you are moving. When you come to the **Set Up Private Key** page, click **Use existing private key**. Then choose to select a certificate and use its associated private key. This will provide you with ability to use an existing certificate from an old CA.
2. On the **Select Existing Certificate** page, click **Import**, type the path of the .p12 file in the backup folder, type the password that you selected in the previous procedure to protect the backup file, and then click **OK**. When you are prompted for **Public and Private Key Pair**, verify that **Use existing keys** is selected. This is very important, as you want to keep the same root CA certificate.
3. When prompted on the **Certificate Database** page, specify the same location for the certificate database and certificate database log as on the previous CA computer. After you select all these options, wait for the CA setup to finish.
4. After the setup is done, open the Services snap-in to stop the AD CS service. You do this to restore settings from the old CA.
5. Locate the registry file that you saved in the backup procedure, and then double-click it to import the registry settings.
6. After you restore the registry settings, open the CA Management console, right-click the **CA name**, click **All Tasks**, and then click **Restore CA**. This will start the Certification Authority Restore Wizard. In the wizard, you should select the **Private key and CA certificate** and the **Certificate database and certificate database log** check boxes. This specifies that you want to restore these objects from backup. Next, provide a backup folder location and verify the settings for the restore. The **Issued Log** and **Pending Requests** settings should be displayed.
7. When the restore process is done, select to restart the AD CS service.
8. If you restored an enterprise CA, restore the certificate templates from AD DS that you recorded in the previous procedure.

Lab A: Deploying and Configuring CA Hierarchy

Scenario

As A. Datum Corporation has expanded, its security requirements have also increased. The security department is particularly interested in enabling secure access to critical websites, and in providing additional security for features. To address these and other security requirements, A. Datum has decided to implement a PKI using the AD CS role in Windows Server 2012.

As one of the senior network administrators at A. Datum, you are responsible for implementing the AD CS deployment.

Objectives

After completing this lab, you will be able to:

- Deploy a stand-alone root CA.
- Deploy an enterprise Subordinate CA.

Lab Setup

Estimated Time: 50 minutes

Virtual machines	20412C-LON-DC1 20412C-LON-SVR1 20412C-LON-SVR2 20412C-LON-CA1
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 **20412C-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 two and three for 20412C-LON-SVR1, 20412C-LON-SVR2, and 20412C-LON-CA1. Do not sign in until instructed to do so.

Exercise 1: Deploying a Stand-Alone Root CA

Scenario

A. Datum wants to start using certificates for various purposes. You need to install the appropriate CA infrastructure. Because they are using Windows Server 2012 AD DS, you decided to implement the AD CS role. When you were reviewing available designs, you decided to implement a stand-alone root CA. This CA will be taken offline after it issues a certificate for subordinate CA.

The main tasks for this exercise are as follows:

1. Install and configure Active Directory Certificate Services on LON-CA1
2. Creating a DNS host record for LON-CA1 and configure sharing

▶ Task 1: Install and configure Active Directory Certificate Services on LON-CA1

1. Sign in to LON-CA1 as **Administrator** with the password **Pa$$w0rd**.
2. Use the Add Roles and Features Wizard to install the **Active Directory Certificate Services** role.
3. After installation completes successfully, click the text **Configure Active Directory Certificate Services on the destination server**.
4. Configure the AD CS role as a stand-alone root CA. Name it **AdatumRootCA**.
5. Set the key length to **4096**, and then accept all other values as default.
6. On LON-CA1, open the **Certification Authority** console.
7. Open the **Properties** dialog box for **AdatumRootCA**.
8. Configure new locations for the CDP to be **http://lon-svr1.adatum.com/CertData/<CaName><CRLNameSuffix><DeltaCRLAllowed>.crl**.
9. Select the following options:
   - **Include in the CDP extension of issued certificates**
   - **Include in CRLs**. **Clients use this to find Delta CRL locations**
10. Configure new locations for AIA to be on **http://lon-svr1.adatum.com/CertData/<ServerDNSName>_<CaName><CertificateName>.crt**
11. Select the **Include in the AIA extension of issued certificates** check box.
12. Publish the CRL on LON-CA1.
13. Export the root CA certificate, and then copy the .cer file to **\\lon-svr1\C$**.
14. Copy the contents of folder **C:\Windows\System32\CertSrv\CertEnroll** to **\\lon-svr1\C$**.

▶ Task 2: Creating a DNS host record for LON-CA1 and configure sharing

1. On LON-DC1, open the DNS Manager console.
2. Create a host record for LON-CA1 in the Adatum.com forward lookup zone.
3. Use IP address **172.16.0.40** for the LON-CA1 host record.
4. On LON-CA1, from the Network and Sharing Center, turn on file and printer sharing on guest and public networks.

Results: After completing this exercise, you will have deployed a root stand-alone CA.

Exercise 2: Deploying an Enterprise Subordinate CA

Scenario

After you deploy the stand-alone root CA, the next step is to deploy an enterprise subordinate CA. A. Datum wants to use an enterprise subordinate CA to utilize AD DS integration. In addition, because the root CA is stand-alone, you want to publish its certificate to all clients.

The main tasks for this exercise are as follows:

1. Install and configure AD CS on LON-SVR1
2. Install a subordinate CA certificate
3. Publish the root CA certificate through Group Policy
4. Prepare for the next lab

▶ Task 1: Install and configure AD CS on LON-SVR1

1. Sign in to LON-SVR1 as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. Install the Active Directory Certificate Services role on LON-SVR1. Include the **Certification Authority** and **Certification Authority Web Enrollment** role services.
3. After installation is successful, click **Configure Active Directory Certificate Services on the destination server**.
4. Select the **Certification Authority** and **Certification Authority Web Enrollment** role services.
5. Configure LON-SVR1 to be an **Enterprise CA**.
6. Configure the CA Type to be a **Subordinate CA**.
7. For the CA Name, type **Adatum-IssuingCA**.
8. Save the request file to the local drive.

▶ Task 2: Install a subordinate CA certificate

1. On LON-SVR1, install the **C:\RootCA.cer** certificate to the Trusted Root Certification Authority store.
2. Navigate to **Local Disk (C:),** and copy the **AdatumRootCA.crl** and **LON-CA1_AdatumRootCA.crt** files to **C:\inetpub\wwwroot\CertData**.
3. Copy the **LON-SVR1.Adatum.com_Adatum-LON-SVR1-CA.req** request file to **\\LON-CA1\C$**.
4. Switch to LON-CA1.
5. From the Certification Authority console on LON-CA1, submit a new certificate request by using the .req file that you copied in step 3.
6. Issue the certificate, and then export it to .p7b format with a complete chain. Save the file to **\\lon-svr1\C$\SubCA.p7b**.
7. Switch to LON-SVR1.
8. Install the subordinate CA certificate on LON-SVR1 by using the Certification Authority console.
9. Start the service.

▶ Task 3: Publish the root CA certificate through Group Policy

1. On LON-DC1, from the Server Manager, open the Group Policy Management Console.
2. Edit the Default Domain Policy.
3. Publish the **RootCA.cer** file from **\\lon-svr1\C$** to the Trusted Root Certification Authorities store, which is located in **Computer Configuration\Policies\Windows Settings\Security Settings\Public Key Policies**.

▶ Task 4: Prepare for the next lab

- Keep all virtual machines running for the next lab. Do not revert any virtual machines.

Results: After completing this exercise, you will have deployed and configured an enterprise subordinate CA.

Question: Why is it not recommended to install just an enterprise root CA?

Lesson 4
Deploying and Managing Certificate Templates

Certificate templates define how a certificate can be requested and how it can be used. Templates are configured on the CA, and they are stored in the Active Directory database. There are different versions of templates. The Windows 2000 Server Enterprise CA supports version 1 certificate templates, the Windows Server 2003 Enterprise edition supports versions 1 and 2 templates, and the Windows Server 2008 Enterprise edition supports versions 1, 2, and 3 certificate templates. Windows Server 2012 introduces version 4 templates, but still also supports all three previous template versions.

Two types of certificate template categories are users and computers, and each can be used for multiple purposes. You can assign Full Control, Read, Write, Enroll, and Autoenroll permissions to certificate templates. You can also update certificate templates by modifying the original certificate template, copying a template, or superseding existing certificate templates. In this lesson, you will learn how to manage and deploy certificate templates.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe certificate and certificate templates.
- Describe certificate template versions in Windows Server 2012.
- Configure certificate template permissions.
- Configure certificate template settings.
- Describe options for updating a certificate template.
- Explain how to modify and enable a certificate template.

What Are Certificate and Certificate Templates?

A *certificate* is a small amount of data that contains several pieces of information about its owner. This data can include the owner's email address, the owner's name, the certificate usage type, the validity period, and the URLs for AIA and CDP locations. A certificate also contains the key pair, the private key and its related public key. These keys are used in processes of validating identities, digital signatures, and encryption. The key pair that is generated with each certificate works under the following conditions:

A certificate contains information about users, devices, usage, validity, and a key pair

A certificate template defines:

- The format and contents of a certificate
- The process for creating and submitting a valid certificate request
- The security principals that are allowed to read, enroll, or use autoenrollment for a certificate that will be based on the template
- The permissions required to modify a certificate template

- When content is encrypted with the public key, it can be decrypted only with the private key.
- When content is encrypted with the private key, it can be decrypted only with the public key.
- There is no other key that is in the same relation with the keys from a single key pair.
- The private key cannot be derived in a reasonable amount of time from a public key, or vice versa.

During the enrollment process, key pair is generated on the client, and then the public key is sent with the certificate signing request (CSR) to the CA. The CA validates the CSR, and then signs the public key with

the CAs private key. The signed public key is returned to the requestor. This ensures that the private key never leaves the system (or smart card) and that the certificate is trusted by a CA because the CA signed the public key of the certificate. Certificates provide a mechanism for gaining confidence in the relationship between a public key and the entity that owns the corresponding private key.

You can think of a certificate as being similar to a driver's license. A driver's license is accepted by numerous businesses as a form of identification because the license issuer (a government institution) is accepted by the community as trustworthy. Because businesses understand the process by which someone can obtain a driver's license, they trust that the issuer has verified the identity of the individual to whom the license was issued. Therefore, the driver's license can be accepted as a valid form of identification. A certificate trust is established in a similar way.

Certificate Templates

Certificate templates allow administrators to customize the distribution method of certificates, define certificate purposes, and mandate the type of usage allowed by a certificate. Administrators can create templates and then can deploy them quickly to an enterprise by using built-in GUI or command-line management utilities.

Associated with each certificate template is its DACL, which defines which security principals have permissions to read and configure the template, and which security principals can enroll or use autoenrollment for certificates based on the template. Certificate templates and their permissions are defined in AD DS and are valid within the forest. If more than one enterprise CA is running in the AD DS forest, permission changes will affect all CAs.

When you define a certificate template, the definition of the certificate template must be available to all CAs in the forest. This is accomplished when you store the certificate template information in the configuration naming context of AD DS. The replication of this information depends on the AD DS replication schedule, and the certificate template might not be available to all CAs until replication completes. Storage and replication occur automatically.

Note: Prior to Windows Server 2008 R2, only the Enterprise editions of Windows Server supported management of certificate templates. In Windows Server 2008 R2 and Windows Server 2012, you also can manage certificate templates in the Standard editions.

Certificate Template Versions in Windows Server 2012

The CA in Windows Server 2012 CA supports four versions of certificate templates. Certificate templates versions 1, 2 and 3 are legacy from previous versions of Windows Server, while version 4 is new to Windows Server 2012.

Certificate template versions correspond to the Windows Server operating system version. Windows 2000 Server, Windows Server 2003, Windows Server 2008, and Windows Server 2012 correspond to version 1, version 2, version 3, and version 4, respectively.

Version 1:
- Introduced in Windows 2000 Server, provides for backward compatibility in newer versions
- Creates by default when a CA is installed
- Cannot be modified (except for permissions) or removed, but can be duplicated to become version 2 or 3 templates, which can then be modified

Version 2:
- Default template introduced with Windows Server 2003
- Allows customization of most settings in the template
- Several preconfigured templates are provided when a CA is installed

Version 3:
- Supports advanced Suite B cryptographic settings
- Includes advanced options for encryption, digital signatures, key exchange, and hashing
- Only supports Windows Server 2008 and Windows Server 2008 R2 servers
- Only supports Windows Vista and Windows 7 client computers

Version 4:
- Available only for Windows Server 2012 and Windows 8 clients
- Supports both CSPs and KSPs
- Supports renewal with the same key

Aside from corresponding with Windows Server operating system versions, certificate template versions also have some functional differences, as follows:

- Windows 2000 Advanced Server operating system provides support for version 1 certificate templates. The only modification allowed to version 1 templates is changing permissions to either

allow or disallow enrollment of the certificate template. When you install an enterprise CA, version 1 certificate templates are created by default. As of July 13, 2010, Windows 2000 Server is no longer supported by Microsoft.

- Windows Server 2003 Enterprise Edition operating systems provide support for version 1 and version 2 templates. You can customize several settings in the version 2 templates. The default installation provides several preconfigured version 2 templates. You can add version 2 templates based on the requirements of your organization. Alternatively, you can duplicate a version 1 certificate template to create a new version 2 of the template. You can then modify and secure the newly created version 2 certificate template. When new templates are added to a Windows Server 2003 Enterprise CA, they are version 2 by default.
- Windows Server 2008 Enterprise operating systems bring support for new version 3 certificate templates. Additionally, support for version 1 and version 2 is provided. Version 3 certificate templates support several features of a Windows Server 2008 enterprise CA, such as CNG. CNG provides support for Suite B cryptographic algorithms such as elliptic curve cryptography (ECC). In a Windows Server 2008 Enterprise, you can duplicate default version 1 and version 2 templates to bring them up to version 3.
- Windows Server 2008 provides two new certificate templates by default: Kerberos authentication and OCSP Response Signing. The Windows Server 2008 R2 operating system version was also able to support certificate templates. When you use version 3 certificate templates, you can use CNG encryption and hash algorithms for the certificate requests, issued certificates, and protection of private keys for key exchange and key archival scenarios.
- Windows Server 2012 operating systems provide support for version 4 certificate templates, and for all other versions from earlier editions of Windows Server. These certificate templates are available only to Windows Server 2012 and Windows 8. To help administrators separate the features supported by each operating system version, the Compatibility tab was added to the certificate template Properties tab. It marks options as unavailable in the certificate template properties, depending upon the selected operating system versions of certificate client and CA. Version 4 certificate templates also support both CSPs and key storage providers. They can also be configured to require renewal with a same key.

Upgrading certificate templates is a process that applies only in situations where the CA has been upgraded from Windows Server 2008 or Windows Server 2008 R2 to Windows Server 2012. After the upgrade, you can upgrade the certificate templates by launching the CA Manager console and clicking Yes at the upgrade prompt.

Configuring Certificate Template Permissions

To configure certificate template permissions, you need to define the DACL on the Security tab for each certificate template. The permissions that are assigned to a certificate template will define which users or groups can read, modify, enroll, or autoenroll for that certificate template.

Permissions	Description
Full Control	Allows a designated user, group, or computer to modify all attributes—including ownership and permissions
Read	Allows a designated user, group, or computer to read the certificate in AD DS when enrolling
Write	Allows a designated user, group, or computer to modify all attributes except permissions
Enroll	Allows a designated user, group, or computer to enroll for the certificate template
Autoenroll	Allows a designated user, group, or computer to receive a certificate through the autoenrollment process

You can assign the following permissions to certificate templates:

- Full Control. The Full Control permission allows a security principal to modify all attributes of a certificate template, which includes permissions for the certificate template itself. It also includes permission to modify the security descriptor of the certificate template.
- Read. The Read permission allows a user or computer to view the certificate template when enrolling for certificates. The Read permission is also required by the certificate server to find the certificate templates in AD DS.
- Write. The Write permission allows a user or computer to modify the attributes of a certificate template, which includes permissions assigned to the certificate template itself.
- Enroll. The Enroll permission allows a user or computer to enroll for a certificate based on the certificate template. However, to enroll for a certificate, you must also have Read permissions for the certificate template.
- Autoenroll. The Autoenroll permission allows a user or computer to receive a certificate through the autoenrollment process. However, the Autoenroll permission requires the user or computer to also have both Read and Enroll permissions for a certificate template.

As a best practice, you should assign certificate template permissions to global or universal groups only. This is because the certificate template objects are stored in the configuration naming context in AD DS. You should avoid assigning certificate template permissions to individual users or computer accounts.

As a best practice, keep the Read permission allocated to the Authenticated Users group. This permission allocation enables all users and computers to view the certificate templates in AD DS. This permission assignment also enables the CA that is running under the System context of a computer account to view the certificate templates when assigning certificates. This permission, however, does not grant Enroll rights, so it is safe to have it configured this way.

Configuring Certificate Template Settings

Besides configuring security settings for certificate templates, you can also configure several other settings for each template. Be aware however, that the number of configurable options depends on the certificate template version. For example, the version 1 certificate templates do not allow modification of any settings except for security, while certificate templates from higher versions allow you to configure most of the available options.

For each certificate template, you can customize several settings, such as validity time, purpose, CSP, private key exportability, and issuance requirements

Category	Single purpose examples	Multiple purpose examples
Users	• Basic EFS • Authenticated session • Smart card logon	• Administrator • User • Smart card user
Computers	• Web server • IPsec	• Computer • Domain controller

Windows Server 2012 provides several default certificate templates for these purposes: code-signing (for digitally signing software), EFS (for encrypting data), and the ability for users to sign in with a smart card. To customize a template for your company, duplicate the template and then modify the certificate configuration.

For example, you can configure the following:

- Format and content of a certificate based on the certificate's intended use.

Note: The intended use of a certificate may relate to users or to computers, based on the types of security implementations that are required to use the PKI.

- Process of creating and submitting a valid certificate request.
- CSP supported.
- Key length.
- Validity period.
- Enrollment process or enrollment requirements.

You can also define a certificate purpose in certificate settings. Certificate templates can have the following purposes:

- Single Purpose. A single purpose certificate serves a single purpose, such as allowing users to sign in with a smart card. Organizations utilize single purpose certificates in cases where the certificate configuration differs from other certificates that are being deployed. For example, if all users will receive a certificate for smart card logon but only a couple of groups will receive a certificate for EFS, organizations will generally keep these certificates and templates separate to ensure that users only receive the required certificates.
- Multiple Purposes. A multipurpose certificate serves more than one purpose (often unrelated) at the same time. While some templates (such as the User template) serve multiple purposes by default, organizations will often modify templates to serve additional purposes. For example, if a company intends to issue certificates for three purposes, those purposes can be combined into a single certificate template to ease administrative effort and maintenance.

Options for Updating a Certificate Template

The CA hierarchy in most organizations has one certificate template for each job function. For example, there may be a certificate template for file encryption and another for code signing. Additionally, there may be a few templates that cover functions for most of the common groups of subjects.

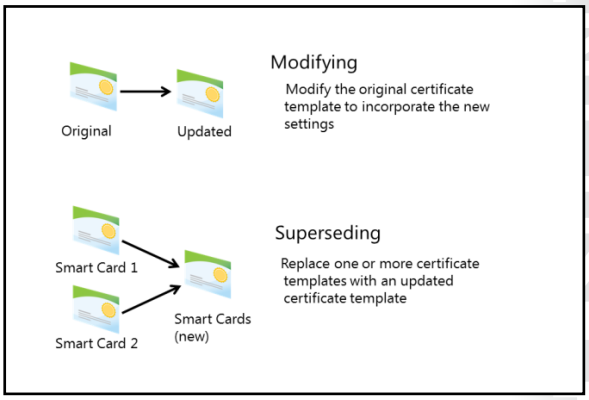

As an IT administrator, you may need to modify an existing certificate template because of incorrect settings or other issues in the original certificate template. You may also need to merge multiple existing certificate templates into a single template.

You can update a certificate template by either modifying the template or superseding the existing template:

- Modify the original certificate template. To modify a certificate template of version 2, 3, or 4, you need to make changes and then apply them to that template. After this, any certificate issued by a CA based on that certificate template will include the modifications that you made.
- Supersede existing certificate templates. The CA hierarchy of an organization may have multiple certificate templates that provide the same or similar functionality. In such a scenario, you can supersede or replace the multiple certificate templates by using a single certificate template. You can make this replacement in the Certificate Templates console by designating that a new certificate template supersedes, or replaces, the existing certificate templates. Another benefit of superseding the template is that the new version will be used when a certificate expires.

Demonstration: Modifying and Enabling a Certificate Template

In this demonstration, you will see how to modify and enable a certificate template.

Demonstration Steps

Modify and enable a certificate template

1. On LON-SVR1, open the Certificate Templates console.
2. Review the list of available templates.
3. Open the **IPsec** certificate template **Properties**, and review available settings.
4. Duplicate the **Exchange User** certificate template. Name it **Exchange User Test1**, and then configure it to supersede the **Exchange User** template.
5. Allow Authenticated Users to enroll for the **Exchange User Test1** template.
6. Publish the template on LON-SVR1.

Lesson 5
Implementing Certificate Distribution and Revocation

One of the steps in deploying PKI in your organization is to define methods for certificate distribution and enrollment. In addition, during the certificate management process, there will be times that you need to revoke certificates. Reasons for revoking certificates might include a key becoming compromised, or someone leaving the organization. You need to ensure that network clients can determine which certificates are revoked before accepting authentication requests. To ensure scalability and high availability, you can deploy the AD CS Online Responder, which provides certificate revocation status. In this lesson, you will learn about methods for certificate distribution and certificate revocation.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe options for certificate enrollment.
- Describe how autoenrollment works.
- Describe the Enrollment Agent.
- Explain how to configure the Restricted Enrollment Agent.
- Describe the NDES.
- Explain how certificate revocation works.
- Describe considerations for publishing AIAs and CDPs.
- Describe an Online Responder.
- Configure Online Responder.

Options for Certificate Enrollment

In Windows Server 2012, you can use several methods to enroll for a user or computer certificate. The use of these methods depends on specific scenarios. For example, you would most likely use autoenrollment to mass-deploy certificates to a large number of users or computers, while you would most likely use manual enrollment for certificates dedicated to specific security principals.

Method	Use
Autoenrollment	• To automate the request, retrieval, and storage of certificates for domain-based computers
Manual enrollment	• To request certificates by using the Certificates Templates console or Certreq.exe when the requestor cannot communicate directly with the CA
CA Web enrollment	• To request certificates from a website that is located on a CA • To issue certificates when autoenrollment is not available
Enroll on behalf	• To provide IT staff with the right to request certificates on behalf of another user (Enrollment Agent)

The following list describes the different enrollment methods, and when to use them:

- Autoenrollment. Using this method, the administrator defines the permissions and the configuration of a certificate template. These definitions help the requestor to request, retrieve, and renew certificates automatically without end-user interaction. This method is used for AD DS domain computers. The certificate must be configured for autoenrollment through Group Policy.
- Manual enrollment. Using this method, the private key and a certificate request are generated on a device, such as a web service or a computer. The certificate request is then transported to the CA to generate the certificate being requested. The certificate is then transported back to the device for

installation. Use this method when the requestor cannot communicate directly with the CA, or if the device does not support autoenrollment.

- CA Web enrollment. Using this method, you can enable a website CA so that users can obtain certificates. To use CA Web enrollment, you must install IIS and the web enrollment role on the CA of AD CS. To obtain a certificate, the requestor logs on to the website, selects the appropriate certificate template, and then submits a request. The certificate is issued automatically if the user has the appropriate permissions to enroll for the certificate. The CA Web enrollment method should be used to issue certificates when autoenrollment cannot be used. This can happen in the case of an Advanced Certificate request. However, there are cases where autoenrollment can be used for certain certificates, but not for all certificates.
- Enrollment on behalf (Enrollment Agent). Using this method, a CA administrator creates an Enrollment Agent account for a user. The user with Enrollment Agent rights can then enroll for certificates on behalf of other users. You would use this method, for example, if you need to allow a manager to preload logon certificates of new employees on to smart cards.

How Does Autoenrollment Work?

One of the most common methods for deploying certificates in an Active Directory environment is to use autoenrollment. This method provides an automated way to deploy certificates to both users and computers within the PKI. You can use autoenrollment in environments that meet specific requirements, such as the use of certificate templates and Group Policy in AD DS. It is important to note, however, that you cannot use autoenrollment with a stand-alone CA. You must have an AD DS integrated enterprise CA available to make use of autoenrollment.

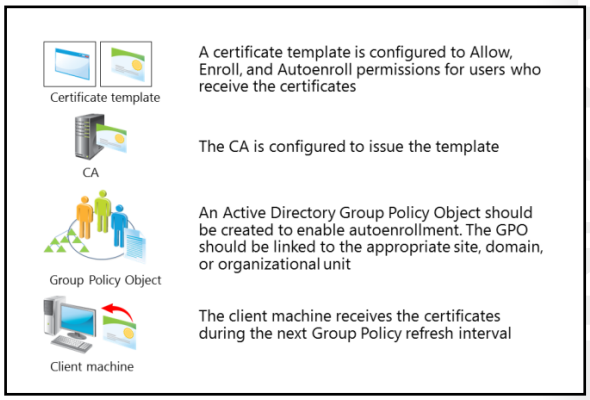

You can use autoenrollment to automatically deploy public key–based certificates to users and computers in an organization. The Certificate Services administrator duplicates a certificate template, and then configures the permissions to allow Enroll and Autoenroll permissions for the users or computers who will receive the certificates. Domain-based Group Policies, such as computer-based and user-based policies, can activate and manage autoenrollment.

By default, Group Policy is applied when you restart computers, or at logon for users. Also by default, Group Policy is refreshed every 90 minutes on domain members. This Group Policy setting is named Certificate Services Client - AutoEnrollment.

An internal timer triggers autoenrollment every eight hours after the last autoenrollment activation. The certificate template might specify user interaction for each request. For such a request, a pop-up window displays approximately 60 seconds after the user logs on.

Many certificates can be distributed without the client even being aware that enrollment is occurring. These include most types of certificates that are issued to computers and services, as well as many certificates issued to users.

To enroll clients automatically for certificates in a domain environment, you must:

- Have membership in Domain Admins or Enterprise Admins, or equivalent, which is the minimum required to complete this procedure.

- Configure a certificate template with Autoenroll permissions.
- Configure an autoenrollment policy for the domain.

What Is Credential Roaming?

Credential Roaming allows organizations to store certificates and private keys in AD DS, separately from the application state or configuration information. Credential Roaming uses existing logon and autoenrollment mechanisms to download certificates and keys to a local computer whenever a user logs on and, if desired, remove them when the user logs off. In addition, the integrity of these credentials is maintained under any conditions, such as when certificates are updated, or when users sign in to more than one computer at a time. This avoids the scenario where a user is autoenrolled for a certificate on each new machine to which he or she logs on.

Credential Roaming is triggered any time a private key or certificate in the user's local certificate store changes, whenever the user locks or unlocks the computer, and whenever Group Policy is refreshed.

All certificate-related communication between components on the local computer and between the local computer and AD DS is signed and encrypted. Credential Roaming is supported in Windows 7 and newer Windows operating systems.

Enrollment Agent Overview

In the Windows Server 2012 CA, it is possible to configure certificate enrollment on behalf of another user. To do this, you must have a specific certificate issued. This certificate is based on the *Enrollment Agent* template. When a user gets a certificate based on an Enrollment Agent template, he or she has the ability to enroll for a certificate on behalf of another user. Unlike a Certificate Manager, an Enrollment Agent can only process the enrollment request and cannot approve pending requests or revoke issued certificates.

An Enrollment Agent is a user who has the appropriate certificate assigned and and has the ability to request certificates on behalf of other users or computers

The restricted Enrollment Agent has limited permissions:

- Limits permissions of the Enrollment Agent:
  - For specific group of users
  - For specific certificate templates
- Requires Windows Server 2008 Enterprise edition or Windows Server 2012 CA

Note: Enrollment Agent is a certificate with very high security risk. A person who has an Enrollment Agent certificate can impersonate other users, as he or she is able to issue a certificate for other users—for example, smart card certificates. Make sure that you protect this certificate template.

Windows Server 2012 includes the following three certificate templates that enable different types of Enrollment Agents:

- Enrollment Agent. Used to request certificates on behalf of another subject.
- Enrollment Agent (Computer). Used to request certificates on behalf of another computer subject.
- Exchange Enrollment Agent (Offline Request). Used to request certificates on behalf of another subject and supply the subject name in the request. This template is used by the NDES for its Enrollment Agent certificate.

When you create an Enrollment Agent, you can further refine the agent's ability to enroll for certificates on behalf of others by a group and by a certificate template. For example, you might want to implement a

restriction that the Enrollment Agent can enroll for smart card logon certificates only and just for users in a certain office or organizational unit (OU) that is the basis for a security group.

In older versions of Windows Server CA, it was not possible to permit an Enrollment Agent to enroll only a certain group of users. As a result, every user with an Enrollment Agent certificate was able to enroll on behalf of any user in an organization.

The restricted Enrollment Agent is functionality that was introduced in the Windows Server 2008 Enterprise edition operating system. This functionality allows you to limit the permissions for users who are designated as Enrollment Agents in enrolling smart card certificates on behalf of other users.

Typically, one or more authorized individuals within an organization are designated as Enrollment Agents. The Enrollment Agent needs to be issued an Enrollment Agent certificate, which enables the agent to enroll for certificates on behalf of users. Enrollment agents typically are members of corporate security, IT security, or help desk teams, because these individuals already have been entrusted with safeguarding valuable resources. In some organizations, such as banks that have many branches, help desk and security workers might not be conveniently located to perform this task. In this case, designating a branch manager or another trusted employee to act as an Enrollment Agent is required to enable smart card credentials to be issued from multiple locations.

On a Windows Server 2012 CA, the restricted Enrollment Agent features allow an Enrollment Agent to be used for one or many certificate templates. For each certificate template, you can choose the users or security groups on behalf of whom the Enrollment Agent can enroll. You cannot constrain an Enrollment Agent based on a certain Active Directory OU or container; instead, you must use security groups.

 **Note:** Using restricted Enrollment Agents will affect the performance of the CA. To optimize performance, you should minimize the number of accounts that are listed as Enrollment Agents. You minimize the number of accounts in the Enrollment Agent's certificate template ACL. As a best practice, use group accounts in both lists instead of individual user accounts.

Demonstration: Configuring the Restricted Enrollment Agent

In this demonstration, you will see how to configure the restricted Enrollment Agent.

Demonstration Steps

Configure the Restricted Enrollment Agent

1. On LON-SVR1, open the Certificate Templates console.
2. Configure **Allie Bellew** with permissions to enroll for an Enrollment Agent certificate.
3. Publish the Enrollment Agent certificate template.
4. Sign in to LON-CL1 as **Adatum\Allie** with the password **Pa$$w0rd**.
5. Open a Microsoft Management console (MMC), and add the Certificates snap-in.
6. Request the Enrollment Agent certificate.
7. Switch to LON-SVR1, and open the properties of **AdatumRootCA**.
8. Configure the restricted Enrollment Agent so that Allie can only issue certificates based on the **User** template, and only for the **Marketing** security group.

What Is NDES?

The Network Device Enrollment Service (NDES) is the Microsoft implementation of Simple Certificate Enrollment Protocol (SCEP). SCEP is a communication protocol that makes it possible for software that is running on network devices such as routers and switches—which cannot otherwise be authenticated on the network—to enroll for X.509 certificates from a CA.

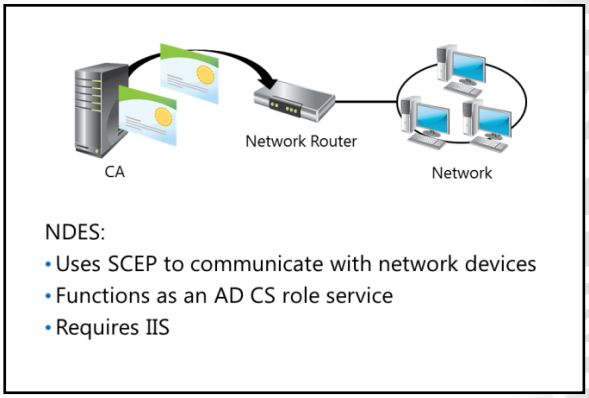

You can use NDES as an Internet Server API filter on IIS to perform the following functions:

- Create and provide one-time enrollment passwords to administrators.
- Retrieve awaiting requests from the CA.
- Collect and process SCEP enrollment requests for the software that runs on network devices.

This feature applies to organizations that have PKIs with one or more Windows Server 2012–based CAs and that want to enhance the security for their network devices. Port security, based on 802.1x, requires certificates be installed on switches and access points. Secure Shell, instead of Telnet, requires a certificate on the router, switch, or access point. NDES is the service that allows administrators to install certificates on devices using SCEP.

Adding support for NDES can enhance the flexibility and scalability of an organization's PKI. Therefore, PKI architects, planners, and administrators may be interested in this feature.

Before installing NDES, you must decide:

- Whether to set up a dedicated user account for the service, or use the Network Service account.
- The name of the NDES registration authority and the country/region to use. This information is included in any SCEP certificates that are issued.
- The CSP to use for the signature key that is used to encrypt communication between the CA and the registration authority.
- The CSP to use for the encryption key that is used to encrypt communication between the registration authority and the network device.
- The key length for each of these keys.

In addition, you need to create and configure the certificate templates for the certificates that are used in conjunction with NDES.

When you install NDES on a computer, this creates a new registration authority and deletes any preexisting registration authority certificates on the computer. Therefore, if you plan to install NDES on a computer where another registration authority has already been configured, any pending certificate requests should be processed and any unclaimed certificates should be claimed before you install NDES.

How Does Certificate Revocation Work?

Revocation is the process by which you disable validity of one or more certificates. By initiating the revoke process, you actually publish a certificate thumbprint in the corresponding CRL.

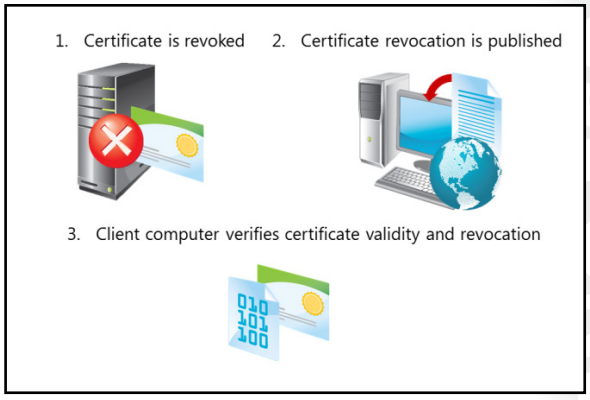

An overview of the certificate revocation life cycle is outlined as follows:

- A certificate is revoked from the CA Microsoft Management Console (MMC) snap-in. During revocation, a reason code and a date and time are specified. This is optional, but it is recommended.
- The CRL is published using the CA MMC snap-in (or the scheduled revocation list is published automatically based on the configured value). CRLs can be published in AD DS, in some shared folder location, or on a website.
- When Windows client computers are presented with a certificate, they use a process to verify revocation status by querying the issuing CA. This process determines whether the certificate is revoked, and then presents the information to the application that requested the verification. The Windows client computer uses one of the CRL locations specified in certificate to check its validity.

The Windows operating systems include a CryptoAPI, which is responsible for the certificate revocation and status checking processes. The CryptoAPI utilizes the following phases in the certificate checking process:

- Certificate discovery. Certificate discovery collects CA certificates, AIA information in issued certificates, and details of the certificate enrollment process.
- Path validation. Path validation is the process of verifying the certificate through the CA chain (or *path*) until the root CA certificate is reached.
- Revocation checking. Each certificate in the certificate chain is verified to ensure that none of the certificates is revoked.
- Network retrieval and caching. Network retrieval is performed by using OCSP. CryptoAPI is responsible for checking the local cache first for revocation information and if there is no match, making a call using OCSP, which is based on the URL provided by the issued certificate.

Considerations for Publishing AIAs and CDPs

When you manage and issue certificates, it is important that you properly configure the certificate extensions that are used to verify the certificate of the CA and the certificate that is used by the user, computer, or device. These extensions—AIA and CDP—are part of each certificate. They must point to proper locations, or the PKI might not function correctly.

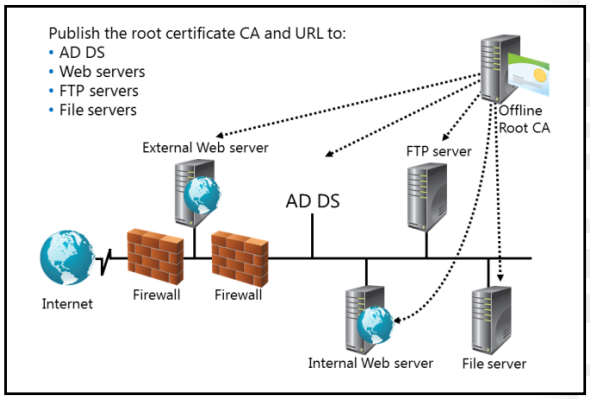

Note: One common cause of CA hierarchy malfunctions and downtime is improperly configured AIA and CDP extensions. Make sure that you configure these options properly before you put your CA hierarchy in production.

What Is AIA?

AIA addresses are the URLs in the certificates that a CA issues. These addresses tell the verifier of a certificate where to retrieve the CA's certificate. AIA access URLs can be HTTP, File Transfer Protocol (FTP), Lightweight Directory Access Protocol (LDAP), or FILE addresses.

What Is CDP?

The CDP is a certificate extension that indicates from where the CRL for a CA can be retrieved. It can contain none, one, or many HTTP, FTP, FILE, or LDAP addresses.

Each certificate that you issue from your CA contains information about the CDP and AIA location. Each time a certificate is used, these locations are checked. The AIA location is checked to verify the validity of the CA certificate, while the CDP location is checked to verify content of the CRL for that CA. At least one AIA and one CDP location must be available for each certificate. If they are not available, the system will presume that the certificate is not valid, the revocation check will fail, and you will not be able to use that certificate for any purpose.

AIA and CDP Publishing

If you only use an online CA, these values are configured by default locally on the CA. However, if you want to deploy an offline root CA, or if you want to publish AIA and CDP to an Internet-facing location, you must reconfigure these values so that they apply to all the certificates issued by the root CA. The AIA and CDP extensions define where client applications can locate AIA and CDP information for the root CA. The formatting and publishing of AIA and CDP extension URLs are generally the same for root CAs and subordinate CAs. You can publish the root CA certificate and the CRL to the following locations:

- AD DS
- Web servers
- FTP servers
- File servers

Publication Points

To ensure accessibility to all computers in the forest, publish the offline root CA certificate and the offline Root CA's CRL to AD DS by using the certutil command. This places the root CA certificate and the CRL in the Configuration naming context, which is then replicated to all domain controllers in the forest.

For computers that are not members of an AD DS domain, place the CA certificate and the CRL on web servers by using the HTTP protocol. Locate the web servers on the internal network, and on the external network if external client computers—or the internal clients from the external networks—require access. This is very important if you use internally issued certificates outside of your company.

You also can publish certificates and CRLs to the ftp:// and file:// URLs, but we recommend that you use only the LDAP and HTTP URLs because they are the most widely supported URL formats for interoperability purposes. The order in which you list the CDP and AIA extensions is important, because the certificate-chaining engine searches the URLs sequentially. If your certificates are mostly used internally in an AD DS domain, place the LDAP URL first in the list.

Note: Besides configuring CDP and AIA publication points, you also should make sure that the CRL is valid. An online CA will automatically renew the CRL periodically, but an offline root CA will not. If the offline root CA CRL expires, the revocation check will fail. To prevent failure, make

sure that you configure the validity period for the offline root CA CRL to be long enough, and set a reminder to turn that CA on and issue a new CRL before the old one expires.

What Is an Online Responder?

By using OCSP, an Online Responder provides clients with an efficient way to determine the revocation status of a certificate. OCSP submits certificate status requests using HTTP.

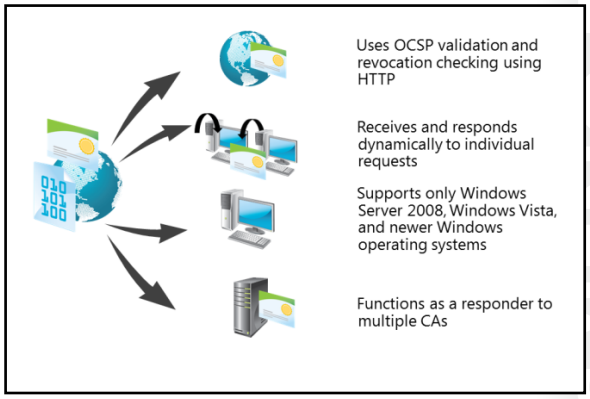

Clients access CRLs to determine the revocation status of a certificate. CRLs might be large, and clients might utilize a large amount of time to search through these CRLs. An Online Responder can dynamically search these CRLs for the clients, and respond only to the requested certificate.

You can use a single Online Responder to determine revocation status information for certificates that are issued by a single CA, or by multiple CAs. However, you can use more than one Online Responder to distribute CA revocation information.

You can install an Online Responder on any computer that runs Windows Server 2008 Enterprise or Windows Server 2012. You should install an Online Responder and a CA on different computers.

The following operating systems can use Online Responder for validation of certificate status:

- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012
- Windows Vista
- Windows 7
- Windows 8

For scalability and high availability, you can deploy the Online Responder in a load-balanced array using Network Load Balancing, which processes certificate status requests. You can monitor and manage each member of the array independently. To configure the Online Responder, you must use the Online Responder management console.

You must configure the CAs to include the URL of the Online Responder in the AIA extension of issued certificates. The OCSP client uses this URL to validate the certificate status. You must also issue the OCSP Response Signing certificate template, so that the Online Responder can also enroll that certificate.

How to Install and Configure Online Responder

You can install Online Responders on computers that are running Windows Server 2008 R2 or Windows Server 2012. You should install Online Responders after the CAs, but prior to issuing any client certificates. The certificate revocation data is derived from a published CRL. The published CRL can come from a CA on a computer that is running Windows Server 2008 or newer, or Windows Server 2003, or from a CA other that Microsoft.

Before configuring a CA to support the Online Responder service, the following must be present:

- IIS must be installed on the computer during the Online Responder installation. When you install an Online Responder, the correct configuration of IIS for the Online Responder is installed automatically.
- An OCSP Response Signing certificate template must be configured on the CA, and autoenrollment used to issue an OCSP Response Signing certificate to the computer on which the Online Responder will be installed.
- The URL for the Online Responder must be included in the AIA extension of certificates issued by the CA. This URL is used by the Online Responder client to validate certificate status.

After an Online Responder has been installed, you need to create a revocation configuration for each CA and CA certificate that is served by an Online Responder. A revocation configuration includes all of the settings that are needed to respond to status requests regarding certificates that have been issued using a specific CA key. These configuration settings include:

- CA certificate. This certificate can be located on a domain controller, in the local certificate store, or imported from a file.
- Signing certificate for the Online Responder. This certificate can be selected automatically for you, selected manually (which involves a separate import step after you add the revocation configuration), or you can use the selected CA certificate.
- Revocation provider. The revocation provider will provide the revocation data used by this configuration. This information is entered as one or more URLs where the valid base and delta CRLs can be obtained.

Demonstration: Configuring an Online Responder

In this demonstration, you will see how to configure an Online Responder.

Demonstration Steps

Configure an Online Responder

1. On LON-SVR1, use the Server Manager to add an Online Responder role service to the existing AD CS role.
2. Configure a new AIA distribution location on **AdatumRootCA** to be **http://lon-svr1/ocsp**.
3. On **AdatumRootCA**, publish the OCSP Response signing certificate template, and allow Authenticated users to enroll.
4. Open the Online Responder Management console.
5. Add revocation configuration for **AdatumRootCA**.
6. Enroll for an OCSP Response signing certificate.
7. Ensure that the revocation configuration status displays as working.

Lesson 6
Managing Certificate Recovery

During the certificate life cycle, certificate or key recovery is one of the most important management tasks. If you lose your public and private keys, you use a key archival and recovery agent for data recovery. You can also use automatic or manual key archival and key recovery methods to ensure that you can gain access to data in the event that your keys are lost. In this lesson, you will learn how to manage key archival and recovery in AD CS in Windows Server 2012.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the process of key archival and recovery.
- Configure Automatic Key Archival.
- Configure a CA for Key Archival.
- Explain key recovery.
- Recover a lost private key.

Overview of Key Archival and Recovery

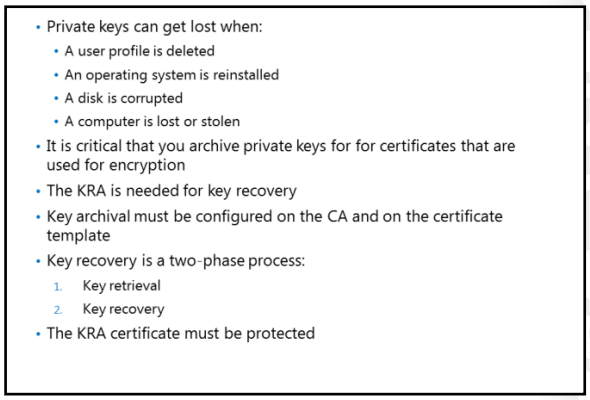

Keeping the certificate and the corresponding key pair secure can be critical in some scenarios. For example, if you use a certificate to perform content encryption of emails or documents, and you lose your public and private keys, you will not be able to access any data that is encrypted by using the certificate's public key. This data can include EFS-encrypted data and S/MIME–protected emails. Therefore, archival and recovery of public and private keys are important. You can archive, or back up, your private key by exporting a certificate with a private key and storing it in a secure location, such as an alternative media, or certain cloud-based storage. However, this approach requires that each user backs up his or her private key, which usually is not a reliable backup method. An alternative method is to centralize private key archival on the CA.

Note: In regular operations, CA does not have access to a user's private key, because it is generated on the client side. Because of this, the archival of private keys must be enabled explicitly on each certificate template where you want to have this functionality.

Conditions for Losing Keys

You might lose key pairs due to the following conditions:

- A user profile is deleted or corrupted. A CSP encrypts a private key and stores the encrypted private key in the local file system and registry in the user profile folder. Deletion or corruption of the profile results in the loss of the private key material.

- An operating system is reinstalled. When you reinstall the operating system, the previous installations of the user profiles are lost, including the private key material. In this scenario, the computer's certificates also are lost.
- A disk is corrupted. If a hard disk becomes corrupted and the user profile is unavailable, the private key material is lost automatically, in addition to installed computer certificates.
- A computer is lost or stolen. If a user's computer is lost or stolen, the user profile with the private key material is unavailable.

Note: Losing a key pair (certificate) is not always a critical situation. For example, if you lose a certificate used for digital signing or logging, you simply can issue a new one, and no data will be affected. However, losing a certificate that was used for data encryption will result in the inability to access data. For that reason, archival and recovery is critical.

Key Archival and Recovery Agents

To use private key archival, you must enable this functionality on both the CA and specific certificate templates, such as EFS. This functionality is not enabled by default on the CA or on any certificate template. To be able to archive private keys from certificates, you also must define a KRA.

Note: Key archival on the CA works from the moment you fully configure it. It does not apply, however, to certificates issued before this functionality was enabled.

You use key archival and KRAs for data recovery in scenarios where the private key is lost. The KRA is the user with the KRA certificate issued who is able to decrypt private keys stored in a Certificate Services database. When key archival is enabled on the CA and on certificate templates, each private key is encrypted with a KRS's public key and then stored in the CA database. As a result, a KRA's private key must be used to decrypt the private key on any user. KRAs are designated to users who are able to retrieve the original certificate, private key, and public key that were used to encrypt the data.

Note: The KRA should not be mixed with the Data Recovery Agent. The Data Recovery Agent is able to decrypt data encrypted with EFS directly in case the originating user's private key is not available. The KRA, on the other hand, does not decrypt any data directly; it just decrypts archived private keys. Data Recovery Agent functionality is discussed earlier in this module.

To become a KRA, you must enroll with a certificate based on the KRA template. After this certificate is issued to the designated user, a public key from the KRA's certificate is imported on the CA, and key archival is enabled. From that moment, each certificate that is issued based on a template where key archival is enabled will have its private key stored in the CA database and encrypted with the KRA's public key.

During the key recovery process, the Certificate Manager or CA administrator retrieves the encrypted file that contains the certificate and private key from the CA database. Next, a KRA uses its private key to decrypt the private key from the encrypted file and then returns the certificate and private key to the user.

Note: Key recovery is a two-phase process. First, the encrypted key is retrieved from CA database. Second, the KRA decrypts the key and certificate. For security reasons, we recommend that these two phases be performed by different people. By default, the KRA does not have permission to retrieve encrypted keys from a CA database.

Security for Key Archival

When you have a configured CA to issue a KRA certificate, any user with Read and Enroll permission on the KRA certificate template can enroll and become a KRA. Domain Admins and Enterprise Admins receive permissions by default. However, you must ensure the following:

- Only trusted users are allowed to enroll for this certificate.
- The KRA's private key is stored in a secure manner.
- The server where the keys are archived is in a separate, physically secure location.

After the KRA certificate is issued, we recommend that you remove this template from the CA. Also, we recommend that you import the KRA certificate only when a key recovery procedure should be performed.

Understanding Key Archival and Recovery

Key recovery implies that the private key portion of a public-private key pair might be archived and recovered. Private key recovery does not recover any data or messages. It merely enables a user to retrieve lost or damaged keys, or for an administrator to assume the role of a user for data access or data recovery purposes. In many applications, data recovery cannot occur without first performing key recovery.

The key recovery procedure is as follows:

1. The user requests a certificate from a CA and provides a copy of the private key as part of the request. The CA, which processes the request, archives the encrypted private key in the CA database and issues a certificate to the requesting user.
2. An application such as EFS can use the issued certificate to encrypt sensitive files.
3. If, at some point, the private key is lost or damaged, the user can contact the organization's Certificate Manager to recover the private key. The Certificate Manager, with the help of the KRA, recovers the private key, stores it in a protected file format, and then sends it back to the user.
4. After the user stores the recovered private key in the user's local keys store, the key once again can be used by an application such as EFS to decrypt previously encrypted files or to encrypt new ones

Configuring Automatic Key Archival

Before you can use key archival, you must perform several configuration steps. The key archival feature is not enabled by default, and you should configure both CA and certificate templates for key archival and key recovery.

The following steps describe the automatic key archival process:

1. Configure the KRA certificate template. Only Enterprise Admins or Domain Admins are allowed to request a KRA certificate. If you want to enroll some other user with a KRA certificate, you must specify it on the template DACL.

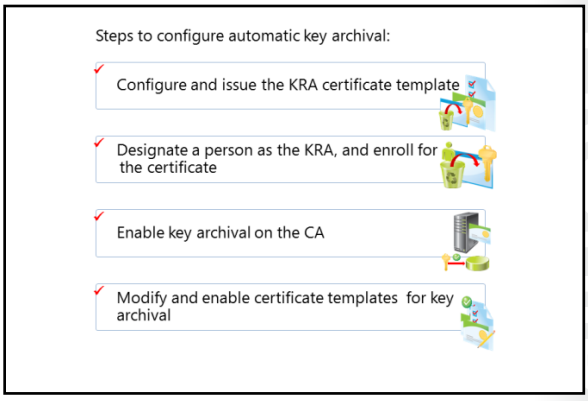

2. Configure Certificate Managers.
    - CA enforces a person to be a Certificate Manager, if defined. The Certificate Manager usually holds a private key for valid KRA certificates. By default, the CA Administrator is a Certificate Manager for all users, except for cases with another explicit definition. However, as a best practice, you should separate these two roles if possible.
    - A CA Officer is defined as a Certificate Manager. This user has the security permission to issue and manage certificates. The security permissions are configured on a CA in the CA MMC snap-in, in the CA Properties dialog box, from the Security tab.
    - A KRA is not necessarily a CA Officer or a Certificate Manager. These roles may be segmented as separate roles. A KRA is a person who holds a private key for a valid KRA certificate.
3. Enable KRA.
    - Sign in as the Administrator of the server, or as the CA Administrator if role separation is enabled.
    - In the CA console, right-click the CA name, and then click **Properties**. To enable key archival, on the **Recovery Agents** tab, click **Archive the key**.
    - By default, the CA uses one KRA. However, you must first select the KRA certificate for the CA to begin archival by clicking **Add**.
    - The system finds valid KRA certificates, and then displays available KRA certificates. These are generally published to AD DS by an enterprise CA during enrollment. KRA certificates are stored under the KRA container in the Public Key Services branch of the configuration partition in AD DS. Because CA issues multiple KRA certificates, each KRA certificate will be added to the multivalued user attribute of the CA object.
    - Select one certificate, and then click **OK**. Ensure that you have selected the intended certificate.
    - After you have added one or more KRA certificates, click **OK**. KRA certificates are only processed at service start.
4. Configure user templates.
    - In the Certificate Templates MMC, right-click the key archival template, and then click **Properties**.
    - To always enforce key archival for the CA, in the **Properties** dialog box, on the **Request Handling** tab, select the **Archive subject's encryption private key** check box. In Windows Server 2008 or newer CAs, select the **Use advanced symmetric algorithm to send the key to the CA** option.

Demonstration: Configuring a CA for Key Archival

In this demonstration, you will see how to configure a CA for key archival.

Demonstration Steps

Configure automatic key archival

1. Configure **adatumRootCA** to issue Key Recovery Agent certificates without approval.
2. Enroll Administrator for Key Recovery Agent certificate.
3. Configure **adatumRootCA** to use the certificate enrolled in step 2 as Key Recovery Agent.

4. Configure the **Exchange User Test 1** certificate template to allow key archival.
5. Configure **adatumRootCA** to allow key archival.

Recovering a Lost Key

Key recovery consists of several steps, and you must strictly follow the procedure to recover archived keys. The procedure for key recovery is as follows:

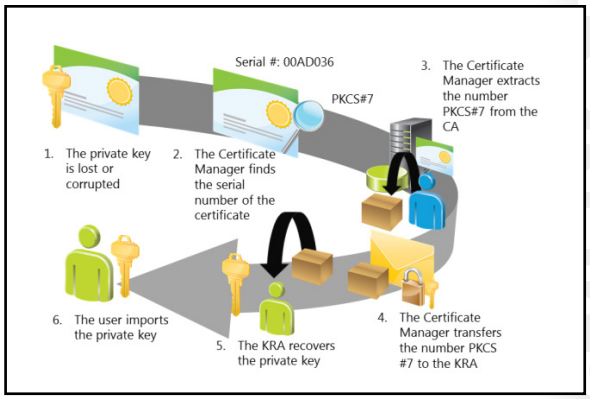

1. Find recovery candidates. You will require two pieces of information to perform key recovery. First, the Certificate Manager or the CA Administrator locates the correct certificate entry in the CA database. Next, the Certificate Manager or the CA Administrator obtains the serial number of the correct certificate entry and the KRA certificate required for key recovery.
2. Retrieve PKCS #7 BLOB from the database. This is the first half of the key recovery step. A Certificate Manager or a CA Administrator retrieves the correct BLOB from the CA database. The certificate and the encrypted private key to be recovered are present in PKCS #7 BLOB. The private key is encrypted alongside the public key of one or more KRAs.
3. Recover key material and save to PKCS #12 (.pfx). This is the second half of the key recovery step. The holder of one of the KRA private keys decrypts the private key to be recovered. The holder also generates a password-protected .pfx file that contains the certificate and private key.
4. Import recovered keys. The password-protected .pfx file is delivered to the end user. This user imports the .pfx file into the local user certificate store. Alternatively, the KRA or an administrator can perform this part of the procedure on behalf of the user.

Demonstration: Recovering a Lost Private Key

In this demonstration, you will see how to recover a lost private key.

Demonstration Steps

Recover a lost private key

1. Enroll Administrator for Exchange User Test1 certificate.
2. Delete the certificate from Administrator personal store to simulate key loss.
3. On LON-SVR1 in the CA console, retrieve the serial number of the lost certificate.

```
Use command Certutil -getkey <serialnumber> outputblob to generate blob file.
Use command Certutil -recoverkey outputblob recover.pfx, to recover the private key.
```

4. Import the private key back to the administrator personal store.

Lab B: Deploying and Managing Certificates

Scenario

As A. Datum Corporation has expanded, its security requirements have also increased. The security department is particularly interested in enabling secure access to critical websites, and in providing additional security for features such as drive encryption, smart cards, and the Windows 7 and Windows 8 DirectAccess feature. To address these and other security requirements, A. Datum has decided to implement a PKI using the AD CS role in Windows Server 2012.

As one of the senior network administrators at A. Datum, you are responsible for implementing the AD CS deployment. You will deploy the CA hierarchy, develop the procedures and process for managing certificate templates, and deploy and revoke certificates.

Objectives

After completing this lab, you will be able to:

- Configure certificate templates.
- Configure certificate enrollment.
- Configure certificate revocation.
- Configure and perform private key archival and recovery.

Lab Setup

Estimated Time: 75 minutes

Virtual machines	20412C-LON-DC1 20412C-LON-SVR1 20412C-LON-SVR2 20412C-LON-CA1 20412C-LON-CL1
User name	**Adatum\Administrator**
Password	**Pa$$w0rd**

For this lab, you will use the available virtual machine environment. All virtual machines needed for this lab should be running from the previous lab.

Exercise 1: Configuring Certificate Templates

Scenario

After deploying the CA infrastructure, the next step is to deploy the certificate templates that are required in the organization. First, A. Datum wants to implement a new Web server certificate and implement smart card certificates for users. They also want to implement new certificates on the LON-SVR2 web server.

The main tasks for this exercise are as follows:

1. Create a new template based on the web server template

2. Create a new template for users that includes smart card logon

3. Configure the templates so that they can be issued
4. Update the web server certificate on the LON-SVR2 web server

▶ Task 1: Create a new template based on the web server template

1. On LON-SVR1, from the Certification Authority console, open the Certificate Templates console.
2. Duplicate the **Web Server** template.
3. Create a new template and name it **Adatum WebSrv**.
4. Configure validity for **3 years**.
5. Configure the private key as exportable.

▶ Task 2: Create a new template for users that includes smart card logon

1. In the Certificate Templates console, duplicate the **User** certificate template.
2. Name the new template **Adatum User**.
3. On the **Subject Name** tab, clear both the **Include e-mail name in subject name** and the **E-mail name** check boxes.
4. Add **Smart Card Logon** to the Application Policies of the new certificate template.
5. Configure this new template to supersede the **User** template.
6. Allow Authenticated Users to **Read**, **Enroll**, and **Autoenroll** for this certificate.
7. Close the Certificate Templates console.

▶ Task 3: Configure the templates so that they can be issued

- Configure LON-SVR1 to issue certificates based on the **Adatum User** and **Adatum WebSrv** templates.

▶ Task 4: Update the web server certificate on the LON-SVR2 web server

1. Sign in to LON-SVR2 as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. Refresh Group Policy, and restart the server if necessary.
3. From the Server Manager, open the **Internet Information Services (IIS) Manager**.
4. Enroll for a domain certificate using the following parameters:
    - Common name: **lon-svr2.adatum.com**
    - Organization: **Adatum**
    - Organizational Unit: **IT**
    - City/locality: **Seattle**
    - State/province: **WA**
    - Country/region: **US**
    - Friendly name: **lon-svr2**
5. Create HTTPS binding for the Default Web Site, and associate it with a new certificate.

Results: After completing this exercise, you will have created and published new certificate templates.

Exercise 2: Configuring Certificate Enrollment

Scenario

The next step in implementing the PKI at A. Datum is to configure certificate enrollment. A. Datum wants to enable different options for distributing the certificates. Users should be able to enroll automatically, and smart card users should get their smart cards from an Enrollment Agent. A. Datum has delegated Enrollment Agent rights for the Marketing department group to user Allie Bellew.

The main tasks for this exercise are as follows:

1. Configure autoenrollment for users
2. Verify autoenrollment
3. Configure the Enrollment Agent for smart card certificates

▶ Task 1: Configure autoenrollment for users

1. On LON-DC1, open Group Policy Management.
2. Edit the Default Domain Policy.
3. Navigate to **User Configuration**/**Policies**/**Windows Settings**/**Security Settings**, and then click to highlight **Public Key Policies**.
4. Enable the **Certificate Services Client – Auto-Enrollment** option, and enable **Renew expired certificates, update pending certificates, and remove revoked certificates** and **Update certificates that use certificate templates**.
5. Enable the **Certificate Services Client – Certificate Enrollment Policy**.
6. Close Group Policy Management Editor and GPMC.

▶ Task 2: Verify autoenrollment

1. On LON-SVR1, open the Windows PowerShell and use **gpupdate /force** to refresh Group Policy.
2. Open an mmc.exe console and add the Certificates snap-in focused on the user account.
3. Verify that you have been issued a certificate based on the **Adatum Smart Card User** template.

▶ Task 3: Configure the Enrollment Agent for smart card certificates

1. On LON-SVR1, from the Certification Authority console, open the Certificate Templates console.
2. Allow **Allie Bellew** to enroll for an Enrollment Agent certificate.
3. Publish the Enrollment Agent certificate template.
4. Sign in to **LON-CL1** as **Allie**, and enroll for an Enrollment Agent certificate.
5. On LON-SVR1, open properties of **Adatum-IssuingCA**, and configure **Restricted Enrollment Agent** so that Allie can only issue certificates based on **Adatum User**, for security group **Marketing**.

Results: After completing this exercise, you will have configured and verified autoenrollment for users, and configured an Enrollment Agent for smart cards.

Exercise 3: Configuring Certificate Revocation

Scenario

As part of configuring the certificate infrastructure, A. Datum wants to configure revocation components on newly established CAs. You will configure CRL and Online Responder components.

The main tasks for this exercise are as follows:

1. Configure Certified Revocation List (CRL) distribution
2. Install and configure an Online Responder

▶ Task 1: Configure Certified Revocation List (CRL) distribution

1. On LON-SVR1, in the Certification Authority console, right-click **Revoked Certificates**, and then click **Properties**.
2. Set the CRL publication interval to **1 Days**, and set the **Delta CRL** publication interval to **1 Hours**.
3. Review CDP locations on **Adatum-IssuingCA**.

▶ Task 2: Install and configure an Online Responder

1. On LON-SVR1, use the Server Manager to add an Online Responder role service to the existing AD CS role.
2. When the message displays that installation succeeded, click **Configure Active Directory Certificate Services on the destination server**.
3. Configure the online responder.
4. On LON-SVR1, open the Certification Authority console.
5. Configure the new AIA distribution location on **Adatum-IssuingCA** to be **http://lon-svr1/ocsp**.
6. On **Adatum-IssuingCA**, publish the OCSP Response signing certificate template, and allow **Authenticated** users to enroll.
7. Open the Online Responder Management console.
8. Add revocation configuration for **Adatum-IssuingCA**.
9. Enroll for an OCSP Response signing certificate.
10. Ensure that revocation configuration status is **Working**.

Results: After completing this exercise, you will have configured certificate revocation settings.

Exercise 4: Configuring Key Recovery

Scenario

As a part of establishing a PKI, you want to configure and test procedures for the recovery of private keys. You want to assign a KRA certificate for an administrator, and configure CA and specific certificate templates to allow key archiving. In addition, you want to test a procedure for key recovery.

The main tasks for this exercise are as follows:

1. Configure the CA to issue KRA certificates
2. Acquire the KRA certificate
3. Configure the CA to allow key recovery
4. Configure a custom template for key archival
5. Verify key archival functionality
6. Prepare for the next module

▶ Task 1: Configure the CA to issue KRA certificates

1. On LON-SVR1, in the Certification Authority console, right-click the **Certificates Templates** folder, and then click **Manage**.
2. In the Certificates Templates console, open the **Key Recovery Agent certificate properties** dialog box.
3. On the **Issuance Requirements** tab, clear the **CA certificate manager approval** check box.
4. On the **Security** tab, notice that only Domain Admins and Enterprise Admins groups have the **Enroll** permission.
5. Right-click the **Certificates Templates** folder, and enable the **Key Recovery Agent** template.

▶ Task 2: Acquire the KRA certificate

1. Create an MMC console window that includes having the Certificates snap-in for the current user loaded.
2. Use the Certificate Enrollment Wizard to request a new certificate, and to enroll the KRA certificate.
3. Refresh the console window, and view the KRA in the personal store.

▶ Task 3: Configure the CA to allow key recovery

1. On LON-SVR1, open the Certification Authority console.
2. On LON-SVR1, in the Certification Authority console, open the **Adatum-IssuingCA Properties** dialog box.
3. On the **Recovery Agents** tab, click **Archive the key**, and then add the certificate by using the **Key Recovery Agent Selection** dialog box.
4. Restart Certificate Services when prompted.

▶ Task 4: Configure a custom template for key archival

1. On LON-SVR1, open the Certificates Templates console.
2. Duplicate the User template, and name it **Archive User**.
3. On the **Request Handling** tab, set the option for the **Archive subject's encryption private key**. Using the archive key option, the KRA can obtain the private key from the certificate store.
4. Click the **Subject Name** tab, and clear both the **E-mail name** and **Include e-mail name in subject name** check boxes.
5. Add the **Archive User template** as a new certificate template to issue.

▶ Task 5: Verify key archival functionality

1. Sign in to LON-CL1 as **Adatum\Aidan**, using the password **Pa$$w0rd**.
2. Create an MMC console window that includes the Certificates snap-in.
3. Request and enroll a new certificate based on the **Archive User** template.
4. From the personal store, locate the Archive User certificate.
5. Delete the certificate for Aidan to simulate a lost key.
6. Switch to LON-SVR1.
7. Open the Certification Authority console, expand **Adatum-IssuingCA**, and then click the **Issued Certificates** store.

8. In the Certificate Authority console, note the serial number of the certificate that has been issued for Aidan Delaney.
9. On LON-SVR1, open a command prompt, and type the following command:

```
Certutil –getkey <serial number> outputblob
```

Note: Replace serial number with the serial number that you wrote down.

10. Verify that the **Outputblob** file now displays in the C:\Users\Administrator folder.
11. To convert the **Outputblob** file into an importable .pfx file, at the command prompt, type the following command:

```
Certutil–recoverkey outputblob aidan.pfx
```

12. Enter the password **Pa$$w0rd** for the certificate.
13. Verify the creation of the recovered key in the C:\Users\Administrator folder.
14. Switch to LON-CL1 machine.
15. Open Network and Sharing Center item in Control panel, and then enable file and printer sharing for guest or public network profiles.
16. Switch back to LON-SVR1, and then copy and paste the **aidan.pfx** file to the root of drive C on LON-CL1.
17. Switch to LON-CL1, and import the **aidan.pfx** certificate.
18. Verify that the certificate displays in the Personal store.

▶ Task 6: Prepare for the next module

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

1. On the host computer, start **Hyper-V Manager**.
2. On the **Virtual Machines** list, right-click 20412C-LON-DC1, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps two and three for 20412C-LON-CL1, 20412C-LON-SVR1, 20412C-LON-CA1, and 20412C-LON-SVR2.

Results: After completing this exercise, you will have implemented key archival and tested private key recovery.

Question: What is the main benefit of OCSP over CRL?

Question: What must you do to recover private keys?

Module Review and Takeaways

Review Questions

Question: What are some reasons that an organization would utilize PKI?

Question: What are some reasons that an organization would use an enterprise root CA?

Question: List the requirements to use autoenrollment for certificates.

Question: What are the steps to configure an Online Responder?

Real-world Issues and Scenarios

Contoso, Ltd. wants to deploy PKI to support and secure several services. They have decided to use Windows Server 2012 Certificate Services as a platform for PKI. Certificates will be primarily used for EFS, digital signing, and for web servers. Because documents that will be encrypted are important, it is crucial to have a disaster recovery strategy in case of key loss. In addition, clients that will access secure parts of the company website must not receive any warning in their browsers.

1. What kind of deployment should Contoso select?
2. What kind of certificates should Contoso use for EFS and digital signing?
3. What kind of certificates should Contoso use for a website?
4. How will Contoso ensure that EFS-encrypted data is not lost if a user loses a certificate?

Tools

- Certificate Authority console
- Certificate Templates console
- Certificates console
- Certutil.exe

Best Practice:

- When you deploy CA infrastructure, deploy a stand-alone (non-domain-joined) root CA, and an enterprise subordinate CA (issuing CA). After the enterprise subordinate CA receives a certificate from root CA, take root CA offline.
- Issue a certificate for root CA for a long period of time, such as 15 or 20 years.
- Use autoenrollment for certificates that are widely used.
- Use a Restricted Enrollment Agent whenever possible.
- Use Virtual Smart Cards to improve logon security.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
The location of the CA certificate that is specified in the AIA extension is not configured to include the certificate name suffix. Clients may not be able to locate the correct version of the issuing CA's certificate to build a certificate chain, and certificate validation may fail.	Use the Certification Authority snap-in to configure the AIA extension to include the certificate name suffix in each location.
CA is not configured to include CRL distribution point locations in the extensions of issued certificates. Clients may not be able to locate a CRL to check the revocation status of a certificate, and certificate validation may fail.	Use the CA snap-in to configure the CRL distribution point extension and to specify the network location of the CRL. The default locations of the CRL are added to the CRL distribution point extension settings during CA installation, and the CA is configured to include the default locations in the extensions of all issued certificates.
CA was installed as an enterprise CA, but Group Policy settings for user autoenrollment have not been enabled. An enterprise CA can use autoenrollment to simplify certificate issuance and renewal. If autoenrollment is not enabled, certificate issuance and renewal may not occur as expected.	Use the Group Policy Management Console to configure user autoenrollment policy settings, and use the Certificate Templates snap-in to configure autoenrollment settings on the certificate template.

Module 7

Implementing Active Directory Rights Management Services

Contents:

Module Overview

Active Directory® Rights Management Services (AD RMS) provides a method for protecting content that goes beyond encrypting storage devices using Windows® BitLocker® Drive Encryption, or encrypting individual files using Encrypting File System (EFS). AD RMS provides a method to protect data in transit and at rest, and ensures that it is accessible only to authorized users for a specific duration.

This module introduces you to AD RMS. It also describes how to deploy AD RMS, how to configure content protection, and how to make AD RMS-protected documents available to external users.

Objectives

After completing this module, you will be able to:

- Provide an overview of AD RMS.
- Deploy and manage an AD RMS infrastructure.
- Configure AD RMS content protection.
- Configure external access to AD RMS.

Lesson 1
AD RMS Overview

Before you deploy AD RMS, you need to know how AD RMS works, what components are included in an AD RMS deployment, and how you should deploy AD RMS. You must also understand the concepts behind various AD RMS certificates and licenses.

This lesson provides an overview of AD RMS, and reviews the scenarios in which you can use it to protect your organization's confidential data.

Lesson Objectives

After completing this lesson you will be able to:

- Describe AD RMS.
- Explain the scenarios in which you can use AD RMS.
- List the AD RMS components.
- List the different AD RMS certificates and licenses.
- Explain how AD RMS works.

What Is AD RMS?

AD RMS is an information protection technology that is designed to minimize the possibility of data leakage. *Data leakage* is the unauthorized transmission of information—either to people within the organization or people outside the organization—who should not be able to access that information. AD RMS integrates with existing Microsoft products and operating systems including Windows Server®, Microsoft Exchange Server®, Microsoft SharePoint® Server, and the Microsoft Office Suite.

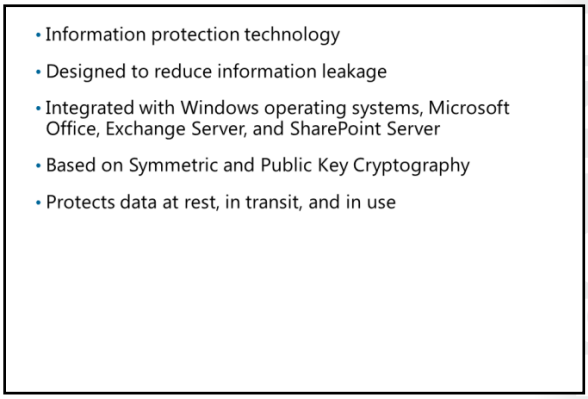

AD RMS can protect data in transit and at rest. For example, AD RMS can protect documents that are sent as email messages by ensuring that a message cannot be opened even if it is accidentally addressed to the wrong recipient. You can also use AD RMS to protect data that is stored on devices such as removable USB drives. A drawback of file and folder permissions is that once the file is copied to another location, the original permissions no longer apply. A file that is copied to a USB drive will inherit the permissions on the destination device. Once copied, a file that was read-only can be made editable by altering the file and folder permissions.

With AD RMS, the file can be protected in any location, irrespective of file and folder permissions that grant access. With AD RMS, only the users who are authorized to open the file will be able to view the contents of that file. The author can decide which permissions (read, write, print) do apply to whom and for which timeframe.

Usage Scenarios for AD RMS

The primary use for AD RMS is to control the distribution of sensitive information. You can use AD RMS in combination with encryption techniques to secure data when it is in storage or in transit. There may be many reasons to control the distribution of sensitive information, such as needing to ensure that only authorized staff members have access to a file, ensuring that sensitive email messages cannot be forwarded, or ensuring that details of an unreleased project are not made public. Consider the following scenarios:

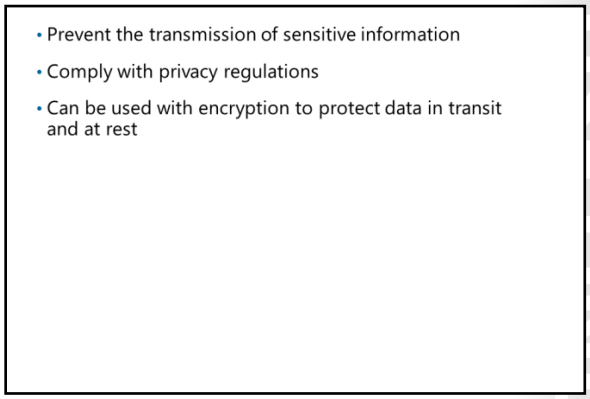

Scenario 1

The chief executive officer (CEO) copies a spreadsheet file containing the compensation packages of an organization's executives from a protected folder on a file server to the CEO's personal USB drive. During the commute home, the CEO leaves the USB drive in the limousine, where someone with no connection to the organization finds it. Without AD RMS, whoever finds the USB drive can open the file. With AD RMS, it is possible to ensure that the file cannot be opened by unauthorized users.

Scenario 2

An internal document should be viewable by a group of authorized people within the organization. These people should not be able to edit or print the document. While you can use the native functionality of Microsoft Office Word to restrict these features, by using a password for each document, you must remember different passwords for potentially hundreds of documents. With AD RMS, you can configure these permissions based on existing accounts in Active Directory® Domain Services (AD DS) or even share with business partners through other means.

Scenario 3

People within the organization should not be able to forward sensitive email messages that have been assigned a particular classification. With AD RMS, you can allow a sender to assign a particular classification to a new email message, and that classification will ensure that the recipient cannot forward the message.

Overview of the AD RMS Components

The AD RMS root certification cluster is the first AD RMS server that you deploy in a forest. The AD RMS root certification cluster manages all licensing and certification traffic for the domain in which it is installed. AD RMS stores configuration information either in a Microsoft SQL Server® database or in the Windows Internal Database (WID). In large environments, the SQL Server database is hosted on a server that is separate from the server that hosts the AD RMS role.

AD RMS licensing-only clusters are used in distributed environments. Licensing-only clusters

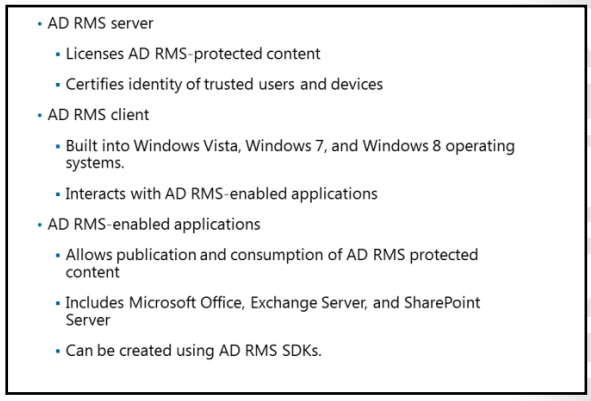

do not provide certification, but do allow the distribution of licenses that are used for content consumption and publishing. Licensing-only clusters are often deployed to large branch offices in organizations that use AD RMS.

AD RMS Server

AD RMS servers must be members of an AD DS domain. When you install AD RMS, information about the location of the cluster is published to AD DS to a location known as the service connection point. Computers that are members of the domain query the service connection point to determine the location of AD RMS services.

AD RMS Client

AD RMS client is built into the Windows Vista®, Windows 7®, and Windows 8® operating systems. The AD RMS client allows AD RMS-enabled applications to enforce the functionality dictated by the AD RMS template. Without the AD RMS client, AD RMS-enabled applications would be unable to interact with AD RMS-protected content.

AD RMS Enabled Applications

AD RMS-enabled applications allow users to create and consume AD RMS–protected content. For example, Microsoft Outlook® allows users to view and create protected email messages. Office Word allows uses to view and create protected word processing documents. Microsoft provides an AD RMS software development kit (SDK) to allow developers to enable their applications to support AD RMS protection of content.

AD RMS Certificates and Licenses

To understand how AD RMS works, you need to be familiar with its various certificates and license types. Each of these certificates and licenses functions in a different way. Some certificates, such as the server licensor certificate (SLC), are critically important, and you must back them up on a regular basis.

- AD RMS certificate and licenses include:
  - Server licensor certificate
  - AD RMS machine certificate
  - Rights Account Certificate
  - Client licensor certificate
  - Publishing license
  - End-user license

SLC

The SLC is generated when you create the AD RMS cluster. It has a validity of 250 years. The SLC allows the AD RMS cluster to issue:

- SLCs to other servers in the cluster.
- Rights Account Certificates to clients.
- Client licensor certificates.
- Publishing licenses.
- Use licenses.
- Rights policy templates.

The SLC public key encrypts the content key in a publishing license. This allows the AD RMS server to extract the content key and issue end-user licenses against the publishing key.

AD RMS Machine Certificate

The AD RMS machine certificate is used to identify a trusted computer or device. This certificate identifies the client computer's lockbox. The machine certificate public key encrypts the Rights Account Certificate private key. The machine certificate private key decrypts the Rights Account Certificates.

Rights Account Certificate

The Rights Account Certificate (RAC) identifies a specific user. The default validity time for a RAC is 365 days. RACs can only be issued to users in AD DS whose user accounts have email addresses that are associated with them. A RAC is issued the first time a user attempts to access AD RMS-protected content. You can adjust the default validity time using the Rights Account Certificate Policies node of the Active Directory Rights Management Services console.

A temporary RAC has a validity time of 15 minutes. Temporary RACs are issued when a user is accessing AD RMS-protected content from a computer that is not a member of the same or trusted forest as the AD RMS cluster. You can adjust the default validity time using the Rights Account Certificate Policies node of the Active Directory Rights Management Services console.

AD RMS supports the following additional RACs:

- Active Directory® Federation Services (AD FS) RACs are issued to federated users. They have a validity of seven days.
- Two types of Windows Live® ID RACs are supported. Windows Live ID RACs used on private computers have a validity of six months; Windows Live ID RACs used on public computers are valid until the user logs off.

Client Licensor Certificate

A client licensor certificate allows a user to publish AD RMS-protected content when the client computer is not connected to the same network as the AD RMS cluster. The client licensor certificate public key encrypts the symmetric content key and includes it in the publishing license that it issues. The client licensor certificate private key signs any publishing licenses that are issued when the client is not connected to the AD RMS cluster.

Client licensor certificates are tied to a specific user's RAC. If another user who has not been issued a RAC attempts to publish AD RMS-protected content from the same client, the user will be unable to do so until the client is connected to the AD RMS cluster and can issue that user with a RAC.

Publishing License

A publishing license (PL) determines the rights that apply to AD RMS-protected content. For example, the publishing license determines if the user can edit, print, or save a document. The publishing license contains the content key, which is encrypted using the public key of the licensing service. It also contains the URL and the digital signature of the AD RMS server.

End-User License

An end-user license is required to consume AD RMS-protected content. The AD RMS server issues one end-user license per user per document. End-user licenses are cached by default.

How AD RMS Works

AD RMS works in the following manner:

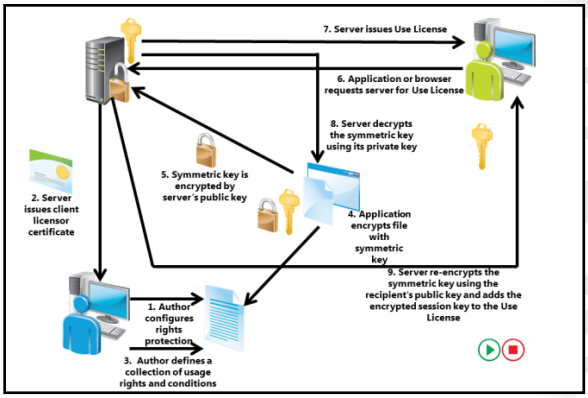

1. The first time the author of the document configures rights protection for the document, a client licensor certificate will be requested from the AD RMS server.
2. The server then issues the client licensor certificate to the client.
3. When the author receives the certificate from the AD RMS server, he or she can configure usage rights on the document.
4. When the author configures usage rights, the application encrypts the file with a symmetric key.
5. This symmetric key is encrypted to the public key of the AD RMS server that is used by the author.
6. The recipient of the file opens it using an AD RMS application or browser. It is not possible to open AD RMS-protected content unless the application or browser supports AD RMS. If the recipient does not have an account certificate on the current device, one will be issued to the user at this point. The application or browser transmits a request to the author's AD RMS server for a Use License.
7. The AD RMS server determines if the recipient is authorized. If the recipient is authorized, the AD RMS server issues a Use License.
8. The AD RMS server decrypts the symmetric key that was encrypted in step 3, using its private key.
9. The AD RMS server re-encrypts the symmetric key using the recipient's public key, and adds the encrypted session key to the Use License.

Lesson 2
Deploying and Managing an AD RMS Infrastructure

Before you deploy AD RMS, it is important to have a deployment plan that is appropriate for your organization's environment. AD RMS deployment in a single-domain forest is different from AD RMS deployment in scenarios where you need to support the publication and consumption of content across multiple forests, to trusted partner organizations, or across the public Internet. Before you deploy AD RMS, you also need to have an understanding of the client requirements, and an appropriate strategy for backing up and recovering AD RMS.

This lesson provides an overview of AD RMS deployment, and the steps you need to take to back up, recover, and decommission an AD RMS infrastructure.

Lesson Objectives

After completing this lesson you will be able to:

- Describe AD RMS deployment scenarios.
- Configure the AD RMS cluster.
- Explain how to install the first server of an AD RMS cluster.
- Describe AD RMS client requirements.
- Explain how to implement an AD RMS backup and recovery strategy.
- Explain how to decommission and remove AD RMS.

AD RMS Deployment Scenarios

An AD RMS deployment consists of one or more servers known as a *cluster*. An AD RMS cluster is not a high-availability failover cluster. When you are deploying AD RMS, you should host the server so that it is highly available. AD RMS is commonly deployed as a highly available virtual machine.

- Deployment scenarios for AD RMS are:
  - AD RMS in a single forest
  - AD RMS in multiple forests
  - AD RMS used on an extranet
  - AD RMS integrated with AD FS

When you deploy AD RMS in a single forest, you have a single AD RMS cluster. This is the most common form of AD RMS deployment. You add servers to the AD RMS cluster as needed, to provide additional capacity.

When you deploy AD RMS across multiple forests, each forest must have its own AD RMS root cluster. It is necessary to configure AD RMS Trusted Publishing Domains to ensure that AD RMS content can be protected and consumed across the multiple forests.

You can also deploy AD RMS to extranet locations. In this deployment, the AD RMS licensing server is accessible to hosts on the Internet. You use this type of deployment to support collaboration with external users.

You can deploy AD RMS with AD FS or the Microsoft Federation Gateway. In this scenario, users leverage federated identity to publish and consume rights-protected content.

As a best practice, you should not deploy AD RMS on a domain controller. You can only deploy AD RMS on a domain controller if the service account is a member of the Domain Admins group.

Configuring the AD RMS Cluster

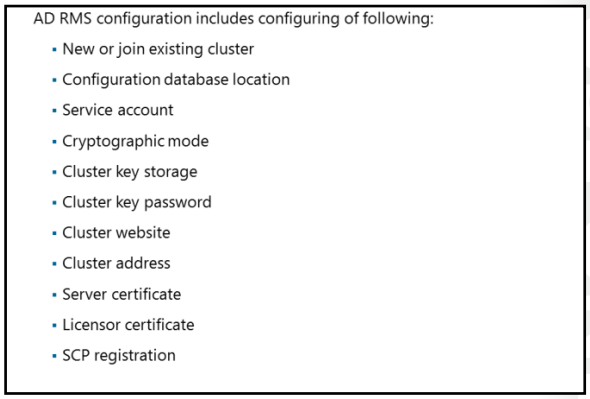

Once you have deployed the AD RMS server role, you need to configure the AD RMS cluster before you can use AD RMS. To configure the AD RMS cluster, you must configure the following components:

1. AD RMS cluster. Choose whether to create a new AD RMS root cluster, or join an existing cluster.
2. Configuration database. Select whether to use an existing SQL Server instance to store the AD RMS configuration database, or to configure and install the Windows Internal Database locally. You can use SQL Server 2008, SQL Server 2008 R2, or SQL Server 2012 to support an AD RMS deployment in Windows Server 2012. As a best practice, use a SQL Server database that is hosted on a separate server.
3. Service account. We recommend that you use a standard domain user account with additional permissions. You can use a managed service account as the AD RMS service account.
4. Cryptographic mode. Choose the strength of the cryptography used with AD RMS.
   - Cryptographic Mode 2 uses RSA 2048-bit keys and SHA-256 hashes.
   - Cryptographic Mode 1 uses RSA 1045-bit keys and SHA-1 hashes.
5. Cluster key storage. Choose where the cluster key is stored. You can either store it within AD RMS, or use a special cryptographic service provider (CSP). If you choose to use a CSP and you want to add additional servers, you need to distribute the key manually.
6. Cluster key password. This password encrypts the cluster key, and is required if you want to join other AD RMS servers to the cluster, or if you want to restore the cluster from backup.
7. Cluster website. Choose the website on the local server that will host the AD RMS cluster website.
8. Cluster address. Specify the fully qualified domain name (FQDN) to be used with the cluster. You have the option of choosing between a Secure Sockets Layer (SSL)-encrypted and non-SSL-encrypted website. If you choose non-SSL-encrypted, you will be unable to add support for identity federation. Once you set the cluster address and port, you cannot change them without completely removing AD RMS.
9. Licensor certificate. Choose the friendly name that the SLC will use. It should represent the function of the certificate.
10. Service connection point registration. Choose whether the service connection point is registered in AD DS when the AD RMS cluster is created. The service connection point allows computers that are members of the domain to locate the AD RMS cluster automatically. Only users that are members of the Enterprise Admins group are able to register the service connection point. You can perform this step after the AD RMS cluster is created; you do not have to perform it during the configuration process. Failure to do so, however, could result in an Error ID of 189 or 190, failure to delete or create the SCP registration, when a client attempts to find the SCP.

Demonstration: Installing the First Server of an AD RMS Cluster

In this demonstration, you will see how to deploy AD RMS on a computer that is running Windows Server 2012.

Demonstration Steps

Configure Service Account

1. Use the Active Directory Administrative Center to create an organizational unit (OU) named Service Accounts in the adatum.com domain.
2. Create a new user account in the Service Accounts OU with the following properties:
    - First name: **ADRMSSVC**
    - User UPN logon: **ADRMSSVC**
    - Password: **Pa$$w0rd**
    - Confirm Password: **Pa$$w0rd**
    - Password never expires: **Enabled**
    - User cannot change password: **Enabled**

Prepare DNS

- Use the DNS Manager console to create a host (A) resource record in the adatum.com zone with the following properties:
    - Name: **adrms**
    - IP Address: **172.16.0.21**

Install the AD RMS role

1. Sign in to LON-SVR1 with the **Adatum\Administrator** account using the password **Pa$$w0rd**.
2. Use the Add Roles and Features Wizard to add the AD RMS role to LON-SVR1 using the following option:
    - Role services: **Active Directory Rights Management Server Services**

Configure AD RMS

1. In Server Manager, from the **AD RMS** node, click **More** to start post deployment configuration of AD RMS.
2. In the AD RMS Configuration Wizard, provide the following information:
    - **Create a new AD RMS root cluster**
    - **Use Windows Internal Database on this server**
    - **Use Adatum\ADRMSSVC as the service account**
    - Cryptographic Mode: **Cryptographic Mode 2**
    - Cluster Key Storage: **Use AD RMS centrally managed key storage**
    - Cluster Key Password: **Pa$$w0rd**
    - Cluster Web Site: **Default Web Site**
    - Connection Type: **Use an unencrypted connection**
    - Fully Qualified Domain Name: **http://adrms.adatum.com**

- Port: **80**
- Licensor Certificate: **Adatum AD RMS**
- Register AD RMS Service Connection Point: **Register the SCP Now**

3. Sign out of LON-SVR1.

Note: You must sign out before you can manage AD RMS.

AD RMS Client Requirements

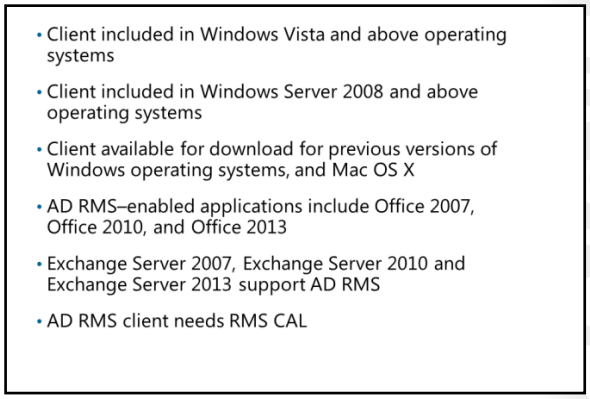

AD RMS content can only be published and consumed by computers that are running the AD RMS client. All versions of Windows Vista and newer client operating systems include AD RMS client software. Windows Server 2008 and newer operating systems also include the AD RMS client. These operating systems do not require additional configuration to consume and publish AD RMS-protected content.

AD RMS client software is available for download to computers that are running the Microsoft Windows XP operating system and Mac OS X. This client software must be installed before users of these operating systems are able to consume and publish AD RMS-protected content.

AD RMS requires compatible applications. Server applications that support AD RMS include the following:

- Microsoft Exchange Server 2007
- Exchange Server 2010
- Exchange Server 2013
- Microsoft Office SharePoint Server 2007
- SharePoint Server 2010
- SharePoint Server 2013

Client applications, such as those included in Microsoft Office 2003, Office 2007, Office 2010, and Office 2013, can publish and consume AD RMS-protected content. You can use the AD RMS SDK to create applications that can publish and consume AD RMS-protected content. XPS Viewer and Windows Internet Explorer® are also able to view AD RMS-protected content.

Note: Microsoft has released the new version of the AD RMS Client Software—AD RMS Client 2.0. You can download it from the Microsoft Download Center. Among other things, the new version provides a new SDK that you can also download from Microsoft Download Center. The new AD RMS SDK provides a simple mechanism for developers to create applications and solutions that protect and consume critical content. With the new SDK it is now possible to rights-enable applications and solutions much faster and easier than before.

AD RMS Client Licensing

To use Rights Management Services in your AD DS environment, you must have Windows Rights Management Client Access Licenses (CALs). These CALs are different from classic Windows Server CALs that you need to connect the client to the server. Each user that will be creating or using right-protected files will need to have an RMS User CAL. Alternatively, you can also use RMS Device CALs for computers that will be used for creating and viewing RMS-protected content.

If you need to share your RMS-protected content outside your organization, you will have to acquire an RMS External Connector License. This license gives organizations the right to permit an unlimited number of external users to access or use a single, licensed copy of the RMS server software without the need to acquire CALs for each external user.

Implementing an AD RMS Backup and Recovery Strategy

To prevent data loss, you must ensure that the AD RMS server is backed up in such a way that it can be recovered in the event of file corruption or server failure. If the AD RMS server becomes inaccessible, all AD RMS-protected content also becomes inaccessible.

- Back up private key and certificates
- Ensure that the AD RMS database is backed up regularly
- Export templates to back them up
- Run AD RMS server as a virtual machine, and perform full server backup

A simple strategy for implementing AD RMS backup and recovery is to run AD RMS server as a virtual machine, and then use an enterprise backup product such as Microsoft System Center 2012 Data Protection Manager to perform regular virtual machine backups. Some of the important components that require backups are the private key, certificates, the AD RMS database, and templates. You can also perform a full-server backup, by running AD RMS server on a virtual machine.

As a best practice, you need to back up the AD RMS private key and all certificates used by AD RMS. The simplest method of doing this is to export the certificates to a safe location. You must also back up the AD RMS database on a regular basis. The method you use to do this depends on whether AD RMS uses SQL Server or the Windows Internal Database. To back up templates, configure the templates to be exported to a shared folder, and then back up these templates.

When you perform recovery of the AD RMS role, it may be necessary to delete the ServiceConnectionPoint object from AD DS. You need to do this if you are recovering an AD RMS root configuration server, and the server attempts to provision itself as a licensing-only server.

Decommissioning and Removing AD RMS

Before you remove an AD RMS server, you should decommission that server. Decommissioning AD RMS puts the cluster into a state where consumers of AD RMS-protected content are able to obtain special keys that decrypts that content, irrespective of the existing restrictions that were placed on the use of that content. If you do not have a decommissioning period, and if you simply remove the AD RMS server, then the AD RMS–protected content will become inaccessible.

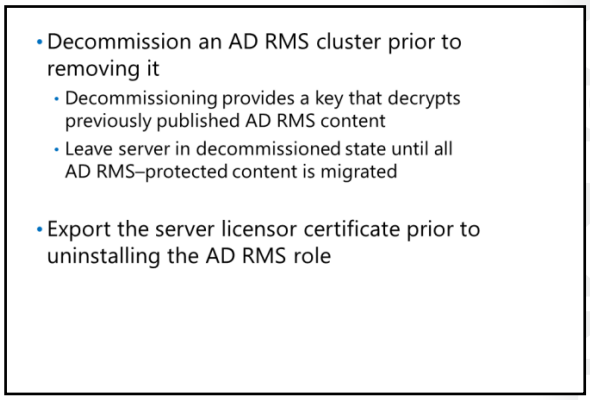

To decommission AD RMS, perform the following steps:

1. Sign in to the server that is hosting AD RMS, and that you wish to decommission.
2. Modify the access control list (ACL) of the file decommissioning.asmx. Grant the **Everyone** group **Read & Execute** permission on the file. This file is stored in the %systemdrive%\inetpub\wwwroot_wmcs\decomission folder.
3. In the Active Directory Rights Management Services console, expand the **Security Policies** node, and then click the **Decommissioning** node.
4. In the Actions pane, click **Enable Decommissioning**.
5. Click **Decommission**.
6. When prompted to confirm that you want to decommission the server, click **Yes**.

After the AD RMS decommissioning process is complete, you should export the server licensor certificate before you uninstall the AD RMS role.

Lesson 3
Configuring AD RMS Content Protection

AD RMS uses rights policy templates to enforce a consistent set of policies to protect content. When you configure AD RMS, you need to develop strategies to ensure that users can still access protected content from a computer that is not connected to the AD RMS cluster. You also need to develop strategies for excluding some users from being able to access AD RMS-protected content, and strategies to ensure that protected content can be recovered in the event that it has expired, the template has been deleted, or the author of the content is no longer available.

Lesson Objectives

After completing this lesson you will be able to:

- Describe the function of rights policy templates.
- Explain how to create a rights policy template.
- Explain how to implement strategies to ensure rights policy templates are available for offline use.
- Describe exclusion policies.
- Explain how to create an exclusion policy to exclude an application.
- Implement an AD RMS Super Users group.

What Are Rights-Policy Templates?

Rights-policy templates allow you to configure standard methods of implementing AD RMS policies across the organization. For example, you can configure standard templates that grant view-only rights, or that block the ability to edit, save, and print. If used with Exchange Server, you can configure templates to block the ability to forward or reply to messages.

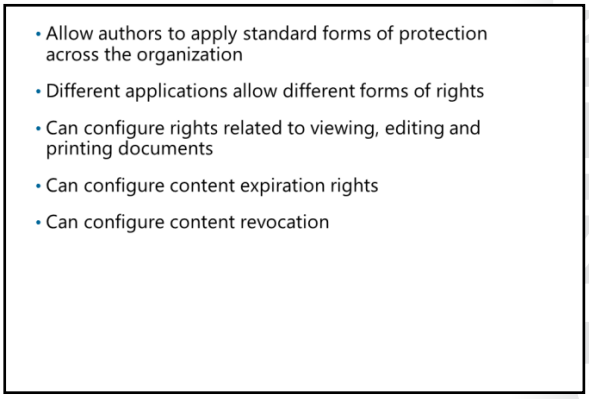

You create rights-policy templates by using the Active Directory Rights Management Services console. The templates are stored in the AD RMS database, and can also be stored in XML format. When content is consumed, the client checks with AD RMS to verify that it has the most recent version of the template.

A document author can choose to protect content by applying an existing template. This is done using an AD RMS-aware application. For example, in Office Word, you apply a template by using the Protect Document function. When you do this, Office Word queries AD DS to determine the location of the AD RMS server. Once the location of the AD RMS server is acquired, templates that are available to the content author can be used.

AD RMS templates support the following rights:

- **Full Control**. Gives a user full control over an AD RMS-protected document.
- **View**. Gives a user the ability to view an AD RMS-protected document.
- **Edit**. Allows a user to modify an AD RMS-protected document.

- **Save**. Allows a user to use the Save function with an AD RMS-protected document.
- **Export (Save as)**. Allows a user to use the Save As function with an AD RMS-protected document.
- **Print**. Allows an AD RMS-protected document to be printed.
- **Forward**. Used with Exchange Server. Allows the recipient of an AD RMS-protected message to forward that message.
- **Reply**. Used with Exchange Server. Allows the recipient of an AD RMS-protected message to reply to that message.
- **Reply All**. Used with Exchange Server. Allows the recipient of an AD RMS-protected message to use the Reply All function to reply to that message.
- **Extract**. Allows the user to copy data from the file. If this right is not granted, the user cannot copy data from the file.
- **Allow Macros**. Allows the user to utilize macros.
- **View Rights**. Allows the user to view assigned rights.
- **Edit Rights**. Allows the user to modify the assigned rights.

Rights can only be granted, and cannot be explicitly denied. For example, to ensure that a user cannot print a document, the template associated with the document must not include the Print right.

Administrators are also able to create custom rights that can be used with custom AD RMS–aware applications.

AD RMS templates can also be used to configure documents with the following properties:

- **Content Expiration**. Determines when the content expires. The options are:
  - **Never**. The content never expires.
  - **Expires on a particular date**. Content expires at a particular date and time.
  - **Expires after**. The content expires a particular number of days after it is created.
- **Use license expiration**. Determines the time interval in which the use license will expire, and a new one will need to be acquired.
- **Enable users to view protected content using a browser add-on**. Allows content to be viewed using a browser add-on. Does not require the user have an AD RMS-aware application.
- **Require a new use license each time content is consumed**. When you enable this option, client-side caching is disabled. This means that the document cannot be consumed when the computer is offline.
- **Revocation policies**. Allows the use of a revocation list. This allows an author to revoke permission to consume content. You can specify how often the revocation list is checked, with the default being once every 24 hours.

Once an AD RMS policy template is applied to a document, any updates to that template will also be applied to that document. For example, if you have a template without a content expiration policy that is used to protect documents, and you modify that template to include a content expiration policy, those protected documents will now have an expiration policy. Template changes are reflected when the end-user license is acquired. If end-user licenses are configured not to expire, and the user who is accessing a document already has a license, then the user may not receive the updated template.

Note: You should avoid deleting templates, because documents that use those templates will become inaccessible to everyone except for members of the Super Users group. As a best practice, archive templates instead of deleting them.

You can view the rights associated with a template by selecting the template within the Active Directory Rights Management Services console, and then in the Actions menu, clicking **View Rights Summary**.

Demonstration: Creating a Rights-Policy Template

In this demonstration, you will see how to create a rights policy template that allows users to view a document, but not to perform other actions.

Demonstration Steps

- In the Active Directory Rights Management Services console, use the Rights Policy Template node to create a Distributed Rights Policy Template with the following properties:
- Language: **English (United States)**
- Name: **ReadOnly**
- Description: **Read-only access. No copy or print**
- Users and rights: **executives@adatum.com**
- Rights for Anyone: **View**
- **Grant owner (author) full control right with no expiration**
- Content Expiration: Expires after **7** days
- Use license expiration: Expires after **7** days
- Require a new use license every time content is consumed (disable client-side caching): **Enabled**

Providing Rights-Policy Templates for Offline Use

If the users will publish AD RMS-connected templates when they are not connected to the network, you need to ensure that they have access to a local copy of the available rights-policy templates.

You can configure computers to acquire and store published rights-policy templates automatically, so that they are available offline. To enable this feature, computers must be running one of the following Windows operating systems:

- Windows Vista with Service Pack 1 (SP1) or newer
- Windows 7
- Windows 8
- Windows 8.1

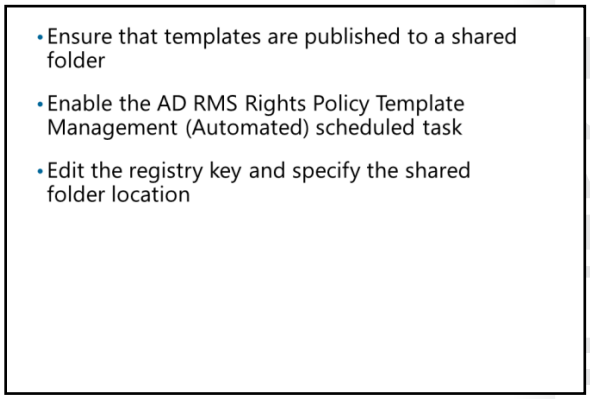

- Windows Server 2008
- Windows Server 2008 R2Windows Server 2012
- Windows Server 2012 R2

To enable this functionality, in the Task Scheduler, enable the AD RMS Rights Policy Template Management (Automated) Scheduled Task, and then edit the following registry key:

```
HKEY_CURRENT_USER\Software\Microsoft\Office\12.0\Common\DRM
```

Provide the following location for templates to be stored:

%LocalAppData%\Microsoft\DRM\Templates

When computers that are running these operating systems are connected to the domain, the AD RMS client polls the AD RMS cluster for new templates, or updates to existing templates.

As an alternative for templates distribution, you can also use shared folders to store templates. You can configure a shared folder for templates by performing the following steps:

1. In the Active Directory Rights Management Services console, right-click the **Rights Policy Templates** node, and then click **Properties**.
2. On the **Rights Policy Templates Properties** dialog box, specify the location of the shared folder to which templates will be published.

What Are Exclusion Policies?

Exclusion policies allow you to prevent specific user accounts, client software, or applications from using AD RMS.

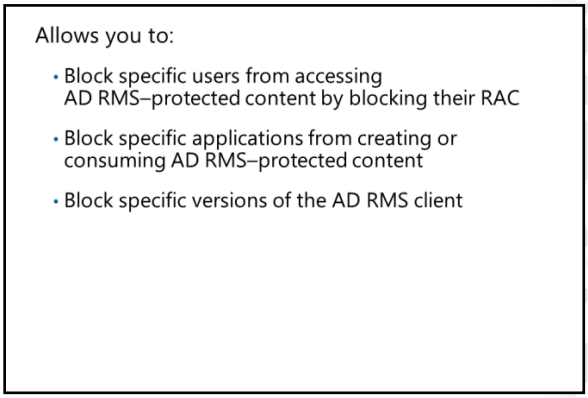

User Exclusion

The User Exclusion policy allows you to configure AD RMS so that specific user accounts, which are identified based on email addresses, are unable to obtain Use Licenses. You do this by adding each user's RAC to the exclusion list. User Exclusion is disabled by default. Once you have enabled User Exclusion, you can exclude specific RACs.

You can use User Exclusion in the event that you need to lock a specific user out of AD RMS-protected content. For example, when users leave the organization, you might exclude their RACs to ensure that they are unable to access protected content. You can block the RACs that are assigned to both internal users and external users.

Application Exclusion

Application Exclusion allows you to block specific applications—such as Office PowerPoint®—from creating or consuming AD RMS-protected content. You specify applications based on executable names. You also specify a minimum and a maximum version of the application. Application Exclusion is disabled by default.

Note: It is possible to circumvent Application Exclusion by renaming an executable file.

Lockbox Exclusion

Lockbox Exclusion allows you to exclude AD RMS clients, such as those used with specific operating systems such as Windows XP and Windows Vista. Lockbox Version Exclusion is disabled by default. Once you have enabled Lockbox Version Exclusion, you must specify the minimum lockbox version that can be used with the AD RMS cluster.

Additional Reading: To find out more about enabling exclusion policies, see Enabling Exclusion Policies at http://go.microsoft.com/fwlink/?LinkId=270031.

Demonstration: Creating an Exclusion Policy to Exclude an Application

In this demonstration, you will see how to exclude the Office PowerPoint application from AD RMS.

Demonstration Steps

1. In the Active Directory Rights Management Services console, enable Application exclusion.
2. In the **Exclude Application** dialog box, enter the following information:
   - Application File name: **Powerpnt.exe**
   - Minimum version: **14.0.0.0**
   - Maximum version: **16.0.0.0**

AD RMS Super Users Group

The AD RMS Super Users Group is a special role, and members of this group have full control over all rights-protected content managed by the cluster. Super Users Group members are granted full owner rights in all use licenses that are issued by the AD RMS cluster on which the Super Users group is configured. This means that members of this group can decrypt any rights-protected content file and remove rights protection from it.

- Super users group members are granted full owner rights in all use licenses that are issued by the AD RMS cluster on which the super users group is configured.
- Super users group:
  - Is not configured by default
  - Can be used as data recovery mechanism for AD RMS–protected content
    - Can recover content that has expired
    - Can recover content if the template is deleted
    - Can recover content without requiring author credentials
  - Must be an Active Directory group with an assigned email address.

The AD RMS super users group provides a data recovery mechanism for AD RMS-protected content. This mechanism is useful in the event that you need to recover AD RMS-protected data, such as when content has expired, when a template has been deleted, or when you do not have access.

Members of the Super Users group are assigned owner use licenses for all content that is protected by the AD RMS cluster on which that particular Super Users group is enabled. Members of the Super Users group are able to reset the AD RMS server's private key password.

As members of the Super Users group can access any AD RMS-protected content, you must be especially careful when you are managing the membership of this group. If you choose to use the AD RMS Super Users group, you should consider implementing restricted-groups policy and auditing to limit group membership, and also record any changes that are made. Super User activity is written to the Application event log.

The Super Users group is disabled by default. You enable the super users group by performing the following steps:

1. In the Active Directory Rights Management Services console, expand the server node, and then click **Security Policies**.
2. In the Security Policies area, under **Super Use**rs, click **Change Super User Settings**.
3. In the Actions pane, click **Enable Super Users**.

To set a particular group as the Super Users group:

1. In the Security Policies\Super Users Super Users area, click **Change super user group**.
2. Provide the email address associated with the Super Users group.

AD RMS Integration with Dynamic Access Control

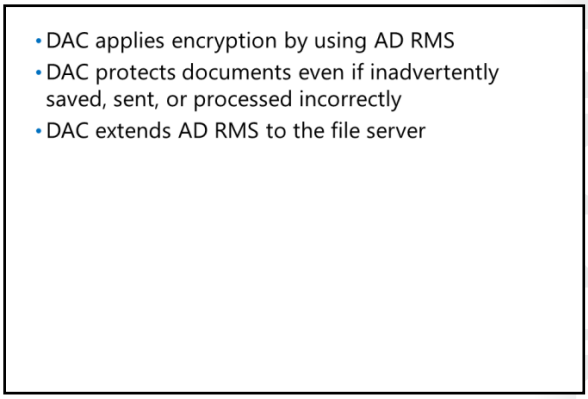

Mitigating organizational risk ultimately is an overarching goal for all IT departments. When organizations deal with sensitive information, and its sharing and dissemination via email and the Internet, that data's security is a foremost concern. Another chief concern is meeting the various compliance regulations, such as the Health Insurance Portability and Accountability Act (HIPAA) or Payment Card Industry Data Security Standard (PCI-DSS). These regulations dictate how information must be encrypted. Additionally, there are numerous business reasons to encrypt sensitive business information. However, encrypting information is expensive, and it can affect business productivity. Accordingly, organizations tend to use different approaches and priorities to encrypt their information.

You can use the data classification feature of Dynamic Access Control (DAC) in Windows Server 2012 to automatically apply RMS templates to sensitive data. You do this by applying RMS template protection for sensitive documents a few seconds after DAC classification identifies the file as sensitive on the file server, and you can enable this by running continuous file-management tasks on the file server.

AD RMS encryption provides an additional layer of protection for files. Even if a person with access to a sensitive file inadvertently sends that file through email, the AD RMS encryption is retained with the file. Users who want to access the file must first authenticate themselves to an AD RMS server to receive the decryption key.

Support for non-Microsoft file formats may be available through independent software vendors (ISVs) using the AD RMS software development Kit (SDK) or who are writing code based on Microsoft attributes, which applies to DAC. After an author protects a file by AD RMS encryption, or file protection occurs automatically via DAC deployment, data-management features, such as search functionality or content-based classification, are no longer available for that file.

AD RMS enables both individuals and administrators, through Information rights management (IRM) policies, to specify access permissions to documents, workbooks, and presentations. This helps prevent sensitive information from being printed, forwarded, or copied by unauthorized people. Once you or DAC restrict a file's permission with IRM, AD RMS enforces access and usage restrictions no matter where the information is, because the file's permission is stored in the document file itself.

File and Storage Services DAC lets you set up and manage file servers that provide central locations on your network, where you can store files and share them with users. File-server administrators can configure file-management tasks that invoke AD RMS protection for sensitive documents a few seconds after AD RMS identifies the file as a sensitive file on the file server. This provides continuous file management tasks on the file server.

DAC enables you to protect sensitive information automatically by using AD RMS. For example, you could apply AD RMS if a file had the word *confidential* in it, as follows:

1. A rule is created to automatically apply RMS protection to any file that contains the word confidential.
2. A user creates a file with the word confidential in the text, and then saves it.
3. The AD RMS DAC classification engine, following rules set in the central access policy, discovers the document with the word confidential, and then initiates AD RMS protection accordingly.
4. AD RMS applies a template and encryption to the document on the file server, and then encrypts and classifies it.

Lesson 4
Configuring External Access to AD RMS

You may often find it necessary to provide users who are not a part of the organization with access to AD RMS-protected content. This could be a situation in which an external user is a contractor who requires access to sensitive materials, or a partner organization where your users will require access to protected content published by their AD RMS server. AD RMS provides many different options for granting external users access to protected content.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the options available for making AD RMS-protected content accessible to external users.
- Explain how to implement Trusted User Domains.
- List the steps necessary to deploy Trusted Publishing Domains.
- Describe the steps necessary to configure AD RMS to share protected content with users who have Windows Live IDs.
- Determine the appropriate solution for sharing AD RMS-protected content with external users.

Options for Enabling External Users with AD RMS Access

Trust policies allow users who are external to your organization the ability to consume AD RMS-protected content. For example, a trust policy can allow users in Bring Your Own Device (BYOD) environments to consume AD RMS-protected content, even though those computers are not members of the organization's AD DS domain. AD RMS trusts are disabled by default, and you must enable them before you can use them. AD RMS supports the following trust policies.

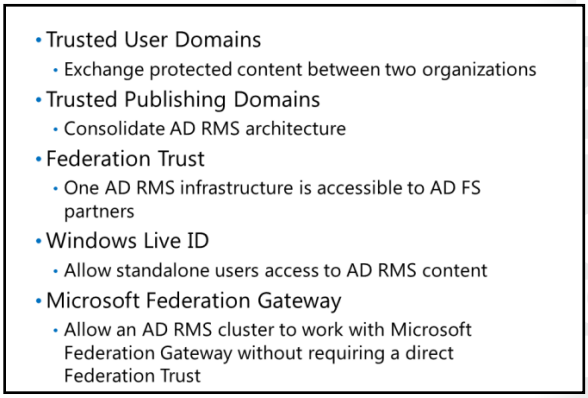

Trusted User Domains

Trusted User Domains (TUDs) allow an AD RMS cluster to process requests for client licensor certificates, or use licenses from people who have RACs issued by a different AD RMS cluster. For example, A. Datum Corporation and Trey Research are separate organizations that have each deployed AD RMS. TUDs allow each organization to publish and consume AD RMS-protected content to and from the partner organization without having to implement AD DS trusts or AD FS.

Trusted Publishing Domains

Trusted Publishing Domains (TPDs) allow one AD RMS cluster to issue end-user licenses to content that uses publishing licenses that are issued by a different AD RMS cluster. TPDs consolidate existing AD RMS infrastructure.

Federation Trust

Federation Trust provides single sign-on (SSO) for partner technologies. Federated partners can consume AD RMS-protected content without deploying their own AD RMS infrastructure. Federation Trust requires AD FS deployment.

Windows Live ID Trust

You can use Windows Live ID to allow standalone users with Windows Live ID accounts to consume AD RMS-protected content generated by users in your organization. However, Windows Live ID users are unable to create content that is protected by the AD RMS cluster.

Microsoft Federation Gateway

Microsoft Federation Gateway allows an AD RMS cluster to process requests to publish and consume AD RMS-protected content from external organizations, by accepting claims-based authentication tokens from the Microsoft Federation Gateway. Rather than configuring a federation trust, each organization has a relationship with the Microsoft Federation Gateway. The gateway acts as a trusted broker.

Additional Reading: To learn more about AD RMS Trust Policies, see http://go.microsoft.com/fwlink/?LinkId=270032.

Implementing Trusted User Domain

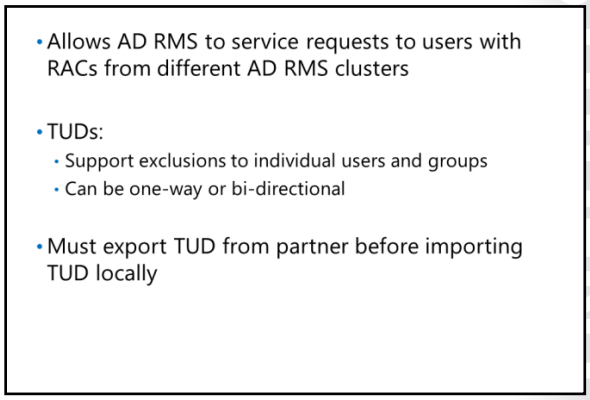

Trusted User Domain (TUD) allows AD RMS to service requests from users who have RACs issued by different AD RMS deployments. You can use exclusions with each TUD to block access to specific users and groups.

To configure AD RMS to support service requests from users who have RACs issued by different AD RMS deployments, you add the organization to the list of TUDs. TUDS can be one-way, where organization A is a TUD of organization B, or bi-directional, where organization A and organization B are TUDs of each other. In one-way deployments, it is possible for the users of the TUD to consume the content of the local AD RMS deployment, but they cannot publish AD RMS-protected content by using the local AD RMS cluster.

You need to enable anonymous access to the AD RMS licensing service in Internet Information Services (IIS) when you use TUD. This is necessary because, by default, accessing the service requires authenticating using Integrated Windows Authentication.

To add a TUD, perform the following steps:

1. The TUD of the AD RMS deployment that you want to trust must have already been exported, and the file must be available. (TUD files use the .bin extension.)
2. In the AD RMS console, expand **Trust Policies**, and then click **Trusted User Domains**.
3. In the Actions pane, click **Import Trusted User Domain**.
4. In the **Trusted User Domain** dialog box, enter the path to the exported TUD file with the .bin extension.
5. Provide a name to identify this TUD. If you have configured federation, you can also choose to extend the trust to federated users of the imported server.

You can also use the Windows PowerShell cmdlet **Import-RmsTUD**, which is part of the ADRMSADMIN Windows PowerShell® module, to add a TUD.

To export a TUD, perform the following steps:

1. In the Active Directory Rights Management Services console, expand **Trust Policies**, and then click **Trusted User Domains**.
2. In the Actions pane, click **Export Trusted User Domain**.
3. Save the TUD file with a descriptive name.

You can also use the Windows PowerShell cmdlet **Export-RmsTUD** to export an AD RMS server TUD.

Implementing TPD

You can use Trusted Publisher Domain (TPD) to set up a trust relationship between two AD RMS deployments. An AD RMS TPD, which is a local AD RMS deployment, can grant end-user licenses for content published using the Trusted Publishing domain's AD RMS deployment. For example, Contoso, Ltd. and A. Datum Corporation are set up as TPD partners. TPD allows users of the Contoso AD RMS deployment to consume content published using the A. Datum AD RMS deployment, by using end-user licenses that are granted by the Contoso AD RMS deployment.

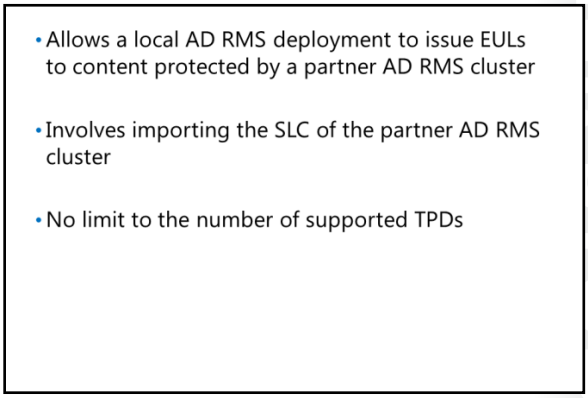

You can remove a TPD at any time. When you do this, clients of the remote AD RMS deployment will not be able to issue end-user licenses to access content protected by your AD RMS cluster.

When you configure a TPD, you import the SLC of another AD RMS cluster. TPDs are stored in .xml format, and are protected by passwords.

To export a TPD, perform the following steps:

1. In the Active Directory Rights Management Services console, expand **Trust Policies**, and then click **Trusted Publishing Domains**.
2. In the **Results** pane, click the certificate for the AD RMS domain that you want to export, and then in the Actions pane, click **Export Trusted Publishing Domain**.
3. Choose a strong password and a filename for the TPD.

When you export a TPD, it is possible to save it as a V1-compatible TPD file. This allows the TPD to be imported into organizations that are using AD RMS clusters on earlier versions of the Windows Server operating system, such as the version available in Windows Server 2003. You can also use the Windows PowerShell cmdlet **Export-RmsTPD** to export a TPD.

To import a TPD, perform the following steps:

1. In the Active Directory Rights Management Services console, expand **Trust Policies**, and then click **Trusted Publishing Domains**.
2. In the Actions pane, click **Import Trusted Publishing Domain**.
3. Specify the path of the Trusted Publishing Domain file that you want to import.
4. Enter the password to open the Trusted Publishing Domain file, and enter a display name that identifies the TPD.

Alternatively, you can also use the Windows PowerShell cmdlet **Import-RmsTPD** to import a TPD.

Additional Reading: To can learn more about importing TPDs, see Add a Trusted Publishing Domain at http://go.microsoft.com/fwlink/?LinkId=270033.

Sharing AD RMS-Protected Documents by Using Windows Live ID

You can use Windows Live ID as a method of providing RACs to users who are not part of your organization.

- Provide RACs to users who are not part of an organization
- Users with Windows Live ID accounts can consume AD RMS–protected content
- Users with Windows Live ID accounts cannot publish AD RMS–protected content

To trust Windows Live ID-based RACs, perform the following steps:

1. In the Active Directory Rights Management Services console, expand **Trust Policies**, and then click **Trusted User Domains**.
2. In the Actions pane, click **Trust Windows Live ID**.

To exclude specific Windows Live ID email domains, right-click the Windows Live ID certificate, click **Properties**, and then click the **Excluded Windows Live IDs** tab. You can then enter the Windows Live ID accounts that you want to exclude from being able to procure RACs.

To allow users with Windows Live ID accounts to obtain RACs from your AD RMS cluster, you need to configure IIS to support anonymous access. To do this, perform the following steps:

1. On the AD RMS server, open the IIS Manager console.
2. Navigate to the **Sites\Default Web Site_wmcs** node, right-click the **Licensing** virtual directory, and then click **Switch to Content View**.
3. Right-click **license.asmx**, and then click **Switch to Content View**.
4. Double-click **Authentication**, and then enable **Anonymous Authentication**.
5. Repeat this step for the file **ServiceLocator.asmx**.

Additional Reading: To can learn more about using Windows Live ID to establish RACs for users, see http://go.microsoft.com/fwlink/?LinkId=270034.

Considerations for Implementing External User Access to AD RMS

The type of external access that you configure depends on the types of external users that need access to your organization's content.

When you are determining which method to use, consider the following questions:

- Does the external user belong to an organization that has an existing AD RMS deployment?
- Does the external user's organization have an existing Federated Trust with the internal organization?

- Use Windows Live ID to issue RACs to users who are not part of organizations, and who need to consume content
- Use TUD for RACs issued by a different AD RMS cluster
- Use TPD to allow local RACs to access remotely published AD RMS content
- Use Federation Trust between organizations that have a federated relationship
- Use Microsoft Federation Gateway when no direct federated relationship exists

- Has the external user's organization established a relationship with the Microsoft Federation Gateway?
- Does the external user need to publish AD RMS-protected content that is accessible to internal RAC holders?

It is possible that organizations may use one solution before deciding to implement another. For example, during initial stages, only a small number of external users may require access to AD RMS-protected content. In this case, using Windows Live ID accounts for RACs may be appropriate. When large numbers of external users from a single organization require access, a different solution may be appropriate. The financial benefit a solution brings to an organization must exceed the cost of implementing that solution.

Windows Azure RMS

Like AD RMS, Windows Azure Active Directory (AD) Rights Management lets you safeguard sensitive information that Office applications and services create or manage, such as email or correspondence that requires confidential treatment. You assign rights to content when it is created, and the content is distributed in an encrypted form of persistent protection wherever it travels. Rights that can be assigned include allowing or deny viewing, printing, and copying messages or documents using template-based assignment.

- Windows Azure AD Rights Management is free
  - Sign up as a free tenant in Windows Azure AD
- Use the Azure viewer app to send a message to an organization with which you wish to employ Rights Management
  - Message will contain simple instructions to obtain tenant status
  - You can then use Rights Management across a B2B partnership
- You could replace your AD RMS infrastructure with Windows Azure AD Rights Management

When you use Azure AD and are connected to AD DS via the Microsoft Directory Synchronization (DirSync) component, the Windows Azure AD Rights Management service benefits from a cross-organization trust. When your organization sets up an Azure AD tenant, and optionally configures DirSync and federation services, Windows Azure AD Rights Management relies on those trusts to enable business-to-business collaboration in a more easily configured and administrated manner.

Before Azure AD, organizations that desired secure collaboration had to use paired federation trusts with every organization with which their organization needed to collaborate. Each business that you wish to share protected documents with must be a Windows Azure AD tenant. There is no cost to become a Windows Azure tenant. You use the Windows Azure viewer app to send someone an email containing a

protected file, and the file provides simple instructions for how to sign up as a new Windows Azure tenant, so that they can start accessing the shared file.

The following table provides a side-by-side comparison of the features and benefits of Windows Azure AD Rights Management and AD RMS.

Windows Azure AD Rights Management	AD RMS
Supports IRM capabilities in Microsoft online services and other online offerings such as Exchange Online and SharePoint Online, and traditional on-premises products, including Microsoft Exchange Server and SharePoint servers.	Works primarily with on-premises Microsoft server products including Microsoft SharePoint Server, Exchange Server and File Classification Infrastructure (FCI).
Enables implicit trust between organizations and users that are current Microsoft® Office 365 subscribers. Content is easily and securely shared between users within the same organization or with valid users in other organizations who have Office 365 tenant accounts.	Explicitly defines trusts in a direct point-to-point relationship between two organizations, by using either trusted user domains (TUDs) or federated trusts created by using AD FS.
Offers a predefined set of rights policy templates. Included are two templates: one delivers read-only viewing of protected content and the other delivers write or modify permissions over the protected content.	Provides the ability to create, define and use your own rights policy templates. Templates for AD RMS can be configured on a more detailed level.
Supports Microsoft Office 2010 and Office 2013 users.	Supports Microsoft Office 2007, Office 2010, and Office 2013 users.
Does not provide support for Windows Vista or Windows XP users. Supports users who are running Windows 7 and Windows 8.	Provides support for users running Windows XP, Windows Vista, Windows 7, and Windows 8.
Supports Cryptographic Mode 2 only.	Supports both Cryptographic Mode 1 and Cryptographic Mode 2.
Supports outbound migration only from Windows Azure AD Rights Management to AD RMS.	Supports migration from Windows Azure AD Rights Management and migration from AD RMS running versions older than Windows Server 2003.

Windows Azure AD Rights Management could replace the entire AD RMS cluster infrastructure. You could do this over an extended period, allowing new content to use the Windows Azure AD Rights Management services while supporting the older content through your existing AD RMS.

For more information on Windows Azure AD Rights Management, go to:
http://go.microsoft.com/fwlink/?LinkID=386641

Lab: Implementing AD RMS

Scenario

Because of the highly confidential nature of the research that is performed at A. Datum Corporation, the security team at A. Datum wants to implement additional security for certain documents that the Research department creates. The security team is concerned that anyone with Read access to the documents can modify and distribute the documents in any way that they choose. The security team would like to provide an extra level of protection that stays with the document even if it is moved around the network or outside the network.

As one of the senior network administrators at A. Datum, you need to plan and implement an AD RMS solution that will provide the level of protection requested by the security team. The AD RMS solution must provide many different options that can be adapted for a wide variety of business and security requirements.

Objectives

In this lab, you will see how to:

- Install and configure AD RMS.
- Configure AD RMS Templates.
- Implement AD RMS Trust Policies.
- Verify AD RMS Deployment.

Lab Setup

Estimated Time: 60 minutes

Virtual machines: 20412C-LON-DC1,

20412C-LON-SVR1,

20412C-LON-CL1,

20412C-TREY-DC1,

20412C-TREY-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 **20412C-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 step 2 for **20412C-LON-SVR1**, **20412C-TREY-DC1**, **20412C-LON-CL1**, and **20412C-TREY-CL1**. Do not sign in until directed to do so.

Exercise 1: Installing and Configuring AD RMS

Scenario

The first step in deploying AD RMS at A. Datum is to deploy a single server in an AD RMS cluster. You will begin by configuring the appropriate DNS records and the AD RMS service account. Then you will install and configure the first AD RMS server. You will also enable the AD RMS Super Users group.

The main tasks for this exercise are as follows:

- Configure Domain Name System (DNS) and the Active Directory® Rights Management Services (AD RMS) service account.
- Install and configure the AD RMS server role.
- Configure the AD RMS Super Users group.

▶ Task 1: Configure Domain Name System (DNS) and the Active Directory® Rights Management Services (AD RMS) service account

1. Sign in to LON-DC1 with the **Adatum\Administrator** account and the password **Pa$$w0rd**.
2. Use the Active Directory Administrative Center to create an OU named **Service Accounts** in the adatum.com domain.
3. Create a new user account in the Service Accounts OU with the following properties:
    - First name: **ADRMSSVC**
    - User UPN logon: **ADRMSSVC**
    - Password: **Pa$$w0rd**
    - Confirm Password: **Pa$$w0rd**
    - Password never expires: **Enabled**
    - User cannot change password: **Enabled**
4. Create a new Global security group in the Users container named **ADRMS_SuperUsers**. Set the email address of this group as **ADRMS_SuperUsers@adatum.com**.
5. Create a new global security group in the Users container named **Executives**. Set the email address of this group as **executives@adatum.com**.
6. Add the user accounts **Aidan Delaney** and **Bill Malone** to the Executives group.
7. Use the DNS Manager console to create a host (**A**) resource record in the adatum.com zone with the following properties:
    - Name: **adrms**
    - IP Address: **172.16.0.21**

▶ Task 2: Install and configure the AD RMS server role

1. Sign in to LON-SVR1 with the **Adatum\Administrator** account and the password **Pa$$w0rd**.
2. Use the Add Roles and Features Wizard to add the Active Directory Rights Management Services role to LON-SVR1 using the following option:
    - Role services: **Active Directory Rights Management Services**
3. From the AD RMS node in the **Server Manager**, click **More** to start post deployment configuration of AD RMS.

4. On the AD RMS Configuration Wizard, provide the following information:
   - **Create a new AD RMS root cluster**
   - **Use Windows Internal Database on this server**
   - Service account: **Adatum\ADRMSSVC**
   - Cryptographic Mode: **Cryptographic Mode 2**
   - Cluster Key Storage: **Use AD RMS centrally managed key storage**
   - Cluster Key Password: **Pa$$w0rd**
   - Cluster Web Site: **Default Web Site**
   - Connection Type: **Use an unencrypted connection**
   - Fully Qualified Domain Name: **http://adrms.adatum.com**
   - Port: **80 (**Note that in production, we would use a encrypted, that is, https connection)
   - Licensor Certificate: **Adatum AD RMS**
   - Register AD RMS Service Connection Point: **Register the SCP Now**
5. Use the Internet Information Services (IIS) Manager console to enable Anonymous Authentication on the **Default Web Site_wmcs** and the **Default Web Site_wmcs\licensing** virtual directories.
6. Sign out of LON-SVR1.

Note: You must sign out before you can manage AD RMS. This lab uses port 80 for convenience. In production environments, you would protect AD RMS using an encrypted connection.

▶ Task 3: Configure the AD RMS Super Users group

1. Sign in to LON-SVR1 with the **Adatum\Administrator** account and the password **Pa$$w0rd**.
2. Open the Active Directory Rights Management Services console.
3. From the Active Directory Rights Management Services console, enable Super Users.
4. Set the **ADRMS_SuperUsers** group as the Super Users group.

Results: After completing this exercise, you will have installed and configured AD RMS.

Exercise 2: Configuring AD RMS Templates

Scenario

After you deploy the AD RMS server, you must configure the rights-policy templates and exclusion policies for the organization. You will then deploy both components.

The main tasks for this exercise are as follows:

- Configure a new rights-policy template.
- Configure the rights-policy template distribution.
- Configure an exclusion policy.

▶ Task 1: Configure a new rights-policy template

- On LON-SVR1, use the Rights Policy Template node of the Active Directory Rights Management Services console to create a Distributed Rights Policy Template with the following properties:
  - Language: **English (United States)**
  - Name: **ReadOnly**
  - Description: **Read only access. No copy or print**
  - Users and rights: **executives@adatum.com**
  - Rights for Anyone: **View**
  - **Grant owner (author) full control right with no expiration**
  - Content Expiration: **7** days
  - Use license expiration: **7** days
  - Require a new use license: **every time content is consumed (disable client-side caching)**

▶ Task 2: Configure the rights-policy template distribution

1. On LON-SVR1, open a Windows PowerShell prompt, issue the following commands, and press Enter after each:

```
New-Item c:\rmstemplates -ItemType Directory
New-SmbShare -Name RMSTEMPLATES -Path c:\rmstemplates -FullAccess ADATUM\ADRMSSVC
New-Item c:\docshare -ItemType Directory
New-SmbShare -Name docshare -Path c:\docshare -FullAccess Everyone
```

2. In the Active Directory Rights Management Services console, set the Rights Policy Templates file location to **\\LON-SVR1\RMSTEMPLATES**.
3. In Windows Explorer, view the c:\rmstemplates folder. Verify that the **ReadOnly.xml** template is present.

▶ Task 3: Configure an exclusion policy

1. In the Active Directory Rights Management Services console, enable Application exclusion.
2. In the **Exclude Application** dialog box, enter the following information:
   - Application File name: **Powerpnt.exe**
   - Minimum version: **14.0.0.0**
   - Maximum version: **16.0.0.0**

Results: After completing this exercise, you will have configured AD RMS templates.

Exercise 3: Implementing the AD RMS Trust Policies

Scenario

As part of the AD RMS deployment, you need to ensure that AD RMS functionality is extended to the Trey Research AD RMS deployment. You will configure the required trust policies, and then validate that you can share protected content between the two organizations.

The main tasks for this exercise are as follows:

- Export the Trusted User Domains policy.
- Export the Trusted Publishing Domains policy.
- Import the Trusted User Domain policy from the partner domain.
- Import the Trusted Publishing Domains policy from the partner domain.

▶ Task 1: Export the Trusted User Domains policy

1. On LON-SVR1, open a Windows PowerShell prompt, issue the following commands, and press Enter after each:

```
New-Item c:\export -ItemType Directory
New-SmbShare -Name Export -Path c:\export -FullAccess Everyone
```

2. Use the Active Directory Rights Management Services console to export the TUD policy to the **\\LON-SVR1\export** share as **ADATUM-TUD.bin**.
3. Sign in to TREY-DC1 with the **TREYRESEARCH\Administrator** account and the password **Pa$$w0rd**.
4. On TREY-DC1, open the Active Directory Rights Management Services console.
5. Export the **Trusted User Domains** policy to the \\LON-SVR1\export share as TREYRESEARCH-TUD.bin.
6. On TREY-DC1, open a Windows PowerShell prompt, issue the following commands, and press Enter after each:

```
Add-DnsServerConditionalForwarderZone -MasterServers 172.16.0.10 -Name adatum.com
```

▶ Task 2: Export the Trusted Publishing Domains policy

1. Switch to LON-SVR1.
2. Use the Active Directory Rights Management Services console to export the TPD policy to the **\\LON-SVR1\export** share as **ADATUM-TPD.xml**. Protect this file by using the password **Pa$$w0rd**.
3. Switch to TREY-DC1.
4. Use the Active Directory Rights Management Services console to export the TPD policy to the **\\LON-SVR1\export** share as **TREYRESEARCH-TPD.xml**. Protect this file by using the password **Pa$$w0rd**.

▶ Task 3: Import the Trusted User Domain policy from the partner domain

1. Switch to LON-SVR1.
2. Import the TUD policy for Trey Research by importing the file **\\LON-SVR1\export\treyresearch-tud.bin**. Use the display name **TreyResearch**.
3. Switch to TREY-DC1.
4. Import the TUD policy for Trey Research by importing the file **\\LON-SVR1\export\adatum-tud.bin**. Use the display name **Adatum**.

▶ Task 4: Import the Trusted Publishing Domains policy from the partner domain

1. Switch to LON-SVR1.
2. Import the Trey Research TPD by importing the file **\\LON-SVR1\export\treyresearch-tpd.xml**, using the password **Pa$$w0rd** and the display name **Trey Research**.

3. Switch to MUN-SVR1.
4. Import the Adatum Trusted Publishing Domain by importing the file **\\LON-SVR1\export\adatum-tpd.xml**, and then using the password **Pa$$w0rd** and the display name **Adatum**.

Results: After completing this exercise, you will have implemented the AD RMS trust policies.

Exercise 4: Verifying the AD RMS Deployment

Scenario

As a final step in the deployment, you will validate that the configuration is working correctly.

The main tasks for this exercise are as follows:

- Create a rights-protected document.
- Verify internal access to protected content.
- Open the rights-protected document as an unauthorized user.
- Open and edit the rights-protected document as an authorized user at Trey Research.
- To prepare for the next module.

▶ Task 1: Create a rights-protected document

1. Sign in to LON-CL1 with the **Adatum\Administrator** account and the password **Pa$$w0rd**.
2. Add Aidan, Bill and Carol as local Remote Desktop Users in the Systems properties.
3. Sign out of LON-CL1.
4. Sign in to LON-CL1 with the **Adatum\Aidan** account and the password **Pa$$w0rd**.
5. Add the http://adrms.adatum.com URL to the Local intranet group in Internet options Security tab using the Advanced button in Sites.

Note: This above step is necessary for the Office program to find the proper AD RMS Cluster URL. It must be in the local intranet sites. It must be done for each user

6. Open **Word 2013**.
7. Create a document named **Executives Only**.
8. In the document, type the following text:

   This document is for executives only, it should not be modified.
9. From the Permissions item, choose to restrict access. Grant **bill@adatum.com** permission to read the document.
10. Save the document in the share **\\lon-svr1\docshare**.
11. Sign out of LON-CL1.

▶ Task 2: Verify internal access to protected content

1. Sign in to LON-CL1 with the **Adatum\Bill** account using the password **Pa$$w0rd**.
2. Add the http://adrms.adatum.com URL to the Local intranet group in Internet options Security tab using the Advanced button in Sites.

3. In the **\\lon-svr1\docshare** folder, open the **Executives Only** document.
4. When prompted, provide the credentials **Adatum\Bill** with the password of **Pa$$w0rd**.
5. Verify that you are unable to modify or save the document.
6. Select a line of text in the document.
7. Right-click the line of text. Verify that you cannot modify this text.
8. View the document permissions.
9. Sign out of LON-CL1.

▶ Task 3: Open the rights-protected document as an unauthorized user

1. Sign in to LON-CL1 as **Adatum\Carol** using the password **Pa$$w0rd**.
2. Add the http://adrms.adatum.com URL to the Local intranet group in Internet options Security tab using the Advanced button in Sites.
3. In the **\\lon-svr1\docshare** folder, attempt to open the **Executives Only** document.
4. Verify that Carol does not have permission to open the document.
5. Sign out of LON-CL1.

▶ Task 4: Open and edit the rights-protected document as an authorized user at Trey Research

1. Sign in to LON-CL1 with the **Adatum\Aidan** account using the password **Pa$$w0rd**.
2. Open Word 2013.
3. Create a new document named **\\LON-SVR1\docshare\TreyResearch-Confidential.docx**.
4. In the document, type the following text:

   This document is for Trey Research only, it should not be modified.
5. Restrict the permission so that **april@treyresearch.net** is able to open the document.
6. Sign in to Trey-CL1 with the **TREYRESEARCH\Administrator** account and the password **Pa$$w0rd**.
7. Add April as local Remote Desktop Users in the Systems properties.
8. Sign out of Trey-CL1.
9. Sign in to TREY-CL1 as **TREYRESEARCH\April** with the password Pa$$w0rd.
10. Add the http://adrms.treyresearch.net URL to the Local intranet group in Internet options Security tab using the Advanced button in Sites.
11. Use Windows Explorer to navigate to **\\LON-SVR1\docshare**. Use the credentials **Adatum\Administrator** and **Pa$$w0rd** to connect.
12. Copy the **TreyReserch-Confidential.docx** document to the desktop.
13. Attempt to open the document. When prompted, enter the following credentials, select the **Remember my credentials** check box, and then click **OK**:
    a. Username: **April**
    b. Password: **Pa$$w0rd**
14. Verify that you can open the document, but that you cannot make modifications to it.
15. View the permissions that the **april@treyresearch.com** account has for the document.

▶ Task 5: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20412C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20412C-LON-SVR1**, **20412C-TREY-DC1**, **20412C-LON-CL1**, and **20412C-TREY-CL1**.

Results: After completing this exercise, you will have verified that the AD RMS deployment is successful.

Question: What steps can you take to ensure that you can use Information Rights Management with the AD RMS role?

Module Review and Takeaways

Best Practces

- Before you deploy AD RMS, you must analyze your organization's business requirements and create the necessary templates. You should meet with users to inform them of AD RMS functionality, and also ask for feedback on the types of templates that they would like to have available.
- Strictly control membership of the Super Users group. Users in this group can access all protected content. Granting a user membership of this group gives them complete access to all AD RMS-protected content.

Review Questions

Question: What are the benefits of having an SSL certificate installed on the AD RMS server when you are performing AD RMS configuration?

Question: You need to provide access to AD RMS-protected content to five users who are unaffiliated contractors, and who are not members of your organization. Which method should you use to provide this access?

Question: You want to block users from protecting Office PowerPoint content by using AD RMS templates. What steps should you take to accomplish this goal?

Module 8

Implementing and Administering AD FS

Contents:

Module Overview

Active Directory Federation Services (AD FS) in the Windows Server® 2012 operating system provides flexibility for organizations that want to enable their users to log on to applications that are located on a local network, at a partner company, or in an online service. With AD FS, an organization can manage its own user accounts, and users only have to remember one set of credentials. However, those credentials can provide access to a variety of applications, which are located in a variety of places.

This module provides an overview of AD FS, and it provides details on how to configure AD FS in both a single-organization scenario and in a partner-organization scenario. Finally, this module describes the Web Application Proxy feature in Windows Server® 2012 R2 that functions as an AD FS proxy and reverse proxy for web-based applications.

Objectives

After completing this module, you will be able to:

- Describe AD FS.
- Describe how to deploy AD FS.
- Describe how to implement AD FS for a single organization.
- Describe how to deploy AD FS in a business-to-business federation scenario.
- Describe how to extend AD FS to external clients.

Lesson 1
Overview of AD FS

AD FS is the Microsoft implementation of an identity federation framework that enables organizations to establish federation trusts and share resources across organizational and Active Directory® Domain Services (AD DS) boundaries. AD FS is compliant with common Web services standards, thus enabling interoperability with identity federation solutions that other vendors provide.

AD FS addresses a variety of business scenarios where the typical authentication mechanisms used in an organization do not work. This lesson provides an overview of the concepts and standards that are implemented in AD FS, and the business scenarios that AD FS can address.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe identity federation.
- Describe claims-based identity.
- Describe Web services.
- Describe AD FS.
- Explain how AD FS enables single sign-on (SSO) in a single organization.
- Explain how AD FS enables SSO in a business-to-business federation.
- Explain how AD FS enables SSO Microsoft Online Services.

What Is Identity Federation?

Identity federation enables you to provide identification, authentication, and authorization across organizational and platform boundaries. You can implement identity federation within a single organization to enable access to diverse web applications, or between two organizations that have an established trust relationship.

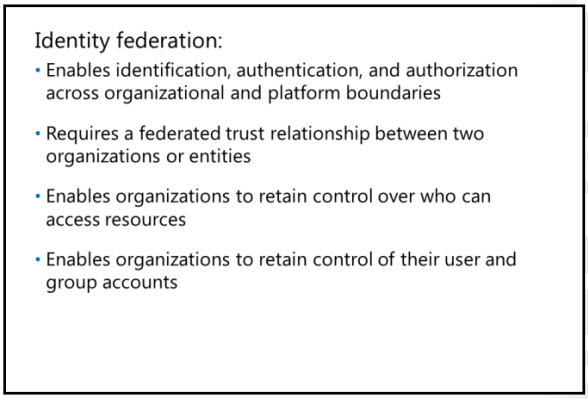

To establish an identity-federation partnership, both partners agree to create a federated-trust relationship. This federated trust is based on an ongoing business relationship, and it enables the organizations to implement business processes that are identified in the business relationship.

Note: A federated trust is not the same as a forest trust that organizations can configure between AD DS forests. In a federated trust, the AD FS servers in two organizations never have to communicate directly with each other. In addition, all communication in a federation deployment occurs over HTTPS, so you do not need to open multiple ports on any firewalls to enable federation.

As a part of the federated trust, each partner defines what resources are accessible to the other organization and how access to those resources is enabled. For example, to update a sales forecast, a sales

representative might need to collect information from a supplier's database that is hosted on the supplier's network. The administrator of the domain for the sales representative is responsible for ensuring that the appropriate sales representatives are members of the group that requires access to the supplier's database. The administrator of the organization where the database is located is responsible for ensuring that the partner's employees only have access to the data they require.

In an identity-federation solution, user identities and their associated credentials are stored, owned, and managed by the organization where the user is located. As part of the trust, each organization also defines how user identities are shared securely to restrict access to resources. Each partner must define the services that it makes available to trusted partners and customers, and which other organizations and users it trusts. Each partner also must define what types of credentials and requests it accepts, and each partner must define its privacy policies to ensure that private information is not accessible across the trust.

Identity federation also can be used within a single organization. For example, an organization might plan to deploy several web-based applications that require authentication. When you use AD FS, the organization can implement one authentication solution for all of the applications, making it easy for users in multiple internal domains or forests to access the application. The solution also can extend to external partners in the future, without changing the application.

What Is Claims-Based Identity?

In most organizations, users sign in to the network and are authenticated by an AD DS domain controller. A user who provides the right credentials to the domain controller is granted a security token. Applications that are running on servers in the same AD DS environment trust the security tokens that are provided by the AD DS domain controllers, because the servers can communicate with the same domain controllers where the users authenticate.

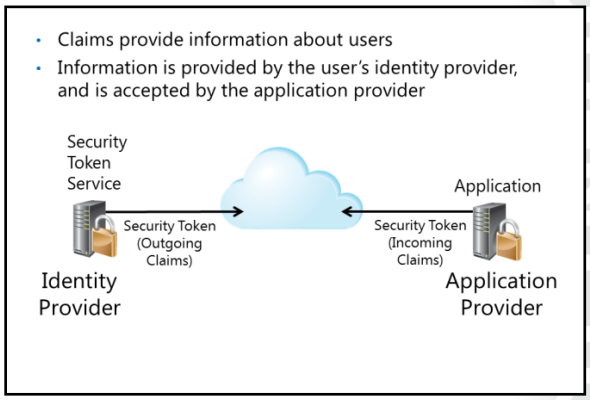

That type of authentication poses a problem because it does not extend outside of AD DS forest boundaries. Although trusts based on the Kerberos protocol or Windows NT LAN Manager (NTLM) can be implemented between two AD DS forests, client computers, and domain controllers on both sides of the trust must communicate with domain controllers in the other forest to make decisions about authentication and authorization. This communication requires network traffic that is sent on multiple ports, so these ports must be open on all firewalls between the domain controllers and other computers. The problem becomes even more complicated when users have to access resources that are hosted in cloud-based systems, such as Windows Azure™ or Microsoft Office® 365.

Claims-based authentication provides a mechanism for separating user authentication and authorization from individual applications. With claims-based authentication, users can authenticate to a directory service that is located within their organization and be granted a claim based on that authentication. The claim then can be presented to an application that is running in a different organization. The application allows user access to information or features based on the claims presented. All communication occurs over HTTPS.

The claim that is used in claims-based authentication is a statement about a user that is defined in one organization or technology and trusted in another. The claim could include a variety of information. For example, the claim could define the user's email address, user principal name (UPN), and information about specific groups to which the user belongs. This information is collected from the authentication mechanism when the user successfully authenticates.

The organization that manages the application defines the types of claims that the application will accept. For example, the application may require the user's email address to verify identity, and it then may use the group membership that is presented inside the claim to determine the level of access the user should have within the application.

Web Services Overview

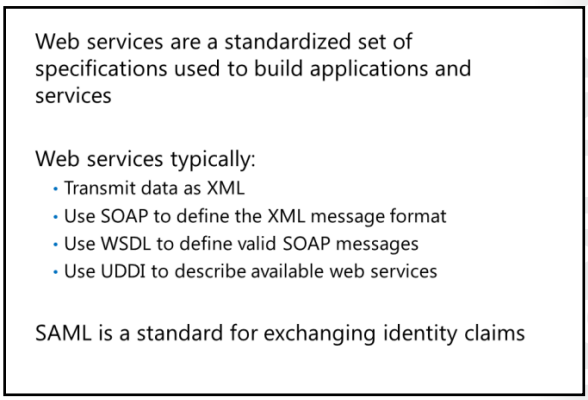

For claims-based authentication to work, organizations must agree on the format for exchanging claims. Rather than have each business define this format, a set of specifications broadly identified as Web services has been developed. Any organization interested in implementing a federated identity solution can use this set of specifications.

Web services are a set of specifications that are used for building connected applications and services, whose functionality and interfaces are exposed to potential users through web technology standards such as Extensible Markup Language (XML), Simple Object Access Protocol (SOAP), Web Services Description Language (WSDL), HTTP, and HTTPS. The goal of creating web applications by using Web services is to simplify interoperability for applications across multiple development platforms, technologies, and networks.

To enhance interoperability, Web services are defined by a set of industry standards. Web services are based on the following standards:

- Most Web services use XML to transmit data through HTTP and HTTPS. With XML, developers can create their own customized tags, thereby facilitating the definition, transmission, validation, and interpretation of data between applications and organizations.
- Web services expose useful functionality to web users through a standard web protocol. In most cases, the protocol SOAP is used; SOAP is the communications protocol for XML Web services. SOAP is a specification that defines the XML format for messages, and it essentially describes what a valid XML document looks like.
- Web services provide a way to describe their interfaces in enough detail to enable a user to build a client application to communicate with the service. This description usually is provided in an XML document called a WSDL document. In other words, a WSDL file is an XML document that describes a set of SOAP messages and how the messages are exchanged.
- Web services are registered, so that potential users can find them easily. This is done with Universal Description, Discovery, and Integration (UDDI). A UDDI directory entry is an XML file that describes a business and the services it offers.

WS-* Security Specifications

There are many components included in Web services specifications, which are also known as WS-* specifications. However, the most relevant specifications for an AD FS environment are the Web Services Security (WS-Security) specifications. The specifications included in WS-Security include the following:

- WS-Security: SOAP Message Security and X.509 Certificate Token Profile. WS-Security describes enhancements to SOAP messaging that provide quality of protection through message integrity, message confidentiality, and single-message authentication. WS-Security also provides a general-purpose, yet extensible, mechanism for associating security tokens with messages, and it provides a

mechanism to encode binary security tokens—specifically, X.509 certificates and Kerberos tickets—in SOAP messages.

- Web Services Trust (WS-Trust). WS-Trust defines extensions that build on WS-Security to request and issue security tokens and to manage trust relationships.
- Web Services Federation (WS-Federation). WS-Federation defines mechanisms that WS-Security can use to enable attribute-based identity, authentication, and authorization federation across different trust realms.
- WS-Federation Passive Requestor Profile (WS-F PRP). This WS-Security extension describes how passive clients, such as web browsers, can acquire tokens from a federation server, and how the clients can submit tokens to a federation server. Passive requestors of this profile are limited to the HTTP or HTTPS protocol.
- WS-Federation Active Requestor Profile (WS-F ARP). This WS-Security extension describes how active clients, such as SOAP-based mobile-device applications, can be authenticated and authorized, and how the clients can submit claims in a federation scenario.

Security Assertion Markup Language

The Security Assertion Markup Language (SAML) is an XML-based standard for exchanging claims between an identity provider and a service or application provider. SAML assumes that a user has been authenticated by an identity provider, and that the identity provider has populated the appropriate claim information in the security token. When the user is authenticated, the identity provider passes a SAML assertion to the service provider. On the basis of this assertion, the service provider can make authorization and personalization decisions within an application. The communication between federation servers is based on an XML document that stores the X.509 certificate for token-signing and the SAML 1.1 or 2.0 token.

What Is AD FS?

AD FS is the Microsoft implementation of an identity federation solution that uses claims-based authentication. AD FS provides the mechanisms to implement both the identity provider and the service provider components in an identity federation deployment.

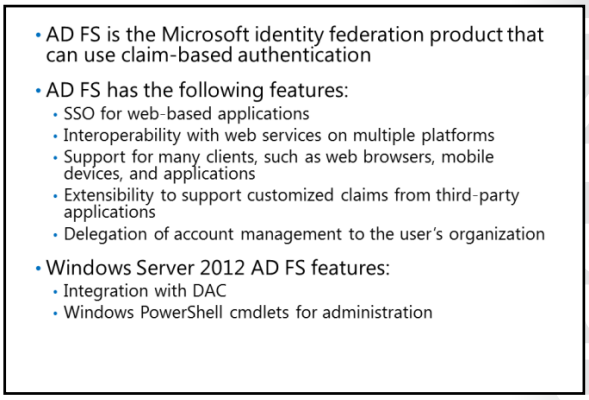

AD FS provides the following features:

- Enterprise *claims provider* for claims-based applications. You can configure an AD FS server as a claims provider, which means that it can issue claims about authenticated users. This enables an organization to provide its users with access to claims-aware applications in another organization by using SSO.
- Federation Service provider for identity federation across domains. This service offers federated web SSO across domains, thereby enhancing security and reducing overhead for information technology (IT) administrators.

Note: The Windows Server 2012 version of AD FS is built on AD FS version 2.0, which is the second generation of AD FS Microsoft released. The first version, AD FS 1.0, required AD FS Web Agents to be installed on all Web servers that were using AD FS, and it provided both claims-

aware and NT token-based authentication. AD FS 1.0 did not support active clients, but it did support SAML tokens.

AD FS Features

The following are some of the key features of AD FS:

- Web SSO. Many organizations have deployed AD DS. After authenticating to AD DS through Integrated Windows authentication, users can access all other resources that they have permission to access within the AD DS forest boundaries. AD FS extends this capability to intranet or Internet-facing applications, enabling customers, partners, and suppliers to have a similar, streamlined user experience when they access an organization's web-based applications.
- Web services interoperability. AD FS is compatible with the Web services specifications. AD FS employs the federation specification of WS-* called WS-Federation. WS-Federation makes it possible for environments that do not use the Windows identity model to federate with Windows environments.
- Passive and smart client support. Because AD FS is based on the WS-* architecture, it supports federated communications between any WS-enabled endpoints, including communications between servers and passive clients, such as browsers. AD FS on Windows Server 2012 also enables access for SOAP-based smart clients, such as mobile phones, personal digital assistants, and desktop applications. AD FS implements the WS-F PRP and some of the WS-F ARP standards for client support.
- Extensible architecture. AD FS provides an extensible architecture that supports various security-token types, including SAML tokens and Kerberos authentication through Integrated Windows authentication, and the ability to perform custom claims transformations. For example, AD FS can convert from one token type to another, or it can add custom business logic as a variable in an access request. Organizations can use this extensibility to modify AD FS to coexist with their current security infrastructure and business policies.
- Enhanced security. AD FS also increases the security of federated solutions by delegating responsibility for account management to the organization closest to the user. Each individual organization in a federation continues to manage its own identities, and each is capable of securely sharing and accepting identities and credentials from other members' sources.

New Features in Windows Server 2012 AD FS

The version of AD FS that ships with Windows Server 2012 includes several new features:

- Integration with the Windows Server 2012 operating system. In Windows Server 2012, AD FS is included as a server role that you can install by using Server Manager. When you install the server role, all required operating system components install automatically.
- Integration with Dynamic Access Control (DAC). When you deploy DAC, you can configure user and device claims that are issued by AD DS domain controllers. AD FS can consume the AD DS claims that domain controllers issue. This means that AD FS can make authorization decisions based on both user accounts and computer accounts.
- Windows PowerShell command-line interface cmdlets for administering AD FS. Windows Server 2012 provides several new cmdlets that you can use to install and configure the AD FS server role.

How AD FS Enables SSO in a Single Organization

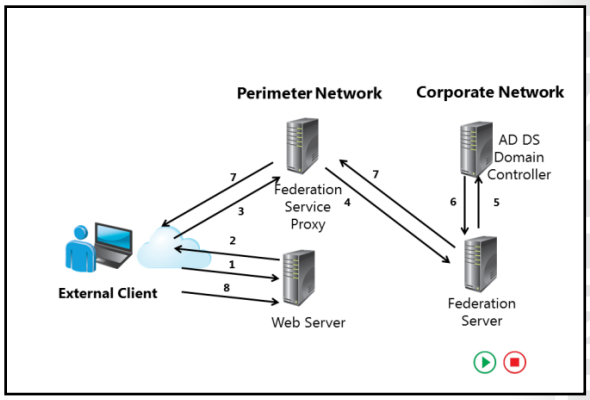

For many organizations, configuring access to applications and services might not require an AD FS deployment. If all users are members of the same AD DS forest, and if all applications run on servers that are members of the same forest, you usually can use AD DS authentication to provide access to the application. However, there are several scenarios where you can use AD FS to optimize the user experience by enabling SSO:

- The applications might not be running on Windows servers or on any servers that support AD DS authentication, or on Windows Server servers that are not domain-joined. The applications might require SAML or Web services for authentication and authorization.
- Large organizations frequently have multiple domains and forests that might result from mergers and acquisitions, or due to security requirements. Users in multiple forests might require access to the same applications.
- Users from outside the office might require access to applications that are running on internal servers. External users might log on to applications from computers that are not part of the internal domain.

Note: Implementing AD FS does not necessarily mean that users are not prompted for authentication when they access applications. Depending on the scenario, users might be prompted for their credentials. However, users always authenticate by using their internal credentials in the trusted account domain, and they never need to remember alternate credentials for the application. In addition, the internal credentials are never presented to the application or to the partner AD FS server.

Organizations can use AD FS to enable SSO in these scenarios. If the organization has a single AD DS forest, the organization only has to deploy a single federation server. This server can operate as the claims provider so that it authenticates user requests and issues the claims. The same server also is the relying party to provide authorization for application access.

Note: The slide and the following description use the terms Federation Service and Federation Service Proxy to describe AD FS role services. The federation server is responsible for issuing claims, and in this scenario, also is responsible for consuming the claims. The Federation Service Proxy is a proxy component that is recommended for deployments where users outside of the network need access to the AD FS environment. These components are covered in more detail in the next lesson.

The following steps describe the communication flow in this scenario:

1. The client computer, which is located outside of the network, must access a web-based application on the Web server. The client computer sends an HTTPS request to the Web server.
2. The Web server receives the request and identifies that the client computer does not have a claim. The Web server redirects the client computer to the Federation Service Proxy.

3. The client computer sends an HTTPS request to the Federation Service Proxy. Depending on the scenario, the Federation Service Proxy might prompt the user for authentication or use Integrated Windows authentication to collect the user's credentials.
4. The Federation Service Proxy passes on the request and the credentials to the federation server.
5. The federation server uses AD DS to authenticate the user.
6. If authentication is successful, the federation server collects AD DS information about the user. That information is then used to generate the user's claims.
7. If the authentication is successful, the authentication information and other information is collected in a security token and passed back to the client computer through the Federation Service Proxy.
8. The client then presents the token to the Web server. The web resource receives the request, validates the signed tokens, and uses the claims in the user's token to provide access to the application.

How AD FS Enables SSO in a Business-to-Business Federation

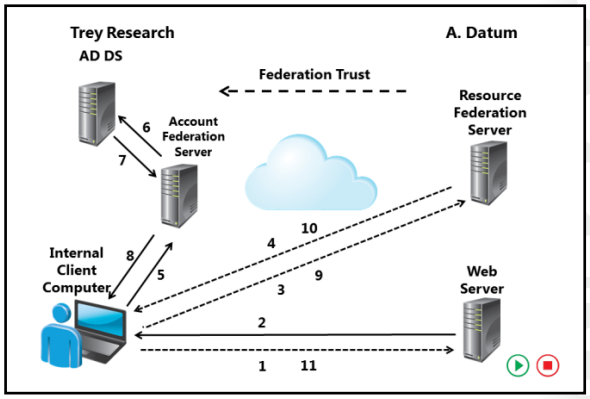

One of the most common scenarios for deploying AD FS is to provide SSO in a business-to-business federation. In this scenario, the organization that requires access to another organization's application or service can manage their own user accounts and define their own authentication mechanisms. The other organization can define which applications and services are exposed to users outside of the organization, and which claims it accepts to provide access to the application. To enable application or service sharing in this scenario, the organizations have to establish a federation trust and then define the rules for exchange claims between them.

The slide for this topic is an animated slide that demonstrates the flow of traffic in a federated business-to-business scenario by using a claims-aware web application. In this scenario, users at Trey Research have to access a web-based application at A. Datum Corporation. The AD FS authentication process for this scenario is as follows:

1. A user at Trey Research uses a web browser to establish an HTTPS connection to the Web server at A. Datum.
2. The web application receives the request, and verifies that the user does not have a valid token stored in a cookie by the web browser. Because the user is not authenticated, the web application redirects the client to the federation server at A. Datum by using an HTTP 302 redirect message.
3. The client computer sends an HTTPS request to A. Datum's federation server. The federation server determines the home realm for the user. In this case, the home realm is Trey Research.
4. The client computer is redirected again to the federation server in the user's home realm, which is Trey Research.
5. The client computer sends an HTTPS request to the Trey Research federation server.
6. If the user is logged on to the domain already, the federation server can take the user's Kerberos ticket and request authentication from AD DS on the user's behalf by using Integrated Windows authentication. If the user is not logged on to its domain, the user is prompted for credentials.

7. The AD DS domain controller authenticates the user and sends the success message back to the federation server, along with other information about the user that can be used to generate the user's claims.
8. The federation server creates the claim for the user based on the rules defined for the federation partner. The claims data is placed in a digitally signed security token, and then it is sent to the client computer, which posts it back to A. Datum's federation server.
9. A. Datum's federation server validates that the security token came from a trusted federation partner.
10. A. Datum's federation server creates and signs a new token, which it sends to the client computer. The client computer then sends the token back to the original URL requested.
11. The application on the Web server receives the request and validates the signed tokens. The Web server issues the client a session cookie, indicating that it has been authenticated successfully. A file-based persistent cookie is issued by the federation server, which is good for 30 days by default, to eliminate the home-realm discovery step during the cookie lifetime. The server then provides access to the application based on the claims provided by the user.

How AD FS Enables SSO with Online Services

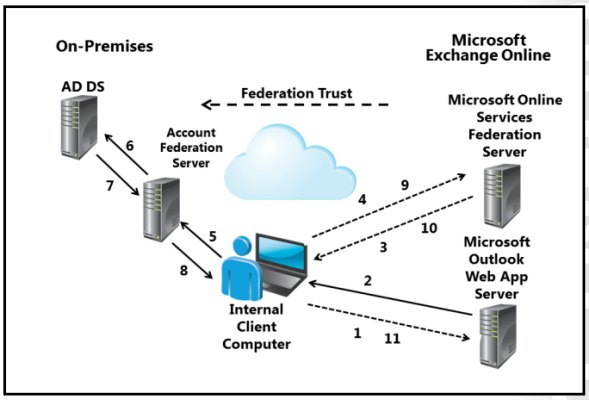

As organizations move services and applications to cloud-based services, it is increasingly important that these organizations have some way to simplify the authentication and authorization experience for their users as they utilize the cloud-based services. Cloud-based services add another level of complexity to the IT environment, as they are located outside the direct administrative control of IT administrators, and they can run on many different platforms.

You can use AD FS to provide an SSO experience to users across various available cloud-based platforms. For example, once users authenticate with AD DS credentials, they then could access Microsoft Online Services, such as Office 365, if they use those domain credentials.

AD FS also can provide SSO to non-Microsoft cloud providers. Because AD FS is based on open standards, it can interoperate with any compliant claims-based system.

The process for accessing a cloud-based application is similar to the business-to-business scenario. A hybrid Exchange Online deployment is an example of a cloud-based service that uses AD FS for authentication. In this type of deployment, an organization deploys some or all of its mailboxes in an Office 365 and an Exchange Online environment. However, the organization manages all of its user accounts in its on-premises AD DS environment. The deployment uses the Microsoft Online Services Directory Synchronization Tool to synchronize user account information from the on-premises deployment to the Exchange Online deployment.

When users try to sign in to their Exchange Online mailbox, users must authenticate by using their internal AD DS credentials. If users try to sign in directly to the Exchange Online environment, they are redirected back to the internal AD FS deployment to authenticate before they are given access.

The following steps describe what happens when a user tries to access his or her online mailbox by using a web browser:

1. The user opens a web browser and sends an HTTPS request to the Exchange Online Microsoft Outlook® Web App server.
2. The Outlook Web App server receives the request and verifies whether the user is part of a hybrid Exchange Server deployment. If this is the case, the server redirects the client computer to the Microsoft Online Services federation server.
3. The client computer sends an HTTPS request to the Microsoft Online Services federation server.
4. The client computer is redirected again to the on-premises federation server. The redirection to the user's home realm is based on the user's UPN suffix.
5. The client computer sends an HTTPS request to the on-premises federation server.
6. If the user is logged on to the domain already, the on-premises federation server can take the user's Kerberos ticket and request authentication from AD DS on the user's behalf by using Integrated Windows authentication. If the user logs on from outside of the network or from a computer that is not a member of the internal domain, the user is prompted for credentials.
7. The AD DS domain controller authenticates the user, and then sends the success message back to the federation server, along with other information about the user that the federation server can use to generate the user's claims.
8. The federation server creates the claim for the user based on the rules defined during the AD FS server setup. The claims data is placed in a digitally signed security token, and then it is sent to the client computer, which posts it back to the Microsoft Online Services federation server.
9. The Microsoft Online Services federation server validates that the security token came from a trusted federation partner. This trust is configured when you configure the hybrid Exchange Server environment.
10. The Microsoft Online Services federation server creates and signs a new token that it sends to the client computer, which then sends the token back to the Outlook Web App server.
11. The Outlook Web App server receives the request and validates the signed tokens. The server issues the client a session cookie indicating that it has authenticated successfully. The user then is granted access to his or her Exchange Server mailbox.

Lesson 2
Deploying AD FS

After you understand how AD FS works, you can deploy the service. Before you deploy AD FS, you must understand the components that you will need to deploy and the prerequisites that you must meet, particularly with regard to certificates. This lesson provides an overview of deploying the AD FS server role in Windows Server 2012 and Windows Server 2012 R2.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe AD FS components.
- Describe AD FS prerequisites.
- Describe the Public Key Infrastructure (PKI) and certificate requirements for AD FS.
- Describe AD FS federation server roles.
- Explain how to install the AD FS server role.

AD FS Components

AD FS is installed as a server role in Windows Server 2012. However, there are many different components that you install and configure in an AD FS deployment. The following table lists the AD FS components.

AD FS components:	
Federation server	Relying parties
Federation server proxy	Claims provider trust
Claims	Relying party trust
Claim rules	Certificates
Attribute store	Endpoints
Claims providers	

Component	What does it do?
Federation server	The federation server issues, manages, and validates requests involving identity claims. All implementations of AD FS require at least one Federation Service for each participating forest.
Federation server proxy	The federation server proxy is an optional component that you usually deploy in a perimeter network. It does not add any functionality to the AD FS deployment, but it is deployed to provide a layer of security for connections from the Internet to the federation server.
Claim	A claim is a statement that is made by a trusted entity about an object such as a user. The claim could include the user's name, job title, or any other factor that might be used in an authentication scenario. With Windows Server 2012, the object also can be a device used in a DAC deployment.

Component	What does it do?
Claim rules	Claim rules determine how claims are processed by federation servers. For example, a claim rule might state that an email address is accepted as a valid claim, or that a group name from one organization is translated into an application-specific role in the other organization. The rules usually are processed in real time as claims are made.
Attribute store	AD FS uses an attribute store to look up claim values. AD DS is a common attribute store and is available by default because the federation server role must be installed on a domain-joined server.
Claims providers	The claims provider is the server that issues claims and authenticates users. A claims provider is one side of the AD FS authentication and authorization process. The claims provider manages user authentication, and then issues the claims that the user presents to a relying party.
Relying party	The relying party is the party where the application is located, and it is the other side of the AD FS authentication and authorization process. The relying party is a web service that consumes claims from the claims provider. The relying party server must have the Microsoft Windows Identity Foundation (WIF) installed or use the AD FS 1.0 claims-aware agent.
Claims provider trust	A claims provider trust configures data that defines rules under which a client might request claims from a claims provider and subsequently submit them to a relying party. The trust consists of various identifiers such as names, groups, and various rules.
Relying-party trust	A relying-party trust defines the claim information about a user or client that will be passed from AD FS to a relying party. It consists of various identifiers, such as names, groups, and various rules.
Certificates	AD FS uses digital certificates when communicating over Secure Sockets Layer (SSL) or as part of the token-issuing process, the token-receiving process, and the metadata-publishing process. Digital certificates also are used for token signing.
Endpoints	Endpoints are Windows Communication Foundation mechanisms that enable access to AD FS technologies, including token issuance and metadata publishing. AD FS comes with built-in endpoints that are responsible for specific functionality.

 **Note:** Many of these components are described in more detail throughout the remainder of this module.

AD FS Prerequisites

Before you deploy AD FS, you must ensure that your internal network meets some basic prerequisites. The configuration of the following network services is critical for a successful AD FS deployment:

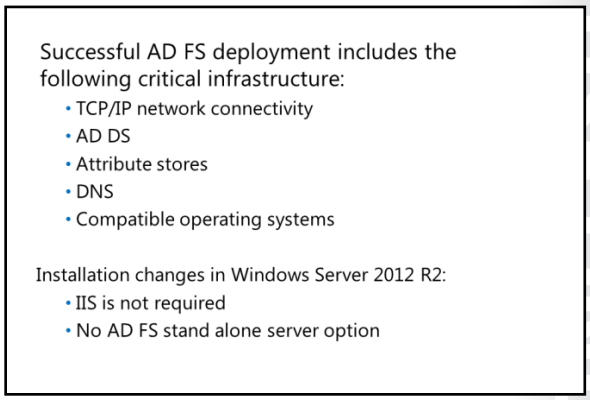

- Network connectivity. The following network connectivity is required:
  - The client computer must be able to communicate with the web application, the resource federation server or federation server proxy, and the account federation server or federation server proxy by using HTTPS.
  - The federation server proxies must be able to communicate with the federation servers in the same organization by using HTTPS.
  - Federation servers and internal client computers must be able to communicate with domain controllers for authentication.
- AD DS. AD DS is a critical piece of AD FS. Domain controllers should run Windows Server 2003 Service Pack 1 as a minimum. Federation servers must be joined to an AD DS domain. The Federation Service Proxy does not have to be domain-joined.
- Attribute stores. AD FS uses an attribute store to build claims information. The attribute store contains information about users, which is extracted from the store by the AD FS server after the user has been authenticated. AD FS supports the following attribute stores:
  - Active Directory Application Mode in Windows Server 2003
  - Active Directory Lightweight Directory Services (AD LDS) in Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012
  - Microsoft® SQL Server 2005, all editions
  - SQL Server 2008, all editions
  - A custom attribute store

Note: You can use AD DS as both the authentication provider and as an attribute store. AD FS can use AD LDS only as an attribute store.

- Domain Name System (DNS). Name resolution allows clients to find federation servers. Client computers must resolve DNS names for all federation servers or AD FS farms to which they connect, and the web applications that the client computer is trying to use. If a client computer is external to the network, the client computer must resolve the DNS name for the Federation Service Proxy, not the internal federation server or AD FS farm. The Federation Service Proxy must resolve the name of the internal federation server or farm. If internal users have to access the internal federation server directly, and external users have to connect through the federation server proxy, you must configure different DNS records in the internal and external DNS zones.
- Operating system prerequisites. You can only deploy the Windows Server 2012 version of AD FS as a server role on a Windows Server 2012 server.

Installation Changes in Windows Server 2012 R2

The version of AD FS included with Windows Server 2012 required the installation of Internet Information Services (IIS). Due in part to the IIS requirement, the installation of AD FS on domain controllers was not recommended for Windows Server 2012. In Windows Server 2012 R2, AD FS does not require the IIS installation , and installation on a domain controller is now acceptable.

During installation of AD FS for Windows Server 2012, you had an option to install AD FS as a stand-alone server. This option was useful for test environments, but it was not recommended for production environments because there were no options for expansion after installation. AD FS installation AD FS in Windows Server 2012 R2 does not include the option to install a stand-alone server. Instead, you can install a single server farm that provides the option for future expansion.

PKI and Certificate Requirements

AD FS enables computers to communicate securely, even though they might be in different locations. In this scenario, most of the communications between computers passes through the Internet. To provide security for the network traffic, all communications are protected by using SSL. This factor means that it is important to correctly choose and assign SSL certificates to the AD FS servers. To provide SSL security, AD FS servers use certificates as service-communication certificates, token-signing certificates, and token-decrypting certificates.

- Certificates used by AD FS:
  - Service communication certificates
  - Token-signing certificates
  - Token-decrypting certificates
- When choosing certificates, ensure that the service communication certificate is trusted by all federation partners and clients
- If you use an internal CA then users must have access to certificate revocation information

Service Communication Certificates

AD FS secures all communication by using SSL, which requires a certificate. The certificate used for service communication must be trusted by all computers that communicate with the AD FS server. If all of the computers and devices that contact your AD FS server are domain-joined, then you can consider using an internally generated certificate for AD FS. However, in most cases, at least some communication is between the AD FS server and external computers or partner organizations, in which case, a certificate from the third-party certification authority (CA) should be used.

In Windows Server 2012, AD FS uses the Default Web Site in IIS to provide Web services. Consequently, service-communication certificate management was performed in IIS Manager.

In Windows Server 2012 R2, IIS is no longer used by AD FS. You can use the certificates snap-in and the AD FS Management console for all management of certificates.

Note: If you change the service communication certificate after initial configuration, you must change it on all nodes in the server farm, and you must ensure that the AD FS service is granted Read permissions to the private key on the certificate on each node.

Token-Signing Certificates

The token-signing certificate is used to sign every token that a federation server issues. This certificate is critical in an AD FS deployment because the token signature indicates which federation server issued the token. The claims provider uses this certificate to identify itself, and the relying party uses it to verify that the token is coming from a trusted federation partner.

The relying party also requires a token-signing certificate to sign the tokens that it prepares for AD FS-aware applications. These tokens must be signed by the relying party's token-signing certificate to be validated by destination applications.

When you configure a federation server, the server assigns a self-signed certificate as the token-signing certificate. In most cases, it is not required to update this certificate with a certificate from a third-party CA. When a federation trust is created, the trust of this certificate is configured at the same time. You can have multiple token-signing certificates configured on the federation server, but only the primary certificate is used to sign tokens.

Token-Decrypting Certificates

Token-decrypting certificates are used to encrypt the entire user token before transmitting the token across the network from the claims provider federation server to the relying party federation server. To provide this functionality, the public key from the relying party federation server certificate is provided to the claims provider federation server. The certificate is sent without the private key. The claims provider server uses the public key from the certificate to encrypt the user token.

When the token is returned to the relying party federation server, it uses the private key from the certificate to decrypt the token. This provides an extra layer of security when transmitting the certificates across an untrusted network such as the Internet.

When you configure a federation server, the server assigns a self-signed certificate as the token-decrypting certificate. In most cases, you do not need to update this certificate with a certificate from a third-party CA. When a federation trust is created, the trust of this certificate is configured at the same time.

 **Note:** The federation server proxies only require a SSL certificate. The certificate is used to enable SSL communication for all client connections.

Choosing a CA

AD FS federation servers can use self-signed certificates, certificates from an internal, private CA, or certificates that have been purchased from an external, public CA. In most AD FS deployments, the most important factor when you choose certificates is that they be trusted by all parties involved. This means that if you configure an AD FS deployment that interacts with other organizations, you almost certainly will use a public CA for the SSL certificate on a federation server proxy, because the certificates issued by the public CA are trusted by all partners automatically.

If you deploy AD FS just for your organization, and all servers and client computers are under your control, you can consider using a certificate from an internal, private CA. If you deploy an internal enterprise CA on Windows Server 2012, you can use Group Policy to ensure that all computers in the organization automatically trust the certificates issued by the internal CA. Using an internal CA can decrease the cost of certificates significantly.

If you use an internal CA, you must also ensure that a certificate revocation can be verified by users from any location. For example, if your users access applications from the Internet, then you must ensure that those users can access certificate revocation information from the Internet. This means that you need to configure a certificate revocation list (CRL) distribution point in your perimeter network.

 **Note:** Deploying an internal CA by using Active Directory Certificate Services is a straightforward process, but it is critical that you plan and implement the deployment carefully.

Federation Server Roles

In Windows Server 2012, when you install the AD FS server role, you can configure the server as either a federation server or a federation server proxy. In Windows Server 2012, the federation server proxy installation option has been removed; it is replaced by the Web Application Proxy role service in the remote access server role. Web Application Proxy functions as a federation server proxy and has additional functionality. The server roles for AD FS are:

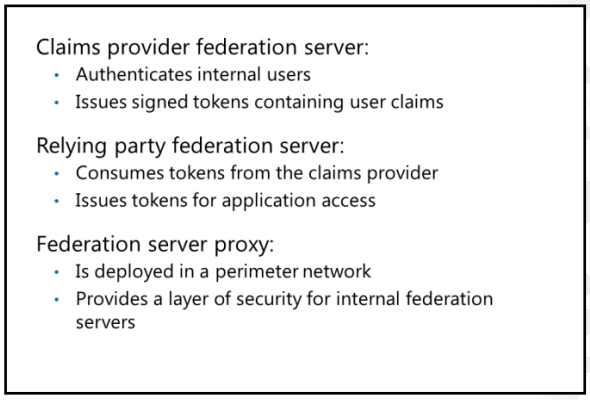

- Claims provider. A claims provider is a federation server that provides users signed tokens that contain claims. Claims provider federation servers are deployed in organizations where user accounts are located. When a user requests a token, the claims provider federation server verifies user authentication by using AD DS, and then it collects information from an attribute store, such as AD DS or AD LDS, to populate the user claim with the attributes required by the partner organization. The server issues tokens in SAML format. The claims provider federation server also protects the contents of security tokens in transit by signing and optionally encrypting them.
- Relying party. A relying party is a federation server that receives security tokens from a trusted claims provider. Relying party federation servers are deployed in organizations that provide application access to claims provider organizations. The relying party accepts and validates the claim, and then it issues new security tokens that the Web server can use to provide appropriate access to the application.

Note: A single AD FS server can operate as both a claims provider and a relying party, even with the same partner organizations. The AD FS server functions as a claims provider when it authenticates users and provides tokens for another organization, but it also can accept tokens from the same or different organizations in a relying-party role.

- Federation server proxy. A federation server proxy provides an extra level of security for AD FS traffic that comes from the Internet to internal AD FS federation servers. Federation server proxies can be deployed in both claims-provider and relying-party organizations. On the claims provider side, the proxy collects the authentication information from client computers and passes it to the claims provider federation server for processing. The federation server issues a security token to the proxy, which sends it to the relying-party proxy. The relying-party federation server proxy accepts these tokens, and then passes them on to the internal federation server. The relying-party federation server issues a security token for the web application, and then it sends the token to the federation-server proxy, which then forwards the token to the client. The federation-server proxy does not provide any tokens or create claims; it only forwards requests from clients to internal AD FS servers. All communication between the federation-server proxy and the federation server uses HTTPS.

Note: A federation-server proxy cannot be configured as a claims provider or a relying party. The claims provider and relying party must be members of an AD DS domain. The federation-server proxy can be configured as a member of a workgroup, or as a member of an extranet forest, and it can be deployed in a perimeter network.

Demonstration: Installing the AD FS Server Role

In this demonstration, you will see how to install and complete the initial configuration of the AD FS server role in Windows Server 2012 R2. The instructor will install the server role, and then run the AD FS Federation Server Configuration Wizard.

Demonstration Steps

Install AD FS

- On LON-DC1, use the Server Manager to install the Active Directory Federation Services role on LON-DC1.Adatum.com.

Add a DNS record for AD FS

- On LON-DC1, use the DNS Manager to add a new host record for AD FS in the **Adatum.com** forward lookup zone with the following settings:
    - Name: **adfs**
    - IP address: **172.16.0.10**

Configure AD FS

1. In the Server Manager notifications, click **Configure the federation services on this server**.
2. Use the following options to configure the AD FS server:
    - **Create the first federation server in a federation server farm**
    - Account for configuration: **Adatum\Administrator**
    - SSL Certificate: **a.adatum.com**
    - Create a Group Managed Service Account: **Adatum\ADFS**
    - **Create a database on this server using Windows Internal Database**

Lesson 3
Implementing AD FS for a Single Organization

The simplest deployment scenario for AD FS is within a single organization. In this scenario, a single AD FS server can operate as both the claims provider and the relying party. All users in this scenario are internal to the organization, as is the application that the users access.

This lesson provides details on the components that are required to configure AD FS in a single-organization AD FS deployment. These components include configuring claims, claim rules, claims-provider trusts, and relying-party trusts. This lesson also provides information about the authentication policies available in AD FS to control the authentication process for users in different locations. Using multifactor authentication to increase security is also discussed.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe AD FS claims.
- Describe AD FS claim rules.
- Describe a claims-provider trust.
- Describe a relying-party trust.
- Describe how to configure claims-provider and relying-party trusts.
- Describe authentication policies.
- Describe multifactor authentication.

What Are AD FS Claims?

AD FS claims provide the link between the claims-provider and relying-party roles in an AD FS deployment. An AD FS claim is a statement made about a particular subject, such as a user, by a trusted entity, such as a claims provider. The claims provider creates the claims, and the relying party consumes the claims. AD FS claims provide a standards-based and flexible way for claims provider organizations to provide specific information about users in their organizations. AD FS claims also provide a way for relying parties to define exactly what information they require to provide application access. The claim information provides the details required by applications to enable access to claims-aware applications.

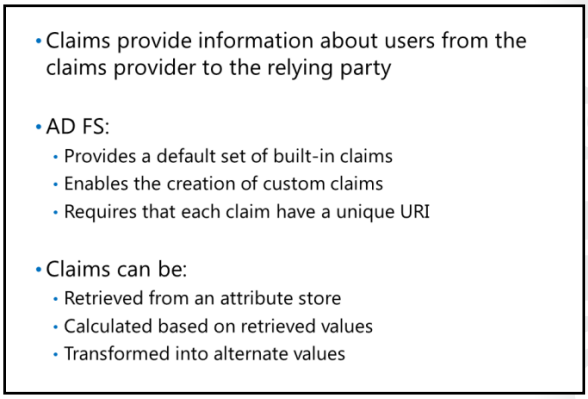

Claim Types

Each AD FS claim has a claim type, such as email address, UPN, or last name. Users can be issued claims based on any defined claim type. Therefore, a user might be issued a claim with a type of Last Name and a value of, for example, Weber. AD FS provides many built-in claim types. Optionally, you can create new ones based on organizational requirements.

Each AD FS claim type is identified by a Uniform Resource Identifier (URI) that uniquely identifies the claim type. This information is provided as part of the AD FS server metadata. For example, if the claims-

provider organization and the relying-party organization decide to use a claim type of AccountNumber, both organizations must configure a claim type with this name. The claim type is published and the claim type URI must be identical on both AD FS servers.

Note: In Windows Server 2012 R2, the number of claims types has increased to support various device types and certificate characteristics.

How Claim Values Are Populated

Claims issued by a claims provider contain the information that is required by the relying party to enable appropriate application access. One of the first steps in planning an AD FS deployment is to define exactly what information the applications must have about each user to provide that user access to the application. Once this information is defined, the claims then are defined on the claims provider federation server. The information required to populate the claim can be obtained in several ways:

- The claim can be retrieved from an attribute store. Frequently, the information required for the claim is already stored in an attribute store that is available to the federation server. For example, an organization might decide that the claim should include the user's UPN, email address, and specific group memberships. This information is stored in AD DS already, so the federation server can retrieve this information from AD DS when it creates the claim. Because AD FS can use AD DS, AD LDS, SQL Server, a non-Microsoft Lightweight Directory Access Protocol (LDAP) directory, or a custom attribute store to populate claims, you can define almost any value within the claim.
- The claim can be calculated based on collected information. Claims-provider federation servers also can calculate information based on information that is gathered from an attribute store. For example, a vendor's database may contain weight of inventory in pounds, while your application requires the weight in kilograms to calculate shipping costs. A calculated claim could make the conversion from pounds to kilograms.
- The claim can be transformed from one value to another. In some cases, the information that is stored in an attribute store does not exactly match the information the application requires when it creates authorization information. For example, the application might have different user roles defined that do not directly match the attributes that are stored in any attribute store. However, the application role might correlate to AD DS group membership. For example, users in the Sales group might correlate to one application role, while users in the Sales Management group might correlate to a different application role. To establish the correlation in AD FS, you can configure a claims transformation that takes the value provided by the claims provider and translates the value into to a claim that is useful to the application in the relying party.
- If you have deployed DAC, a DAC device claim can be transformed into an AD FS claim. This can be used to ensure that users can access an AD FS website only from trusted workstations that have been issued a valid device claim.

What Are AD FS Claim Rules?

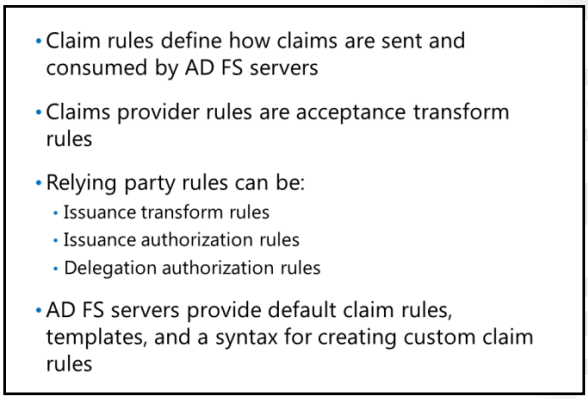

Claim rules define how claims are sent and consumed by AD FS servers. Claim rules define the business logic that is applied to claims that are provided by claims providers, and to claims that are accepted by the relying parties. You can use claim rules to:

- Define which incoming claims are accepted from one or more claims providers.
- Define which outbound claims are provided to one or more relying parties.
- Apply authorization rules to enable access to a specific relying party for one or more users or groups of users.

You can define two types of claim rules:

- Claim rules for a claims provider trust. A claims provider trust is the AD FS trust relationship that is configured between an AD FS server and a claims provider. You can configure claim rules to define how the claims provider processes and issues claims.
- Claim rules for a relying-party trust. A relying-party trust is the AD FS trust relationship that is configured between an AD FS server and a relying party. You can configure claim rules that define how the relying party accepts claims from the claims provider.

Claim rules configured on an AD FS claims provider all are considered *acceptance transform rules*. These rules determine what claim types are accepted from the claims provider and then sent to a relying-party trust. When configuring AD FS within a single organization, there is a default claims provider trust that is configured with the local AD DS domain. This rule set defines the claims that are accepted from AD DS.

There are three types of claim rules for a relying-party trust:

- Issuance transform rules. These rules define the claims that are sent to the relying party that was defined in the relying party trust.
- Issuance authorization rules. These rules define which users are permitted or denied access to the relying party defined in the relying-party trust. This rule set can include rules that explicitly permit access to a relying party, and/or rules that explicitly deny access to a relying party.
- Delegation-authorization rules. These rules define the claims that specify which users can act on behalf of other users when accessing the relying party. This rule set can include rules that explicitly permit delegates for a relying party, or rules that explicitly deny delegates to a relying party.

Note: AD FS servers are preconfigured with a set of default rules and several default templates that you can use to create common claim rules. You can create custom claim rules by using the AD FS claim rule language.

What Is a Claims-Provider Trust?

A *claims-provider trust* is configured on the relying-party federation server. The claims-provider trust identifies the claims provider and describes how the relying party consumes the claims that the claims provider issues. You must configure a claims provider trust for each claims provider. A claims-provider trust for the local AD DS is configured by default. You must configure any additional claims providers.

- Claims provider trusts:
  - Are configured on the relying party federation server
  - Identify the claims provider
  - Configure the claim rules for the claims provider
- In a single-organization scenario, a claims provider trust called Active Directory defines how AD DS user credentials are processed
- Additional claims provider trusts can be configured by:
  - Importing the federation metadata
  - Importing a configuration file
  - Configuring the trust manually

By default, an AD FS server is configured with a claims provider trust named Active Directory. This trust defines the claim rules, which are all acceptance-transform rules that define how the AD FS server accepts AD DS credentials. For example, the default claim rules on the claims-provider trust include rules that pass user names, security identifiers (SIDs), and group SIDs to the relying party. In a single-organization AD FS deployment where AD DS authenticates all users, the default claims-provider trust might be the only claims-provider trust required.

When you expand an AD FS deployment to include other organizations, you must create additional claims-provider trusts for each federated organization that is an identity provider. When you configure a claims-provider trust, you have three options:

- Import data about the claims provider through the federation metadata. If the AD FS federation server or federation server proxy is accessible through the network from your AD FS federation server, you can enter the host name or URL for the partner federation server. Your AD FS federation server connects to the partner server and downloads the federation metadata from the server. The federation metadata includes all the information that is required to configure the claims-provider trust. As part of the federation metadata download, your federation server also downloads the SSL certificate that is used by the partner federation server.
- Import data about the claims provider from a file. Use this option if the partner federation server is not directly accessible from your federation server, but the partner organization has exported its configuration and provided you the information in a file. The configuration file must include configuration information for the partner organization, and the SSL certificate that the partner federation server uses.
- Manually configure the claims provider trust. Use this option if you want to configure all of the settings for the claims-provider trust. When you choose this option, you must provide the features that the claims provider supports and the URL that is used to access the claims provider AD FS servers. You also must add the SSL certificate that the partner organization uses.

What Is a Relying-Party Trust?

A *relying-party trust* is defined on the claims provider federation server. The relying-party trust identifies the relying party and also defines the claim rules that define how the relying party accepts and processes claims from the claims provider.

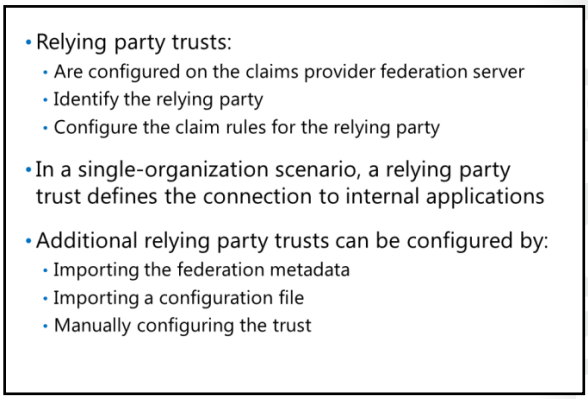

In a single-organization scenario, the relying-party trust defines how the AD FS server interacts with the applications deployed within the organization. When you configure the relying-party trust in a single organization, you provide the URL for the internal application. You can also configure settings such as the URL used by the Web server, the issuance authorization rules for the application, and whether the application supports SAML 2.0 or requires AD FS 1.0 tokens.

The process for configuring a relying-party trust is similar to that used for a claims provider trust. When you expand the AD FS deployment to include other organizations, you must create additional relying-party trusts for each federated organization. When you configure a relying-party trust, you have three options:

- Import data about the relying party through the federation metadata. If the AD FS federation or federation server proxy is accessible through the network from your AD FS federation server, you can enter the host name or URL for the partner federation server. Your AD FS federation server connects to the partner server and then downloads the federation metadata from the server. The federation metadata includes all the information that is required to configure the relying-party trust. As part of the federation metadata download, your federation server also downloads the SSL certificate that the partner federation server uses.
- Import data about the relying party from a file. Use this option if the partner federation server is not accessible from your federation server directly. In this case, the partner organization can export its configuration information to a file and then provide it to you. The configuration file must include configuration information for the partner organization and the SSL certificate that the partner federation server uses.
- Manually configure the claims-provider trust. Use this option if you want to configure all of the settings for the trust.

Demonstration: Configuring Claims Provider and Relying Party Trusts

In this demonstration, you will see how to configure claims-provider trusts and relying-party trusts. The instructor will demonstrate how to edit the default Active Directory claims-provider trust. The instructor also will create a new relying-party trust and demonstrate how to configure the trust.

Demonstration Steps

Configure a Claims Provider Trust

1. On LON-DC1, in the Server Manager, open the AD FS Management tool.
2. Browse to the **Claims Provider Trusts**, and then edit claim rules for **Active Directory**.

3. Add an acceptance transform rule with the following settings:
    - Claim rule template: **Send LDAP Attributes as Claims**
    - Claim rule name: **Outbound LDAP Attributes Rule**
    - Attribute store: **Active Directory**
    - Mapping of LDAP attributes:
        - E-Mail-Addresses: **E-Mail Address**
        - User-Principal-Name: **UPN**

Configure a certificate for a web-based app

1. On LON-SVR1, open Internet Information Services (IIS) Manager and view the server certificates.
2. Create a new domain certificate with the following settings:
    - Common name: **lon-svr1.adatum.com**
    - Organization: **A. Datum**
    - Organizational unit: **IT**
    - City/locality: **London**
    - State/Province: **England**
    - Country/region: **GB**
    - Certification Authority: **AdatumCA**
    - Friendly name: **AdatumTestApp Certificate**
3. Add an https binding for the Default Web Site:
    - SSL certificate: **AdatumTestApp Certificate**

Configure a WIF application for AD FS

1. On LON-SVR1, in the Server Manager, open the Windows Identity Foundation Federation Utility tool.
2. Enter the following in the Federation Utility Wizard:
    - Application configuration location: **C:\inetpub\wwwroot\AdatumTestApp\web.config**
    - Application URI: **https://lon-svr1.adatum.com/AdatumTestApp/**
    - **Use an existing STS**
    - STS WS-Federation metadata document location: **https://adfs.adatum.com/federationmetadata/2007-06/federationmetadata.xml**
    - **Disable certificate chain validation**,
    - **No encryption**

Configure a Relying-Party Trust

1. On LON-DC1, in the AD FS console, add a relying-party trust with the following settings:
    - **Import data about the relying party published online or on a local network**
    - Federation metadata address: **https://lon-svr1.adatum.com/AdatumTestApp/**
    - Display name: **A. Datum Test App**

- **I do not want to configure multi-factor authentication settings for the relying party trust at this time**
- **Permit all users to access this relying party**

2. Leave the Edit Claim Rules for A. Datum Test App window open for the next demonstration.

What Are Authentication Policies?

The AD FS in Windows Server 2012 R2 uses authentication policies to control how authentication is performed by AD FS. You can configure authentication settings globally and through relying-party trust. This allows you to vary some configuration settings for a specific application or when integrating with an external organization. You can configure multifactor authentication as part of an authentication policy.

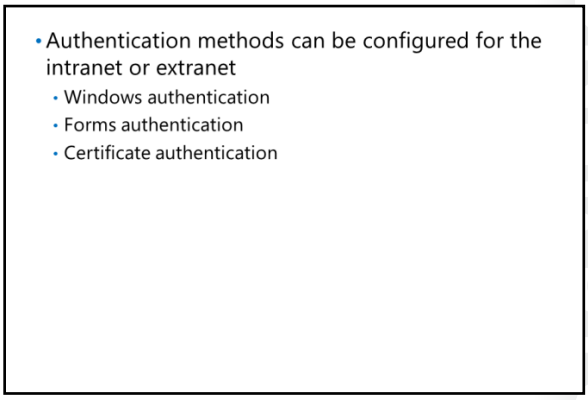

Authentication Methods

You can use the global authentication policy to define which authentication methods are supported for the intranet and extranet. The intranet methods are supported on the AD FS server on the internal network. The extranet methods are supported by the AD FS proxy functionality on the Web Application Proxy server.

The authentication methods are:

- Windows authentication. When you use Windows authentication, workstation credentials can be passed directly to AD FS if the application being accessed is in the local intranet zone of Internet Explorer. Otherwise, the user is prompted for credentials by a pop-up window. This authentication method may experience issues when traversing some firewalls and when used with web browsers other than Internet Explorer. This authentication method is supported only for the intranet.
- Forms authentication. This authentication method presents a web page that is used to enter authentication credentials. Use forms authentication to provide better compatibility for users accessing applications from outside the organization. This authentication method is available for the intranet and extranet.
- Certificate authentication. This authentication method accepts a certificate from the web browser as an alternative to a username and password. You can use certificate authentication to increase security of credentials entering because it is typically more difficult to steal a certificate than a username and password. This authentication method is available for the intranet and extranet.

It is possible to select multiple authentication methods for the intranet or extranet. If you select multiple authentication methods, then any of the selected methods can be used. Browsers that support Windows authentication will use it as the default authentication method.

What Is Multifactor Authentication?

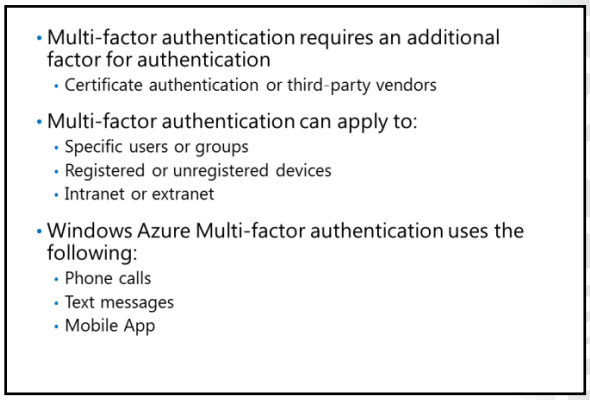

Multifactor authentication increases security for user authentication. Standard authentication is based on a user name and password. Adding a second factor for authentication makes it more difficult for user credentials to be used by an unauthorized person. The second factor is something that the user must possess and not just know.

There are several ways that you can implement multifactor authentication. You can use certificate authentication as an additional authentication method for AD FS. The use of certificate authentication for multifactor authentication is supported by default, and can be done by using smart cards.

There are additional multifactor authentication methods that can be integrated with AD FS from third-party vendors. One common method for multifactor authentication requires users to carry a device with a number that changes periodically. That number needs to be entered when the user authenticates. Microsoft also offers Windows Azure Multi-Factor Authentication, which uses smart phones as the second factor for authentication.

You can require multifactor authentication for:

- Specific users and groups
- Registered or unregistered devices
- Intranet or extranet

Windows Azure Multi-Factor Authentication

Windows Azure Multi-Factor Authentication allows you to use smart phones as the second factor for authentication. If you integrate Windows Azure Multi-Factor Authentication with AD FS, you can implement the following methods for additional authentication:

- Phone calls. When this method is used, you receive a call on your phone to confirm your authentication. You press the # (pound or hash) symbol to confirm after receiving the call.
- Text messages. When this method is used, you receive a text message with a passcode. You respond to the text message and include the passcode.
- Mobile App. When this method is used, an authentication prompt appears in the mobile app that you must acknowledge.

You can use Windows Azure Multi-Factor Authentication for many scenarios other than AD FS authentication. It can be integrated into many situations where increased security is required. For example, you could use it for authentication to virtual private networks (VPNs), cloud-based applications hosted in Windows Azure, RADIUS servers, or AD DS.

To learn about Windows Azure Multi-Factor Authentication, go to:
http://go.microsoft.com/fwlink/?LinkID=386642

Lab A: Implementing AD FS

Scenario

A. Datum Corporation has set up a variety of business relationships with other companies and customers. Some of these partner companies and customers must access business applications that are running on the A. Datum network. The business groups at A. Datum want to provide a maximum level of functionality and access to these companies. The Security and Operations departments want to ensure that the partners and customers can access only the resources to which they require access, and that implementing the solution does not increase the workload for the Operations team significantly. A. Datum also is working on migrating some parts of its network infrastructure to Microsoft Online Services, including Windows Azure and Office 365.

To meet these business requirements, A. Datum plans to implement AD FS. In the initial deployment, the company plans to use AD FS to implement SSO for internal users who access an application on a Web server.

As one of the senior network administrators at A. Datum, it is your responsibility to implement the AD FS solution. As a proof-of-concept, you plan to deploy a sample claims-aware application, and you will configure AD FS to enable internal users to access the application.

Objectives

After completing this lab, you will be able to:

- Install and configure AD FS.
- Configure an internal application for AD FS.

Lab Setup

Estimated Time: 30 minutes

Virtual machines: 20412C-LON-DC1,

20412C-LON-SVR1,

20412C-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 **2041C-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 by using the following credentials:
    - User name: **Adatum\Administrator**
    - Password: **Pa$$w0rd**
5. Repeat steps 2 to 3 for 20412C-LON-SVR1, and 20412C-LON-CL1.
    a. Sign in to 20412C-LON-DC1, and 20412C-LON-SVR1 as **Adatum\Administrator**.
    b. Do not sign in to 20412C-LON-CL1.

Exercise 1: Installing and Configuring AD FS

Scenario

To start the AD FS implementation, you need to install AD FS on an A. Datum domain controller. During the initial deployment, you will configure it as the first server in a farm with the option to expand the farm at a later time. The certificate for AD FS already has been installed on LON-DC1.

The main tasks for this exercise are as follows:

- Create a DNS record for AD FS.
- Create a service account.
- Install AD FS.
- Configure AD FS.
- Verify AD FS functionality.

▶ Task 1: Create a DNS record for AD FS

- On LON-DC1, use the DNS Manager to add a new host record for AD FS:
  - Forward lookup zone: **Adatum.com**
  - Name: **adfs**
  - IP address: **172.16.0.10**

▶ Task 2: Create a service account

1. On LON-DC1, open a Windows PowerShell prompt.
2. Create a new user account:
   - **New-ADUser –Name adfsService**
3. Set a password for adfsService:
   - **Set-ADAccountPassword adfsService**
   - Current password: none (press Enter)
   - Desired password: **Pa$$w0rd**
4. Enable the adfsService account:
   - **Enable-ADAccount adfsService**

▶ Task 3: Install AD FS

- On LON-DC1, in the Server Manager, add the Active Directory Federation Services role.

▶ Task 4: Configure AD FS

1. On LON-DC1, in the Server Manager notifications, click **Configure the federation services on this server**.
2. Use the following options to configure the AD FS server:
   - **Create the first federation server in a federation server farm**
   - Account for configuration: **Adatum\Administrator**
   - SSL Certificate: **adfs.adatum.com**
   - Federation Service Display Name: **A. Datum Corporation**

- Use an existing domain user account or group Managed Service Account
  - **Adatum\adfsService**
  - Password: **Pa$$w0rd**
  - **Create a database on this server using Windows Internal Database**

Note: The adfs.adatum.com certificate was preconfigured for this task. In your own environment, you need to obtain this certificate.

▶ Task 5: Verify AD FS functionality

1. On LON-CL1, sign in as **Adatum\Brad** with the password **Pa$$w0rd**.
2. Use Internet Explorer to access https://adfs.adatum.com/federationmetadata/2007-06/federationmetadata.xml.
3. Verify that the file loads, and then close Internet Explorer.

Results: In this exercise, you installed and configured AD FS. You also verified that it is functioning by viewing the FederationMetaData.xml file contents.

Exercise 2: Configuring an Internal Application for AD FS

Scenario

The first scenario for implementing the proof-of-concept AD FS application is to ensure that internal users can use SSO to access the web application. You plan to configure the AD FS server and a web application to enable this scenario. You also want to verify that internal users can access the application.

The main tasks for this exercise are as follows:

- Configure a certificate for the application.
- Configure the Active Directory claims-provider trust.
- Configure the application to trust incoming claims.
- Configure a relying-party trust for the claims-aware application.
- Configure claim rules for the relying-party trust.
- Test access to the claims-aware application.
- Configure Internet Explorer to pass local credentials to the application automatically.

▶ Task 1: Configure a certificate for the application

1. On LON-SVR1, open Internet Information Services (IIS) Manager and view the server certificates.
2. Create a new domain certificate with the following settings:
   - Common name: **lon-svr1.adatum.com**
   - Organization: **A. Datum**
   - Organizational unit: **IT**
   - City/locality: **London**
   - State/Province: **England**

- Country/region: **GB**
- Certification Authority: **AdatumCA**
- Friendly name: **AdatumTestApp Certificate**

3. Add an https binding for the Default Web Site:
   - SSL certificate: **AdatumTestApp Certificate**

▶ Task 2: Configure the Active Directory claims-provider trust

1. On LON-DC1, in the Server Manager, open the AD FS Management tool.
2. Browse to the **Claims Provider Trusts** and edit claim rules for **Active Directory**.
3. Add an acceptance transform rule with the following settings:
   - Claim rule template: **Send LDAP attributes as claims**
   - Name: **Outbound LDAP Attributes Rule**
   - Attribute store: **Active Directory**
   - Mapping of LDAP attributes:
     - E-Mail-Addresses: **E-Mail Address**
     - User-Principal-Name: **UPN**
     - Display-Name: **Name**

▶ Task 3: Configure the application to trust incoming claims

1. On LON-SVR1, in the Server Manager, open the Windows Identity Foundation Federation Utility tool.
2. Enter the following in the Federation Utility Wizard:
   - Application configuration location: **C:\inetpub\wwwroot\AdatumTestApp\web.config**
   - Application URI: **https://lon-svr1.adatum.com/AdatumTestApp/**
   - **Use an existing STS**
   - STS WS-Federation metadata document location: **https://adfs.adatum.com/federationmetadata/2007-06/federationmetadata.xml**
   - **Disable certificate chain validation**
   - **No encryption**

▶ Task 4: Configure a relying-party trust for the claims-aware application

1. On LON-DC1, in the AD FS console, add a relying-party trust with the following settings:
   - **Import data about the relying party published online or on a local network**
   - Federation Metadata address: **https://lon-svr1.adatum.com/adatumtestapp/**
   - Display name: **A. Datum Test App**
   - **I do not want to configure multi-factor authentication settings for this relying party trust at this time**
   - **Permit all users to access this relying party**
2. Leave the Edit Claim Rules for A. Datum Test App window open for the next task.

▶ Task 5: Configure claim rules for the relying-party trust

1. On LON-DC1, in the Edit Claim Rules for A. Datum Test App window, add a rule on the **Issuance Transform Rules** tab.
2. Complete the Add Transform Claim Rule Wizard with the following settings:
   - Claim rule template: **Pass Through or Filter an Incoming Claim**
   - Claim rule name: **Pass through Windows account name**
   - Incoming claim type: **Windows account name**
   - **Pass through all claim values**
3. Create three more rules to pass through the **E-Mail Address**, **UPN**, and **Name** claim types.

▶ Task 6: Test access to the claims-aware application

1. On LON-CL1, use Internet Explorer to access **https://lon-svr1.adatum.com/AdatumTestApp/**.

Note: It is critical to use the trailing slash in the URL for step 1.

2. Sign in as **Adatum\Brad** with the password **Pa$$w0rd**.
3. Review the claim information that is displayed by the application.
4. Close Internet Explorer.

▶ Task 7: Configure Internet Explorer to pass local credentials to the application automatically

1. On LON-CL1, from the Start screen, open Internet Options.
2. On the **Security** tab, add the following sites to the **Local intranet** zone:
   - **https://adfs.adatum.com**
   - **https://lon-svr1.adatum.com**
3. Use Internet Explorer to access **https://lon-svr1.adatum.com/AdatumTestApp/**.

Note: It is critical to use the trailing slash in the URL for step 3.

4. Notice that you were not prompted for credentials.
5. Review the claim information that is displayed by the application.
6. Close Internet Explorer.

Results: After completing this exercise, you will have configured AD FS to support authentication for an application.

Question: Why was it important to configure adfs.adatum.com to use as a host name for the AD FS service?

Question: How can you test whether AD FS is functioning properly?

Lesson 4

Deploying AD FS in a Business-to-Business Federation Scenario

A second common scenario for implementing AD FS is in a business-to-business federation. In this scenario, users in one organization require access to an application in another organization. AD FS enables SSO in this scenario. This way, users always log on to their home AD DS environment, but they are granted access to the partner application based on the claims acquired from their local AD FS server.

Configuring AD FS in a business-to-business federation scenario is similar to configuring AD FS in a single-organization scenario. The primary difference is that now, both the claims-provider trusts and the relying-party trusts refer to external organizations, rather than to internal AD DS or applications.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe how to configure an account partner for a business-to-business scenario.
- Describe how to configure a resource partner for a business-to-business scenario.
- Explain how to configure claim rules for a business-to-business scenario.
- Explain how home realm discovery works.
- Describe how to configure claims rules.

Configuring an Account Partner

In a business-to-business AD FS scenario, the terminology that you use to describe the two partners involved in an AD FS deployment changes slightly. In this scenario, the claims provider organization also is called the *account partner*. An account-partner organization is an organization in which user accounts are stored in an attribute store. An account partner handles the following tasks:

- An account partner is a claims provider in a business to business federation scenario
- To configure an account partner:
  1. Implement the physical topology
  2. Add an attribute store
  3. Configure a relying party trust
  4. Add a claim description
  5. Prepare client computers for federation

- Gathers credentials from users who are using a web-based service, and then authenticates those credentials.
- Builds up claims for users, and then packages the claims into security tokens. The tokens can then be presented across a federation trust to gain access to federation resources that are located at the resource partner's organization.

To configure the account-partner's organization to prepare for federation, use the following steps:

1. Implement the physical topology for the account-partner deployment. This step could include deciding on the number of federation servers and federation server proxies to deploy, and configuring the required DNS records and certificates.
2. Add an attribute store. Use the AD FS management console to add the attribute store. In most cases, you use the default Active Directory attribute store, which must be used for authentication, but you also can add other attribute stores, if required, to build the user claims. You connect to a resource-partner organization by creating a relying-party trust. The simplest way to do this is to use the

federation metadata URL that is provided by the resource-partner organization. With this option, your AD FS server automatically collects the information required for the relying-party trust.

3. Add a claim description. The claim description lists the claims that your organization provides to the relying partner. This information might include user names, email addresses, group membership information, or other identifying information about a user.
4. Prepare client computers for federation. This might involve two steps:
   - Add the account-partner federation server. In the browsers of client computers, add the account-partner federation server to the local intranet sites list. By adding the account-partner federation server to the local intranet list on client computers, you enable Integrated Windows authentication, which means that users are not prompted for authentication if they are logged on to the domain already. You can use Group Policy Objects (GPOs) to assign the URL to the local intranet site list.
   - Configure certificate trusts. This is an optional step that is required only if one or more of the servers that clients access do not have trusted certificates. The client computer might have to connect to the account-federation servers, resource-federation servers, or federation-server proxies, and the destination Web servers. If any of these certificates are not from a trusted public CA, you might have to add the appropriate certificate or root certificate to the certificate store on the clients. You can do this by using GPOs.

Configuring a Resource Partner

The *resource partner* is the relying party in a business-to-business federation scenario. The resource partner organization is where the resources exist and are made accessible to account-partner organizations. The resource partner handles the following tasks:

- Accepts security tokens that the account-partner federation server produces and then validates them.
- Consumes the claims from the security tokens and then provides new claims to its Web servers after making an authorization decision.

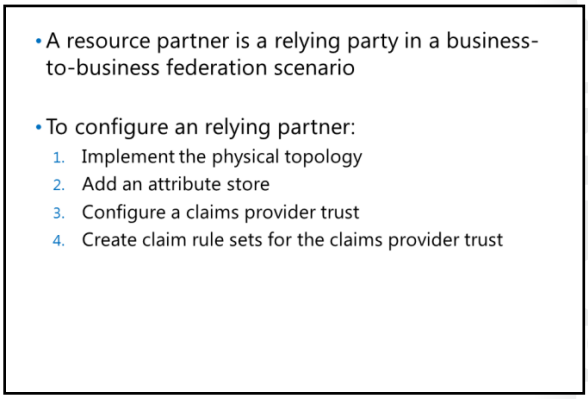

Web servers must have either WIF or the AD FS 1.x Claims-Aware Web Agent role services installed to externalize the identity logic and accept claims. WIF provides a set of development tools that enable developers to integrate claims-based authentication and authorization into their applications. WIF also includes a software development kit and sample applications.

Note: Applications on non-Microsoft Web servers can be integrated with AD FS by using SAML tokens. Additional open-source or third-party software is typically required to support the use of SAML tokens on a non-Microsoft Web server.

Configuring a resource-partner organization is similar to configuring an account-partner organization, and consists of the following steps:

1. Implement the physical topology for the resource-partner deployment. The planning and implementation steps are the same as for the account partner, with the addition of planning the Web server location and configuration.
2. Add an attribute store. On the resource partner, the attribute store is used to populate the claims that are offered to the client to present to the Web server.
3. Connect to an account-partner organization by creating a claims-provider trust.
4. Create claim rule sets for the claims-provider trust.

Configuring Claims Rules for Business-to-Business Scenarios

In a single-organization AD FS deployment, it might be simple to design and implement claims rules. In many cases, you might need to provide only the user or group name that is collected from the claim and presented to the Web server. In a business-to-business scenario, it is more likely that you will have to configure more complicated claims rules to define user access between widely different systems.

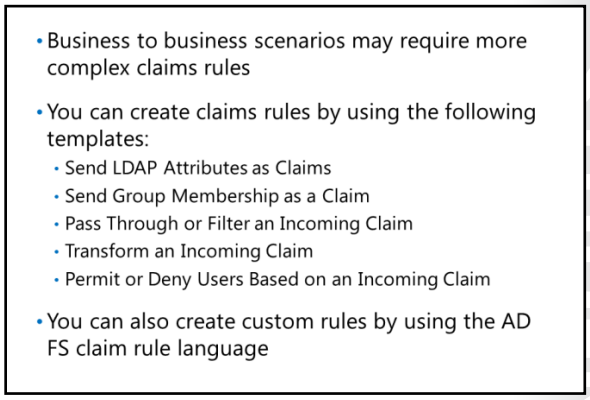

Claim rules define how account partners (claims providers) create claims, and how resource partners (relying parties) consume claims. AD FS provides several rule templates that you can use when you configure claim rules:

- Send LDAP Attributes as Claims. Use this template when you select specific attributes in an LDAP attribute store to populate claims. You can configure multiple LDAP attributes as individual claims in a single claim rule that you create from this template. For example, you can create a rule that extracts the **sn** (surname) and **givenName** AD DS attributes from all authenticated users, and then sends these values as outgoing claims to be sent to a relying party.
- Send Group Membership as a Claim. Use this template to send a particular claim type and an associated claim value that is based on the user's AD DS security group membership. For example, you might use this template to create a rule that sends a group claim type with a value of **SalesAdmin**, if the user is a member of the Sales Manager security group within their AD DS domain. This rule issues only a single claim based on the AD DS group that you select as a part of the template.
- Pass Through or Filter an Incoming Claim. Use this template to set additional restrictions on which claims are submitted to relying parties. For example, you might want to use a user email address as a claim, but only forward the email address if the domain suffix on the email address is adatum.com. When you use this template, you can either pass through whatever claim you extract from the attribute store, or you can configure rules that filter whether the claim is passed on based on various criteria.
- Transform an Incoming Claim. Use this template to map the value of an attribute in the claims-provider attribute store to a different value in the relying-party attribute store. For example, you might want to provide all members of the Marketing department at A. Datum limited access to a purchasing application at Trey Research. At Trey Research, the attribute used to define the limited access level might have an attribute of **LimitedPurchaser**. To address this scenario, you can configure

a claims rule that transforms an outgoing claim where the Department value is Marketing, to an incoming claim where the **ApplicationAccess** attribute is **LimitedPurchaser**. Rules created from this template must have a one-to-one relationship between the claim at the claims provider and the claim at the relying partner.

- Permit or Deny Users Based on an Incoming Claim. This template is available only when you configure issuance-authorization rules or delegation-authorization rules on a relying-party trust. Use this template to create rules that enable or deny access by users to a relying party, based on the type and value of an incoming claim. This claim rule template allows you to perform an authorization check on the claims provider before claims are sent to a relying party. For example, you can use this rule template to create a rule that only permits users from the Sales group to access a relying party, while authentication requests from members of other groups are not sent to the relying party.

If none of the built-in claim rule templates provides the functionality that you require, you can create more complex rules by using the AD FS claim-rule language. By creating a custom rule, you can extract claims information from multiple attribute stores and also combine claim types into a single claim rule.

How Home Realm Discovery Works

Some resource-partner organizations that host claims-aware applications might want to enable multiple account partners to access their applications. In this scenario, when users connect to the web application, there must be some mechanism for directing the users to the AD FS federation server in their home domain, rather than to another organization's federation server. The process for directing clients to the appropriate account partner is called *home realm discovery*.

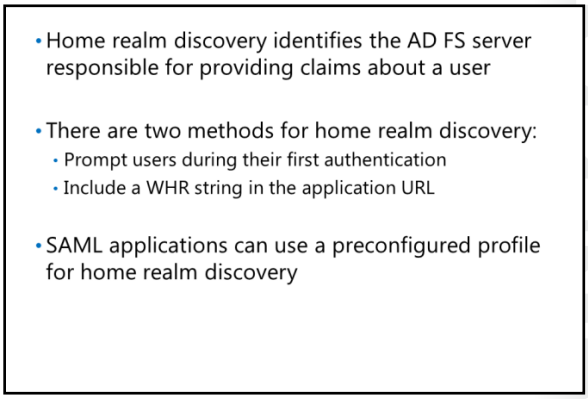

Home realm discovery occurs after the client connects to the relying party's website, and the client is redirected to the relying party's federation server. At this point, the relying party's federation server must redirect the client to the federation server in the client's home realm so that the user can authenticate. If there are multiple claims providers configured on the relying party's federation server, it has to know to which federation server to redirect the client.

In general, there are two ways to implement home realm discovery:

- Ask users to select their home realm. With this option, when users are redirected to the relying party's federation server, the federation server can display a web page that asks them to identify their company. Once users select the appropriate company, the federation server can use that information to redirect client computers to the appropriate home federation server for authentication.
- Modify the link for the web application to pass the WHR parameter that contains the user's home realm. The relying-party's federation server uses this parameter to redirect the user to the appropriate home realm automatically. This means that the user does not have to be prompted to select the home realm because the WHR parameter in the URL that the user clicks includes the needed information for the relying-party's federation server. The modified link might look something like the following: https://www.adatum.com/OrderApp/?whr=urn:federation:TreyResearch.

Note: One of the options available for home realm discovery with SAML 2.0-compliant applications is a SAML profile called IdPInitiated SSO. This SAML profile configures users to access

their local claims provider first, which can prepare the user's token with the claims required to access the partner's web application. The Windows Server 2012 version of AD FS does not implement the IdPInitiated SSO profile fully, but it provides some of the same functionality by implementing a feature named RelayState.

To learn more about the Supporting Identity Provider Initiated RelayState, go to:
http://go.microsoft.com/fwlink/?LinkId=269666

Note: The home realm discovery process occurs the first time a user tries to access a web application. After the user authenticates successfully, a home realm discovery cookie is issued to the client, so the user does not have to go through the process the next time. This home realm discovery cookie expires after a month, unless the cookie cache is cleared prior to expiration.

Demonstration: Configuring Claim Rules

In this demonstration, you will see how to configure claim rules on a relying-party trust that forwards a group name as part of the claim. You also will see how to configure a claims rule that limits access to an application for members of a particular group only.

Demonstration Steps

1. On LON-DC1, in the AD FS Manager, in the Edit Claim Rules for A. Datum Test App window, add an **Issuance Transform Rule** with the following settings:
   - Claim rule template: **Pass Through or Filter an Incoming Claim**
   - Claim rule name: **Send Group Name Rule**
   - Incoming claim type: **Group**
   - **Pass through all claim values**
2. Remove the **Permit Access to All Users** issuance authorization rule.
3. Add a new issuance authorization rule with the following settings:
   - Claim rule template: **Permit or Deny Users Based on an Incoming Claim**
   - Claim rule name: **Permit Production Group Rule**
   - Incoming claim type: **Group**
   - Incoming claim value: **Production**
   - **Permit access to users with this incoming claim**
4. Add a new issuance authorization rule with the following settings:
   - Claim rule template: **Permit or Deny Users Based on an Incoming Claim**
   - Claim rule name: **Allow A. Datum Users**
   - Incoming claim type: **UPN**
   - Incoming claim value: **@adatum.com**
   - **Permit access to users with this incoming claim**
5. View the rule language for the **Allow A. Datum Users** rule.

Lesson 5
Extending AD FS to External Clients

Many organization need to extend the AD FS infrastructure beyond private networks and onto the Internet. To enhance security for AD FS and AD FS applications, you can use Web Application Proxy. It is important to consider high availability for AD FS because it is a critical service. Finally, Workplace Join provides a way to regulate access to applications based on registration of non-domain joined devices.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe Web Application Proxy.
- Describe how to configure an application for Web Application Proxy.
- Describe Web Application Proxy and AD FS.
- Describe high availability for AD FS with Web Application Proxy.
- Install and configure Web Application Proxy.
- Describe Workplace Join.
- Describe the Workplace Join process.

What Is Web Application Proxy?

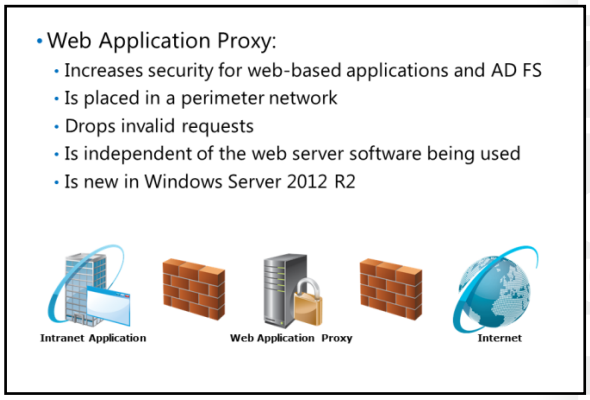

Web Application Proxy in Windows Server 2012 R2 is a role service in the remote access role. You can use it to secure remote access to web-based applications on your internal network. It functions as a reverse proxy for web-based applications. It also functions as an AD FS proxy.

You should place Web Application Proxy in a perimeter network. External clients that access web-based applications or AD FS initiate connections with Web Application Proxy. Web Application Proxy then connects to the web-based application or AD FS on the internal network. No client-specific configuration is required to use Web Application Proxy.

When you implement Web Application Proxy, you enhance security for web-based applications or AD FS by isolating them from direct contact with the Internet. This can help protect the internal, web-based application or AD FS from any malformed packets or requests that might result in a security breach. For example, Web Application Proxy can protect against a zero-day vulnerability that uses malformed requests, which could result in a denial-of-service attack on a server that hosts a web-based application. Web Application Proxy will drop invalid requests before they reach the web-based application on an internal network.

Web Application Proxy is completely independent of the Web server software being used. Because of this, it is unlikely that Web Application Proxy is vulnerable to the same denial of service attack as a web-based application.

Web Application Proxy is not available in Windows Server 2012. Windows Server 2012 included an AD FS proxy option that could be installed as part of deploying AD FS. This option did not provide reverse proxy functionality for web-based applications. It was a reverse proxy only for AD FS.

Configuring an Application for Web Application Proxy

Web Application Proxy is used to protect web applications and AD FS when they are accessible from the Internet. You should place the Web Application Proxy server in a perimeter network. To install Web Application Proxy, AD FS must be implemented in your organization already. All configuration information for Web Application Proxy is stored in AD FS.

When you use Web Application Proxy as a reverse proxy for web applications, you need to configure each application. For each application, you need to configure the type of preauthentication for the application and URLs.

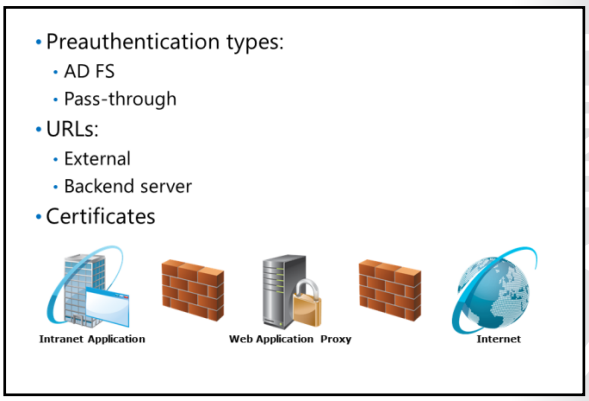

Pass-Through Preauthentication

When you use pass-through preauthentication, no preauthentication is performed and valid requests are passed to web-based applications on an internal network without performing authentication on a user. All authentication for an application is performed by the application only after a user is connected. You can use pass-through preauthentication for any web application.

A web application protected by pre-authentication is protected from malformed packets that could cause a denial-of-service attack. However, the web application would not be protected from application-level threats where the application mishandles valid data. For example, an HTTPS request with valid HTTP commands would be passed through to the application even if the actions requested by the HTTP commands may cause the web application to fail.

AD FS Preauthentication

You can configure Web Application Proxy to use AD FS preauthentication or pass-through authentication. When you use AD FS for preauthentication, a user request is authenticated by AD FS before it is passes to an internal, web-based application. This ensures that only authorized users can send data to a web-based application. AD FS preauthentication provides a higher level of protection than pass-through authentication because unauthenticated users cannot submit requests to the application.

Only a claims-aware application that uses AD FS for authentication can use AD FS preauthentication. The claims-aware application must be configured in AD FS as a relying party and is selected from a list when Web Application Proxy is configured. Web Application Proxy is aware of the relying parties configured in AD FS because of the integration between AD FS and Web Application Proxy.

URLs

For each application that you publish, you must configure an external URL and backend server URL. The external URL is used by external users when accessing the application. The back-end server URL is used by the Web Application Proxy server to access the application on behalf of external users.

If you are using split DNS, it is possible to leave the external URL and the back-end server URL as the same value. Some applications experience errors when the external URL and the back-end server URL are different. When the external URL and the back-end server URL are different, only the host name in the

URL can change. The path to the application must remain the same. For example, if the back-end URL for an application is https://server1.adatum.com/app1, then you cannot have an external URL of https://extranet.adatum.com/application1.

Certificates

When you define the external URL, you also need to select a certificate that contains the host name in the external URL. This certificate must be installed on the local server. However, it does not need to match the certificate used on the back-end server that hosts the application. You can have one certificate for each host name used on the Web Application Proxy server, or a single certificate with multiple names.

Web Application Proxy and AD FS

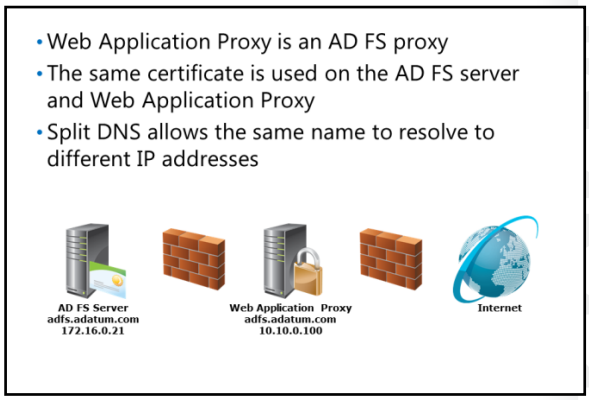

Many organizations need to provide authentication for users and devices that are located on a network that is external to the organization. In most cases, allowing clients to access an AD FS server located on an internal network directly from the Internet is an unacceptable security risk. To allow clients on the Internet to access AD FS, we strongly recommend an AD FS proxy.

An AD FS proxy is a reverse proxy, located in a perimeter network that is specifically for AD FS. Clients from the Internet communicate with the AD FS proxy in the perimeter network instead of directly with the AD FS server. The AD FS proxy mitigates the risks associated with Internet connectivity for AD FS.

In Windows Server 2012, you can install an AD FS proxy as part of an AD FS installation. In Windows Server 2012 R2, you can configure Web Application Proxy as an AD FS proxy.

Authentication Process

An internal AD FS server uses Windows authentication to prompt for authentication. This works well for internal, domain-joined computers that can pass workstation credentials automatically to AD FS and automate authentication. This prevents users from seeing a request for authentication credentials.

When computers that are not domain-joined communicate with AD FS, users encounter a logon prompt that is presented by the web browser. This logon prompt asks for a user name and password, but provides no context.

When you use an AD FS proxy, an authentication web page is provided for computers that are not domain-joined. This provides better compatibility than browser-based Windows authentication for AD FS clients that use non-Microsoft operating systems. You also can customize the web page to provide more context for users, such as a company logo.

DNS Resolution

To provide seamless movement between internal and external networks, the same host name is used when AD FS is accessed internally and externally. On the internal network, the AD FS host name resolves to the IP address of the internal AD FS server. On the external network, the AD FS host name resolves to the IP address of the AD FS proxy. In both cases, the AD FS host name is different than the host name of the computers that host the AD FS roles.

Certificates

The certificate an internal AD FS server uses has a subject name that is the same as the host name for AD FS—for example, adfs.adatum.com. Because the same host name is used to access AD FS internally and externally through the AD FS proxy, you must configure the AD FS proxy with the same certificate as the AD FS server. If the certificate subject does not match the host name, AD FS authentication will fail.

Note: Export the certificate from the AD FS server an import it on the Web Application Proxy server to ensure that you have a certificate with the same subject name. Remember to include the private key when you export the certificate.

High Availability for AD FS

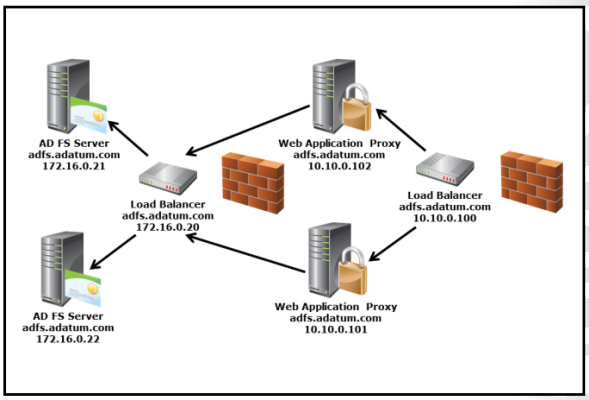

If you use AD FS to provide authentication for applications, it is likely that AD FS is a critical service for your organization. If AD FS is unavailable, access to those applications is lost until AD FS functionality is restored. This can have a critical impact on your business.

The impact will vary, depending on which part of your AD FS infrastructure fails. The loss of the AD FS server will prevent authentication for all applications that use AD FS. The loss of the AD FS proxy would affect external clients on the Internet and partner organizations that access AD FS over the internet. For example, if the AD FS proxy fails, any external applications hosted at a partner site that authenticates by using AD FS will fail because the communication between the AD FS servers in your organizations is done through the AD FS proxy.

AD FS Servers

To provide high availability for AD FS servers, you add an additional AD FS server to the farm. If you selected to use the Windows Internal Database (WID) for the primary AD FS server, the secondary AD FS server also uses WID, and replicates the configuration database from the primary AD FS server.

In scenarios where you require more than five AD FS servers for scalability, you should use a Microsoft SQL Server database as the store for the configuration database rather than WID. All AD FS servers in the farm access a single configuration database. If you use a SQL Server database as the store for the configuration database, then you must configure high availability for the SQL Server database to ensure high availability for AD FS.

You make client access to AD FS highly available by using load balancing. You can use the Network Load Balancing (NLB) feature or hardware load balancing. The federation service name, such as adfs.adatum.com, is configured as the load-balanced host name and IP address. Clients communicate with the load-balanced IP address, and requests are load-balanced across the AD FS servers.

If you use hardware-based load balancers, then the hardware-based load balancer needs to be highly available. This typically is done by using two hardware-based load balancers as a team. The NLB feature is automatically highly available, but it is not as scalable as hardware-based load balancers. NLB also monitors the availability of servers, but not services. Most hardware-based load balancers can monitor the availability of specific services that run on servers.

All AD FS servers in the farm must have the same certificate to support the use of the shared federation service name. Depending on the configuration, you also might need to install the certificate on hardware-based load balancers.

AD FS Proxy Servers

AD FS proxy servers do not contain any configuration information locally. They store all configuration information in the farm's configuration database. This simplifies high availability of AD FS proxy servers when compared to AD FS server. In other respects, the configuration of high availability for AD FS proxy servers is similar to AD FS servers.

To make AD FS proxies highly available, you need to install multiple AD FS proxy servers and provide load balancing between them. The host name for the AD FS proxy should be configured to resolve to the IP address of the load-balanced service. The AD FS proxy servers must be configured to use the same certificate.

Geographic High Availability

It is possible to design AD FS to survive the loss of an entire physical location. To do this, you must have at least one AD FS server at a second location. You also must replicate the configuration database to the second location. If the first location becomes unavailable, the AD FS server in the second location can use the replicated configuration database in the second location to remain functional and to begin servicing clients.

Demonstration: Installing and Configuring Web Application Proxy

In this demonstration, you will see how to install and configure Web Application Proxy. This includes exporting the certificate from the AD FS server and importing it on the Web Application Proxy server.

Demonstration Steps

Install Web Application Proxy

- On LON-SVR2, in the Server Manager, add the remote access server role and the Web Application Proxy role service.

Export the adfs.adatum.com certificate from LON-DC1

1. On LON-DC1, open a Microsoft Management Console, and then add the **Certificates** snap-in for the **Local Computer**.
2. From the Personal folder, export the **adfs.adatum.com** certificate:
    - **Yes, export the private key**
    - File format: **Personal Information Exchange – PKCS #12 (.PFX)**
    - Password: **Pa$$w0rd**
    - File name: **C:\adfs.pfx**

Import the adfs.adatum.com certificate on LON-SVR2

1. On LON-SVR2, open a Microsoft Management Console, and then add the **Certificates** snap-in for the **Local Computer**.

2. From the Personal folder, import the **adfs.adatum.com** certificate.
   - File name: **\\LON-DC1\c$\adfs.pfx**
   - Password: **Pa$$w0rd**
   - Certificate store: **Personal**

Configure Web Application Proxy

1. On LON-SVR2, in the Server Manager, click the **Notifications** icon, and then click **Open the Web Application Proxy Wizard**.
2. In the Web Application Proxy Wizard, provide the following configuration settings:
   - Federation service name: **adfs.adatum.com**
   - User name: **Adatum\Administrator**
   - Password: **Pa$$w0rd**
   - Certificate to be used by the AD FS proxy: **adfs.adatum.com**

What Is Workplace Join?

Workplace Join is a solution, included in Windows Server 2012 R2, which allows you to control access to company resources from non-domain joined computers and devices. This type of control is important because of the trend to allow users to "bring your own devices" (BYOD) and access organizational resources. When you use Workplace Join, you can control which user and device combinations are allowed to access company resources.

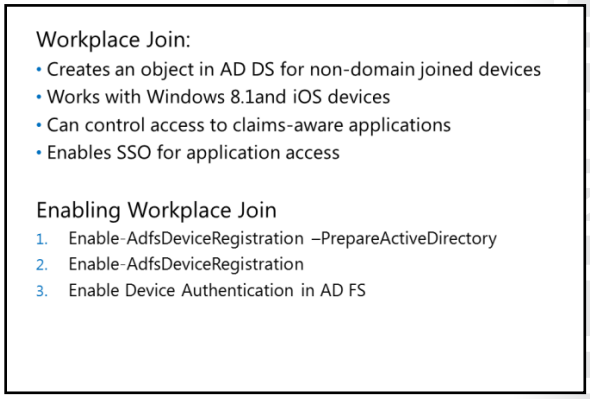

When Workplace Join has been completed for a device, that device is registered. Registration results in a device object being created in AD DS to represent the device. Information about the device can then be used as part of the authentication process for company resources. For example, you can require that only registered devices are allowed to access web-based application. Specific attributes of the device can also be used to control access to the application.

Supported Clients

The only Windows client that supports Workplace Join is Windows 8.1. You cannot use earlier versions of Windows clients for Workplace Join. However, Workplace Join is cross platform, and it supports iOS devices such as iPads and iPhones. Support for Android devices is planned.

Supported Applications

Only claims-aware applications that use AD FS can use device registration information. Device information is provided to the claims aware application by AD FS as part of the authentication process.

Single Sign-on

When you use a workplace-joined device, you have single sign-on (SSO) for your enterprise applications. After you authenticate once to an application, you are not prompted for authentication credentials the second time.

Enabling Workplace Join

You need to enable Workplace Join before any devices can be registered. You need to perform the following steps to enable Workplace Join:

1. Run the following Windows Powershell command:

```
Enable-AdfsDeviceRegistration –PrepareActiveDirectory
```

2. Run the following Windows PowerShell command and provide the name of a service account such as Adatum\ADFS$ when prompted:

```
Enable-AdfsDeviceRegistration
```

3. In the AD FS management console, in the Global Authentication Policy, select the **Enable Device Authentication** check box.

Workplace Join is automatically supported through Web Application Proxy after you have performed these configuration steps.

The Workplace Join Process

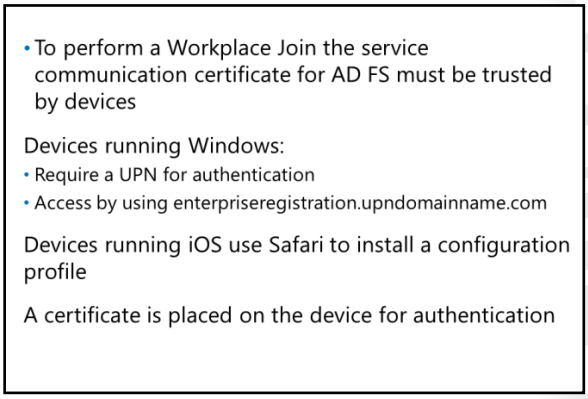

Regardless of the client type, the service communication certificate configured for AD FS must be trusted by the client. Since devices that perform a Workplace Join are not already managed by the organization, you should use a certificate from a trusted third-party certification authority. This avoids the need to configure each device to trust your internal certification authority.

The Workplace Join process requires clients to perform a certificate revocation check on the certificate used by the AD FS server or Web Application Proxy with which they are communicating. If the certificate revocation check fails, then the Workplace Join will also fail. Using a third-party certification authority avoids the need configure a certificate revocation list distribution point for your internal certification authority that is accessible from the Internet.

Workplace Join for Windows Devices

During Workplace Join, you are prompted to provide your email address and password. The required information is actually your UPN and not your email address. To simplify this process, we strongly recommend that the UPN for users match their email address.

Windows devices automatically locate the server for Workplace Join based on the UPN that is provided. The server use for Workplace join is enterpriseregistration.*upndomainname.com*. You need to configure DNS to properly resolve this record to the IP address of your AD FS server or Web Application Proxy that is configured to support Workplace Join.

The certificate for the AD FS server and AD FS proxy functionality of Web Application Proxy need to include the enterpriseregistration.*upndomainname.com* domain name. The configuration process is simpler if you include this name in the certificate used during the installation of AD FS and Web Application Proxy, instead of changing the certificate after installation.

Workplace Join for iOS Devices

To perform a Workplace Join for an iOS device, you need to set up a configure profile on the iOS device. An iOS configuration profile is created by providing an XML file. For a Workplace Join, the XML file is delivered by a web site. This is referred to as over-the air profile delivery.

The website iOS devices use to download the configuration profile is located on the AD FS server where the Device Registration Service is enabled. An example of the URL used to configure an iOS device is https://adfs.contoso.com/enrollmentserver/otaprofile.

On the website, you are prompted to sign in by using your email address as a username. Like the process for devices running windows, your UPN should be entered rather than your email address. After you sign in, you install the profile on the iOS device. If the iOS device requires that a PIN be entered to unlock the device, you are prompted to enter the PIN before the profile is installed.

Certificates on Devices

The Workplace Join process places a certificate on the device. This certificate is used by the device to prove its identity. This certificate is used to authenticate to the object created for the device in AD DS.

Lab B: Implementing AD FS for External Partners and Users

Scenario

A. Datum Corporation has set up a variety of business relationships with other companies and customers. Some of these partner companies and customers must access business applications that are running on the A. Datum network. The business groups at A. Datum want to provide a maximum level of functionality and access to these companies. The Security and Operations departments want to ensure that the partners and customers can access only the resources to which they require access, and that implementing the solution does not increase the workload for the Operations team significantly. A. Datum also plans to migrate some parts of its network infrastructure to Microsoft Online Services, including Windows Azure and Office 365.

Now that you have deployed AD FS for internal users, the next step is to enable access to the same application for external partner organizations and for external users. A. Datum Corporation has entered into a partnership with Trey Research. You need to ensure that Trey Research users can access the internal application. You also need to ensure that A. Datum Corporation users working outside the office can access the application.

As one of the senior network administrators at A. Datum, it is your responsibility to implement the AD FS solution. As a proof-of-concept, you are deploying a sample claims-aware application, and configuring AD FS to enable both Trey Research users and external A. Datum Corporation users to access the same application.

Objectives

After completing this lab, you will be able to:

- Configure AD FS for a federated partner.
- Configure Web Application Proxy for external users.

Lab Setup

Estimated Time: 60 minutes

20412C-LON-DC1

20412C-LON-SVR1

20412C-LON-SVR2

20412C-TREY-DC1

User name: **Adatum\Administrator**

Password: **Pa$$w0rd**

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

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, click **20412C-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 by using the following credentials:
    - User name: **Adatum\Administrator**
    - Password: **Pa$$w0rd**

5. Repeat steps 2 to 4 for 20412C-LON-SVR1, 20413C-LON-SVR2, and 20412C-TREY-DC1.
    - For TREY-DC1, sign in as **TREYRESEARCH\Administrator** with a password of **Pa$$w0rd**.

Exercise 1: Configuring AD FS for a Federated Business Partner

Scenario

The second deployment scenario is to enable Trey Research users to access the web application. You plan to configure the integration of AD FS at Trey Research with AD FS at A. Datum, and then verify that Trey Research users can access the application. You also want to confirm that you can configure access that is based on user groups. You must ensure that all users at A. Datum, and only users who are in the Production group at Trey Research, can access the application.

The main tasks for this exercise are as follows:

- Configure DNS forwarding between TreyResearch.net and Adatum.com.
- Configure certificate trusts between TreyResearch.net and Adatum.com.
- Create a DNS record for AD FS in TreyResearch.net.
- Create a certificate for AD FS.
- Create a service account.
- Install AD FS for TreyResearch.net.
- Configure AD FS for TreyResearch.net.
- Add a claims-provider trust for the TreyResearch.net AD FS server.
- Configure a relying party trust in TreyResearch.net for the Adatum.com application.
- Test access to the application.
- Configure issuance authorization rules.
- Test the application of issuance authorization rules.

▶ Task 1: Configure DNS forwarding between TreyResearch.net and Adatum.com

1. On LON-DC1, use the DNS Manager to create a new conditional forwarder with the following settings:
    - DNS Domain: **TreyResearch.net**
    - IP address of the master server: **172.16.10.10**
    - **Store this conditional forwarder in Active Directory and replicate it as follows: All DNS servers in this forest**
2. On TREY-DC1, use the DNS Manager to create a new conditional forwarder with the following settings:
    - DNS Domain: **Adatum.com**
    - IP address of the master server: **172.16.0.10**
    - **Store this conditional forwarder in Active Directory and replicate it as follows: All DNS servers in this forest**

 **Note:** In a production environment, it is likely that you would use Internet DNS instead of conditional forwarders.

▶ Task 2: Configure certificate trusts between TreyResearch.net and Adatum.com

1. On LON-DC1, use File Explorer to copy **TREY-DC1.TreyResearch.net_TreyResearchCA.crt** from **\\TREY-DC1\CertEnroll** to **C:**.
2. Open Group Policy Management, and then edit the **Default Domain Policy**.
3. In the Default Domain Policy, browse to **Computer Configuration\Policies\Windows Settings\Security Settings\Public Key Policies\Trusted Root Certification Authorities**
4. Import **C:\TREY-DC1.TreyResearch.net_TreyResearchCA.crt** as a trusted root CA.
5. On TREY-DC1, use File Explorer to browse to **\\LON-DC1\CertEnroll**.
6. Right-click **LON-DC1.Adatum.com_AdatumCA.crt**, and then install the certificate into the **Trusted Root Certification Authority** store.
7. On LON-SVR1, run Gpupdate.
8. On LON-SVR2, run Gpupdate.

 **Note:** If you obtain certificates from a trusted certification authority, you do not need to configure a certificate trust between the organizations.

▶ Task 3: Create a DNS record for AD FS in TreyResearch.net

1. On TREY-DC1, use the DNS Manager to add a new host record for AD FS:
   - Forward lookup zone: **TreyResearch.net**
   - Name: **adfs**
   - IP address: **172.16.10.10**

▶ Task 4: Create a certificate for AD FS

1. On TREY-DC1, open Internet Information Services (IIS) Manager and view the server certificates.
2. Create a new domain certificate with the following settings:
   - Common name: **adfs.TreyResearch.net**
   - Organization: **Trey Research**
   - Organizational unit: **IT**
   - City/locality: **London**
   - State/Province: **England**
   - Country/region: **GB**
   - Certification Authority: **TreyResearchCA**
   - Friendly name: **adfs.TreyResearch.net**

▶ Task 5: Create a service account

1. On TREY-DC1, open a Windows PowerShell prompt.
2. Create a new user account:
   - **New-ADUser –Name adfsService**

3. Set a password for adfsService:
   - **Set-ADAccountPassword adfsService**
   - Current password: none (press Enter)
   - Desired password: **Pa$$w0rd**
4. Enable the adfsService account:
   - **Enable-ADAccount adfsService**

▶ Task 6: Install AD FS for TreyResearch.net

- On TREY-DC1, in the Server Manager, add the Active Directory Federation Services role.

▶ Task 7: Configure AD FS for TreyResearch.net

1. In the Server Manager notifications, click **Configure the federation services on this server**.
2. Use the following options to configure the AD FS server:
   - **Create the first federation server in a federation server farm**
   - Account for configuration: **TREYRESEARCH\Administrator**
   - SSL Certificate: **adfs.TreyResearch.net**
   - Federation Service Display Name: **Trey Research**
   - Use an existing domain user account or group Managed Service Account:
     - **TREYRESEARCH\adfsService**
     - Password: **Pa$$w0rd**
   - **Create a database on this server using Windows Internal Database**

▶ Task 8: Add a claims-provider trust for the TreyResearch.net AD FS server

1. On LON-DC1, use the AD FS management console to add a new claims provider trust with the following settings:
   - **Import data about the claims provider published online or on a local network**
   - Federation metadata address: **https://adfs.treyresearch.net**
   - Display name: **Trey Research**
   - **Open the Edit Claim Rules dialog for this claims provider trust when the wizard closes**
2. Create a claim rule for Trey Research by using the following settings:
   - Claim rule template: **Pass Through or Filter an Incoming Claim**
   - Claim rule name: **Pass through Windows account name**
   - Incoming claim type: **Windows account name**
   - **Pass through all claim values**

▶ Task 9: Configure a relying party trust in TreyResearch.net for the Adatum.com application

1. On TREY-DC1, use the AD FS management console to create a new relying-party trust with the following settings:
   - **Import data about the relying party published online or on a local network**

- Federation metadata address: **adfs.adatum.com**
- Display name: **A. Datum Corporation**
- **I do not want to configure multi-factor authentication settings for this relying party trust at this time**
- **Permit all users to access this relying party**
- **Open the Edit Claim Rules dialog box for the relying party trust when the wizard closes**

2. Create a new transform claim rule with the following settings:
   - Claim rule template: **Pass Through or Filter an Incoming Claim**
   - Claim rule name: **Pass through Windows account name**
   - Incoming claim type: **Windows account name**
   - **Pass through all claim values**

▶ Task 10: Test access to the application

1. On TREY-DC1, use Internet Explorer to access https://lon-svr1.adatum.com/adatumtestapp/
2. Select the Trey Research home realm, and then sign in as **TreyResearch\April** with the password **Pa$$w0rd**.
3. Verify that you can access the application.
4. Close Internet Explorer, and then connect to the same website. Verify that you are not prompted for a home realm this time.

Note: You are not prompted for a home realm on the second access. Once users have selected a home realm and have been authenticated by a realm authority, they are issued an _LSRealm cookie by the relying party's federation server. The default lifetime for the cookie is 30 days. Therefore, to sign in multiple times, you should delete that cookie after each logon attempt to return to a clean state.

▶ Task 11: Configure issuance authorization rules

1. On TREY-DC1, in the AD FS management console, remove the issuance authorization rule from the A. Datum Corporation relying party trust that permits access for all users.
2. Add an issuance authorization rule to the A. Datum Corporation relying party trust that allows all users that are members of the Production group:
   - Claim rule template: **Permit or Deny Users Based on an Incoming Claim**
   - Claim rule name: **Allow Production Members**
   - Incoming claim type: **Group**
   - Incoming claim value: **TreyResearch-Production**
   - **Permit access to users with the incoming claim**
3. Add a transform claim rule to the Active Directory claims provider trust to send group membership as a claim:
   - Claim rule template: **Send Group Membership as a Claim**
   - Claim rule name: **Production Group Claim**
   - User's group: **Production**

- Outgoing claim type: **Group**
- Outgoing claim value: **TreyResearch-Production**

▶ Task 12: Test the application of issuance authorization rules

1. On TREY-DC1, use Internet Explorer to access https://lon-svr1.adatum.com/adatumtestapp/
2. Sign in as **TreyResearch\April** with the password **Pa$$w0rd**.
3. Verify that you cannot access the application because April is not a member of the production group.
4. Close Internet Explorer, and then connect to the same website.
5. Sign in as **TreyResearch\Ben** with the password **Pa$$w0rd**.
6. Verify that you can access the application because April is a member of the production group.

Results: After completing this exercise, you will have configured access for a claims-aware application in a partner organization.

Exercise 2: Configuring Web Application Proxy

Scenario

The third scenario for implementing the proof-of-concept AD FS application is to increase security for AD FS authentication by implementing an AD FS proxy for the AD FS and a reverse proxy for the application. You will implement Web Application Proxy to fulfill both of these roles.

The main tasks for this exercise are as follows:

- Install Web Application Proxy.
- Add the adfs.adatum.com certificate to LON-SVR2.
- Add the LON-SVR1.adatum.com certificate to LON-SVR2.
- Configure Web Application Proxy.
- Configure the test application in Web Application Proxy.
- Test Web Application Proxy.
- To prepare for the next module.

▶ Task 1: Install Web Application Proxy

- On LON-SVR2, in the Server Manager, add the **Remote Access** server role and the **Web Application Proxy** role service.

▶ Task 2: Add the adfs.adatum.com certificate to LON-SVR2

1. On LON-DC1, open the Microsoft Management Console, and then add the **Certificates** snap-in for the **Local Computer**.
2. From the Personal folder, export the **adfs.adatum.com** certificate:
   - **Yes, export the private key**
   - File format: **Personal Information Exchange – PKCS #12 (.PFX)**
   - Password: **Pa$$w0rd**
   - File name: **C:\adfs.pfx**

3. On LON-SVR2, open a Microsoft Management Console, and then add the **Certificates** snap-in for the **Local Computer**.
4. From the Personal folder, import the **adfs.adatum.com** certificate:
   - File name: **\\LON-SVR2\c$\adfs.pfx**
   - Password: **Pa$$w0rd**
   - **Mark this key as exportable**
   - Certificate store: **Personal**

▶ Task 3: Add the LON-SVR1.adatum.com certificate to LON-SVR2

1. On LON-SVR1, open the Microsoft Management Console, and then add the **Certificates** snap-in for the **Local Computer**.
2. From the Personal folder, export the **adfs.adatum.com** certificate:
   - **Yes, export the private key**
   - File format: **Personal Information Exchange – PKCS #12 (.PFX)**
   - Password: **Pa$$w0rd**
   - File name: **C:\lon-svr1.pfx**
3. On LON-SVR2, open the Microsoft Management Console, and then add the **Certificates** snap-in for the **Local Computer**.
4. From the Personal folder, import the **adfs.adatum.com** certificate:
   - File name: **\\LON-SVR1\c$\lon-svr1.pfx**
   - Password: **Pa$$w0rd**
   - **Mark this key as exportable**
   - Certificate store: **Personal**

▶ Task 4: Configure Web Application Proxy

1. In the Server Manager, click the **Notifications** icon, and then click **Open the Web Application Proxy Wizard**.
2. In the Web Application Proxy Wizard, provide the following configuration settings:
   - Federation service name: **adfs.adatum.com**
   - User name: **Adatum\Administrator**
   - Password: **Pa$$w0rd**
   - Certificate to be used by the AD FS proxy: **adfs.adatum.com**
3. Leave the Remote Access Management Console open for the next exercise.

▶ Task 5: Configure the test application in Web Application Proxy

- On LON-SVR2, in the Remote Access Management Console, publish a new application with the following settings:
  - Preauthentication: **Pass-through**
  - Name: **A. Datum Test Application**
  - External URL: **http://lon-svr1.adatum.com/adatumtestapp/**

- External certificate: **lon-svr1.adatum.com**
- Back-end server URL: **https://lon-svr1.adatum.com/adatumtestapp/**

▶ Task 6: Test Web Application Proxy

1. On TREY-DC1, open Notepad as Administrator to add the following lines to **C:\Windows\System32\Drivers\etc\hosts**:
   - **172.16.0.22 adfs.adatum.com**
   - **172.16.0.22 lon-svr1.adatum.com**
2. Use Internet Explorer to access https://lon-svr1.adatum.com/adatumtestapp/.
3. Sign in as **TreyResearch\Ben** with the password **Pa$$w0rd**.

Note: You edit the hosts to force TREY-DC1 to access the application through Web Application Proxy. In a production environment, you would do this by using split DNS.

▶ Task 7: To prepare for the next module

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

1. On the host computer, start the Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20412C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 to revert **20412C-LON-SVR1, 20412C-LON-SVR2**, **20412C-LON-CL1** and **20412C-TREY-DC1**.

Results: After completing this exercise, you will have configured Web Application Proxy to secure access to AdatumTestApp from the Internet.

Question: Why would the need to configure certificate trusts between organizations be avoided when you use certificates from a trusted provider on the Internet?

Question: Could you have created authorization rules in Adatum.com and achieved the same result if you had instead created authorization rules in TreyResearch.net?

Module Review and Takeaways

Review Questions

Question: Your organization is planning to implement AD FS. In the short term, only internal clients will be using AD FS to access internal applications. However, in the long run, you will be providing access to web-based applications that are secured by AD FS to users at home. How many certificates should you obtain from a third-party CA?

Question: Your organization has an application for customers that allows them to view their orders and invoices. At the present time, all customers have a user name and password that is managed within the application. To simplify access to the application and reduce support calls, your organization has rewritten the application to support AD FS for authentication. What do you need to configure to support the application?

Question: Your organization has an application for customers that allows them to view their orders and invoices. At the present time, all customers have a user name and password that is managed within the application. To simplify access to the application and reduce support calls, your organization has rewritten the application to support AD FS for authentication. A Web Application Proxy is being configured to support application access over the Internet. Internally, your AD FS server uses the host name adfs.contoso.com and resolves to 10.10.0.99. How will you allow external partners to resolve adfs.contso.com to the external IP address of Web Application Proxy?

Question: Your organization has implemented a single AD FS server and a single Web Application Proxy successfully. Initially, AD FS was used for only a single application, but now it is being used for several business-critical applications. AD FS must be configured to be highly available.

During the installation of AD FS, you selected to use the Windows Internal Database. Can you use this database in a highly available configuration?

Question: Your organization wants to control access to applications that are available from the Internet by using Workplace Join. What DNS changes need to be performed so that devices can locate the Web Application Proxy during the Workplace Join process?

Module 9

Implementing Network Load Balancing

Contents:

Module Overview

Network Load Balancing (NLB) is a feature available to computers that run the Windows Server operating system. NLB uses a distributed algorithm to balance an IP traffic load across multiple hosts. It helps to improve the scalability and availability of business-critical, IP-based services. NLB also provides high availability, because it detects host failures and automatically redistributes traffic to surviving hosts.

To deploy NLB effectively, you must understand its functionality and the scenarios where its deployment is appropriate. The main update to NLB in Windows Server® 2012 and Windows Server® 2012 R2 compared to Windows Server® 2008 R2 is the inclusion of a comprehensive set of Windows PowerShell® cmdlets. These cmdlets enhance your ability to automate the management of Windows Server 2012 NLB and Windows Server 2012 R2 clusters. The Network Load Balancing console, which is also available in Windows Server® 2008 and Windows Server 2008 R2, is also present in Windows Server 2012 and Windows Server 2012 R2.

This module introduces you to NLB, and shows you how to deploy this technology. This module also discusses the situations for which NLB is appropriate, how to configure and manage NLB clusters, and how to perform maintenance tasks on NLB clusters.

Objectives

After completing this module, you will be able to:

- Describe NLB.
- Explain how to configure an NLB cluster.
- Explain how to plan an NLB implementation.

Lesson 1
Overview of NLB

Before you deploy NLB, you need to have a firm understanding of the types of server workloads for which this high availability technology is appropriate. If you do not understand the functionality of NLB, it is possible that you will deploy it in a manner that does not accomplish your overall objectives. For example, you need to understand why NLB is appropriate for web applications, but not for Microsoft® SQL Server databases.

This lesson provides an overview of NLB, and the features new to NLB in Windows Server 2012 and Windows Server 2012 R2. It also describes how NLB works normally, and how it works during server failure and server recovery.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe NLB technology.
- Describe how NLB works.
- Explain how NLB accommodates server failures and recovery.
- Describe new NLB features in Windows Server 2012 and Windows Server 2012 R2.

What Is NLB?

NLB is a scalable, high-availability feature that you can install on all editions of Windows Server 2012 and Windows Server 2012 R2. A *scalable technology* is one that enables you to add additional components, such as additional cluster nodes in this case, to meet an increasing demand. A *node* in a Windows Server 2012 or Windows Server 2012 R2 NLB cluster is a computer, either physical or virtual, that is running the Windows Server 2012 or the Windows Server 2012 R2 operating system.

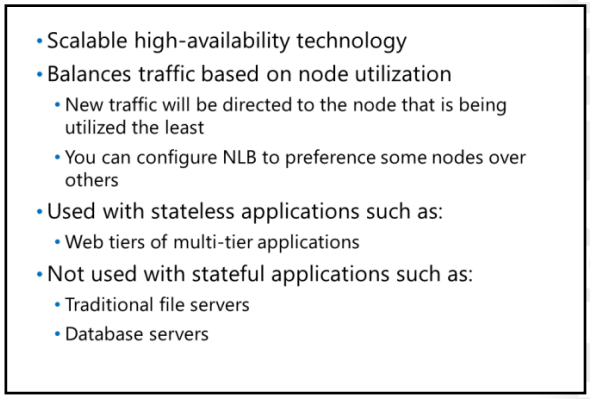

Windows Server 2012 and Windows Server 2012 R2 NLB clusters can have between two and 32 nodes. When you create an NLB cluster, it creates a virtual network address and virtual network adapter. The virtual network adapter has an IP address and a media access control (MAC) address. Network traffic to this address is distributed evenly across the nodes in the cluster. In a basic NLB configuration, each node in an NLB cluster will service requests at a rate that is approximately equal to that of all other nodes in the cluster. When an NLB cluster receives a request, it will forward that request to the node that is currently the least utilized. You can configure NLB to preference certain nodes over others.

NLB is suitable for stateless applications such as the web tier of multi-tier applications because it does not matter which web server a client connects to when connecting to a multi-tier application. NLB is unsuitable for stateful applications such as traditional file servers and database servers as these applications require a persistent connection to a particular server rather than any server handling the connection.

NLB is failure-aware. This means that if one of the nodes in the NLB cluster goes offline, requests will no longer be forwarded to that node, but other nodes in the cluster will continue to accept requests. When the failed node returns to service, incoming requests will be redirected until traffic is balanced across all nodes in the cluster.

How NLB Works

When you configure an application to use NLB, clients address the application using the NLB cluster address rather than the address of nodes that participate in the NLB cluster. The *NLB cluster address* is a virtual address that is shared between the hosts in the NLB cluster.

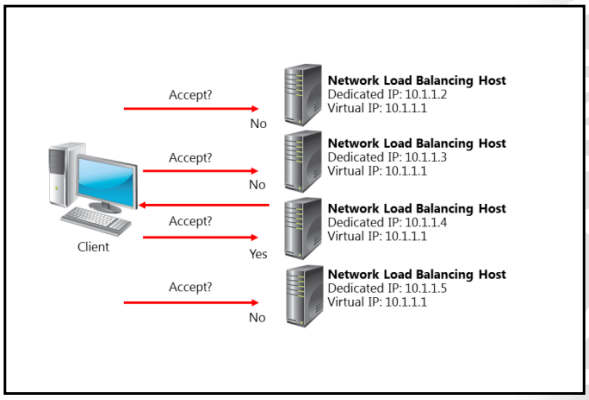

NLB directs traffic in the following manner: All hosts in the NLB cluster receive the incoming traffic, but only one node in the cluster, which is determined through the NLB process, will accept that traffic. All other nodes in the NLB cluster will drop the traffic.

Which node in the NLB cluster accepts the traffic depends on the configuration of port rules and affinity settings. Through these settings, you can determine if traffic that uses a particular port and protocol will be accepted by a particular node, or whether any node in the cluster will be able to accept and respond.

NLB also sends traffic to nodes based on current node utilization. New traffic is directed to nodes that are the least utilized. For example, if you have a four node cluster where three of the nodes are responding to requests from 10 clients and one node is responding to requests from five clients, the node that has fewer clients will receive more incoming traffic until utilization is more evenly balanced across the nodes.

How NLB Works with Server Failures and Recovery

NLB is able to detect the failure of cluster nodes. When a cluster node is in a failed state, it is removed from the cluster, and the hosts in the cluster do not direct new traffic to the node. Failure is detected by using heartbeats. NLB cluster heartbeats are transmitted every second between nodes in a cluster. A node is automatically removed from a NLB cluster if it misses five consecutive heartbeats. Heartbeats are transmitted over a network that is usually different from the network that the client uses to access the cluster. When a node is added or removed from a cluster, a process known as convergence occurs. Convergence allows the cluster to determine its current configuration. Convergence can only occur if each node is configured with the same port rules.

NLB cluster heartbeats are transmitted every second between nodes in a cluster

Convergence occurs when:

- A node misses five consecutive heartbeats, at which time it is automatically removed from an NLB cluster
- A node that was member of a cluster returns to functionality
- An administrator adds or removes a node manually

Nodes can be configured to rejoin a cluster automatically by setting the initial host state setting on the node's properties using the Network Load Balancing Manager. By default, a host that is a member of a cluster will attempt to rejoin that cluster automatically. For example, after applying a software update, if you restart a server that is a member of an NLB cluster the server will rejoin the cluster automatically after the restart process completes.

Administrators can add or remove nodes manually from NLB clusters. When an administrator removes a node, they can choose to perform a Stop or a Drainstop action. The Stop action terminates all existing connections to the cluster node and stops the NLB service. The Drainstop action blocks all new connections without terminating existing sessions. Once all current sessions end, the NLB service is stopped.

NLB can only detect server failure; it cannot detect application failure. This means that if a web application fails but the server remains operational, the NLB cluster will continue to forward traffic to the cluster node that hosts the failed application. One way to manage this problem is to implement a monitoring solution such as Microsoft® System Center 2012 - Operations Manager. With Operations Manager, you can monitor the functionality of applications. You can also configure Operations Manager to generate an alert in the event that an application on a cluster node fails. An alert in turn can configure a remediation action, such as restarting services, restarting the server, or withdrawing the node from the NLB cluster so that the node does not receive further incoming traffic.

NLB Features in Windows Server 2012 and Windows Server 2012 R2

The most substantial change to NLB features in Windows Server 2012 and Windows Server 2012 R2 is the inclusion of Windows PowerShell support. There is no difference between the Windows PowerShell cmdlets available in Windows Server 2012 and Windows Server 2012 R2. The NetworkLoadBalancingClusters module contains 35 NLB–related cmdlets. This module becomes available on a server when the NLB Remote Server Administration Tools (RSATs) are installed. The Windows PowerShell cmdlets have the following nouns:

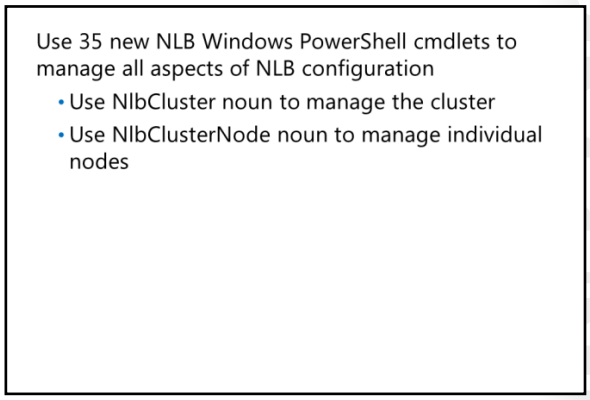

Windows PowerShell NLB Nouns	Description	Windows PowerShell Verbs
NlbClusterNode	Use to manage a cluster node.	Add, Get, Remove, Resume, Set, Start, Stop, and Suspend
NlbClusterNodeDip	Use to configure the cluster node's dedicated management IP.	Add, Get, Remove, and Set
NlbClusterPortRule	Use to manage port rules.	Add, Disable, Enable, Get, Remove, and Set
NlbClusterVip	Use to manage the NLB cluster's virtual IP.	Add, Get, Remove, and Set
NlbCluster	Use to manage the NLB cluster.	Get, New, Remove, Resume, Set, Start, Stop, and Suspend
NlbClusterDriverInfo	Provides information about the NLB cluster driver.	Get

Windows PowerShell NLB Nouns	Description	Windows PowerShell Verbs
NlbClusterNodeNetworkInterface	Use to retrieve information about a cluster node's network interface driver.	Get
NlbClusterIpv6Address	Use to configure the cluster's IPv6 address.	New
NlbClusterPortRuleNodeHandlingPriority	Use to set priority on a per-port rule basis.	Set
NlbClusterPortRuleNodeWeight	Use to set node weight on a per-port rule basis.	Set

 **Note:** To see the list of Windows PowerShell cmdlets for NLB, you can use the **get-command –module NetworkLoadBalancingClusters** command.

Lesson 2
Configuring an NLB Cluster

To deploy NLB successfully, you must first have a firm understanding of its deployment requirements. You must also plan the manner in which you are going to use port rules and affinity settings to ensure that traffic to the application that is being hosted on the NLB cluster is managed appropriately.

This lesson provides you with information about the infrastructure requirements that you must consider before you deploy NLB. It also provides you with important information on how to configure NLB clusters and nodes to best suit your objectives.

Lesson Objectives

After completing this lesson you will be able to:

- Describe NLB deployment requirements.
- Deploy NLB.
- Explain configuration options for NLB.
- Configure NLB affinity and port rules.
- Describe network considerations for NLB.

Deployment Requirements for NLB

NLB requires that all hosts in the NLB cluster reside on the same TCP/IP subnet. Although TCP/IP subnets can be configured to span multiple geographic locations, NLB clusters are unlikely to achieve convergence successfully if the latency between nodes exceeds 250 milliseconds. When you are designing geographically dispersed NLB clusters, you should instead choose to deploy an NLB cluster at each site, and then use Domain Name System (DNS) round robin to distribute traffic between sites.

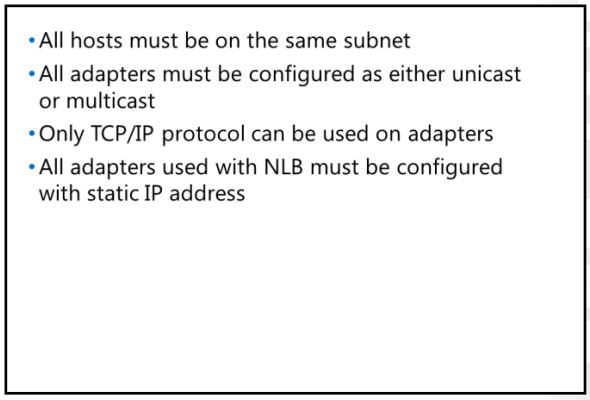

All network adapters within an NLB cluster must be configured as either unicast or multicast. You cannot configure an NLB cluster where there is a mixture of unicast and multicast adapters. When using unicast mode, the network adapter must support changing its MAC address.

You can only use TCP/IP protocol with network adapters that participate in NLB clusters. NLB supports IPv4 and IPv6. The IP addresses of servers that participate in an NLB cluster must be static and must not be dynamically allocated. When you install NLB, Dynamic Host Configuration Protocol (DHCP) is disabled on each interface that you configure to participate in the cluster.

All editions of Windows Server 2012 and Windows Server 2012 R2 support NLB. Microsoft supports NLB clusters with nodes that are running different editions of Windows Server 2012 and Windows Server 2012 R2. However, as a best practice, NLB cluster nodes should be computers with similar hardware specifications, and that are running the same edition of the Windows Server 2012 or Windows Server 2012 R2 operating systems.

Demonstration: Deploying NLB

This demonstration shows how to create a Windows Server 2012 R2 NLB cluster.

Demonstration Steps

Create a Windows Server 2012 R2 NLB cluster

1. Sign in to LON-SVR1 as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. From the Tools menu, open the Windows PowerShell Integrated Scripting Environment (ISE).
3. Enter the following commands, and press Enter after each command:

```
Invoke-Command -Computername LON-SVR1,LON-SVR2 -command {Install-WindowsFeature
NLB,RSAT-NLB}
New-NlbCluster -InterfaceName "Ethernet" -OperationMode Multicast -ClusterPrimaryIP
172.16.0.42 -ClusterName LON-NLB
Add-NlbClusterNode -InterfaceName "Ethernet" -NewNodeName "LON-SVR2" -
NewNodeInterface "Ethernet"
```

4. Open Network Load Balancing Manager from the Tools menu, and view the cluster.

Configuration Options for NLB

Configuring NLB clusters involves specifying how hosts in the cluster will respond to incoming network traffic. How NLB directs traffic depends on the port and protocol that it is using, and whether the client has an existing network session with a host in the cluster. You can configure these settings by using port rules and affinity settings.

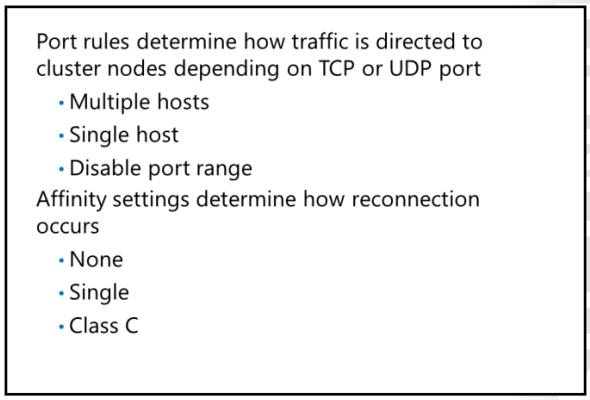

Port Rules

With port rules, you can configure how the NLB cluster directs requests to specific IP addresses and ports. You can load balance traffic on Transmission Control Protocol (TCP) port 80 across all nodes in an NLB cluster, while directing all requests to TCP port 25 to a specific host.

To specify how you want to distribute requests across nodes in the cluster, you configure a filtering mode when creating a port rule. You can do this in the Add/Edit Port Rule dialog box, which you can use to configure one of the following filtering modes:

- Multiple hosts. When you configure this mode, all NLB nodes respond according to the weight assigned to each node. Node weight is calculated automatically, based on the performance characteristics of the host. If a node fails, other nodes in the cluster continue to respond to incoming requests. Multiple host filtering increases availability and scalability, as you can increase capacity by adding nodes, and the cluster continues to function in the event of node failure.
- Single host. When you configure this mode, the NLB cluster directs traffic to the node that is assigned the highest priority. If the node that is assigned the highest priority is unavailable, the host assigned the next highest priority manages the incoming traffic. Single host rules increase availability but do not increase scalability.

Note: The highest priority is the lowest number, with a priority of one being a higher priority than a priority of 10.

- Disable this port range. When you configure this option, all packets for this port range are dropped, without being forwarded to any cluster nodes. If you do not disable a port range and there is no existing port rule, the traffic is forwarded to the host with the lowest priority number.

You can use the following Windows PowerShell cmdlets to manage port rules:

- **Add-NlbClusterPortRule**. Use this cmdlet to add a new port rule.
- **Disable-NlbClusterPortRule**. Use this cmdlet to disable an existing port rule.
- **Enable-NlbClusterPortRule**. Use this cmdlet to enable a disabled port rule.
- **Set-NlbClusterPortRule**. Use this cmdlet to modify the properties of an existing port rule.
- **Remove-NlbClusterPortRule**. Use this cmdlet to remove an existing port rule.

Note: Each node in a cluster must have identical port rules. The exception to this is the load weight (in multiple-hosts filter mode) and handling priority (in single-host filter mode). Otherwise, if the port rules are not identical, the cluster will not converge.

Affinity

Affinity determines how the NLB cluster distributes requests from a specific client. Affinity settings only come into effect when you are using the multiple hosts filtering mode. You can select from the following affinity modes:

- None. In this mode, any cluster node responds to any client request, even if the client is reconnecting after an interruption. For example, the first webpage on a web application might be retrieved from the third node, the second webpage from the first node, and the third webpage from the second node. This affinity mode is suitable for stateless applications.
- Single. When you use this affinity mode, a single cluster node handles all requests from a single client. For example, if the third node in a cluster handles a client's first request, then all subsequent requests are also handled by that node. This affinity mode is useful for stateful applications.
- Class C. When you set this mode, a single node will respond to all requests from a class C network (one that uses the 255.255.255.0 subnet mask). This mode is useful for stateful applications where the client is accessing the NLB cluster through load balanced proxy servers. These proxy servers will have different IP addresses, but they will be within the same class C (24 bit) subnet block.

Host Parameters

You configure the host parameters for a host by clicking the host in the Network Load Balancing Manager console, and then from the Host menu, clicking Properties. You can configure the following host settings for each NLB node:

- Priority. Each NLB node is assigned a unique priority value. If no existing port rule matches the traffic that is addressed to the cluster, traffic will be assigned to the NLB node that is assigned the lowest priority value.
- Dedicated IP address. You can use this parameter to specify the address that the host uses for remote management tasks. When you configure a dedicated IP address, NLB configures port rules so that they do not affect traffic to that address.
- Subnet mask. When you are selecting a subnet mask, ensure that there are enough host bits to support the number of servers in the NLB cluster, and any routers that connect the NLB cluster to the

rest of the organizational network. For example, if you plan to have a cluster that has 32 nodes and supports two routes to the NLB cluster, you will need to set a subnet mask that supports 34 host bits or more—such as 255.255.255.192.

- Initial host state. You can use this parameter to specify the actions the host will take after a reboot. In the default Started state, the host will rejoin the NLB cluster automatically. The Suspended state pauses the host, and allows you to perform operations that require multiple reboots without triggering cluster convergence. The Stopped state stops the node.

Demonstration: Configuring NLB Affinity and Port Rules

In this demonstration, you will see how to:

- Configure affinity for NLB cluster nodes
- Configure NLB port rules

Demonstration Steps

Configure affinity for NLB cluster nodes

1. On LON-SVR2, on the taskbar, click the Windows PowerShell icon.
2. In Windows PowerShell, enter each of the following commands, pressing Enter after each command:

```
Cmd.exe
Mkdir c:\porttest
Xcopy /s c:\inetpub\wwwroot c:\porttest
Exit
New-Website –Name PortTest –PhysicalPath “C:\porttest” –Port 5678
New-NetFirewallRule –DisplayName PortTest –Protocol TCP –LocalPort 5678
```

Configure NLB port rules

1. On LON-SVR1, open the Network Load Balancing Manager console.
2. Remove the **All port** rule.
3. In Network Load Balancing Manager, edit the properties of the LON-NLB cluster.
4. Add a port rule with the following properties:
   - Port range: **80 to 80**
   - Protocols: **Both**
   - Filtering mode: **Multiple Host**
   - Affinity: **None**
5. Create a port rule with the following properties:
   - Port range: **5678 to 5678**
   - Protocols: **Both**
   - Filtering mode: **Single Host**
6. Edit the host properties of LON-SVR1.
7. Configure the port rule for port **5678** and set handling priority to **10**.

Network Considerations for NLB

When you are designing a network to support an NLB cluster, you must consider several factors. The primary decision is whether you want to configure the NLB cluster to use Unicast or Multicast cluster operation mode.

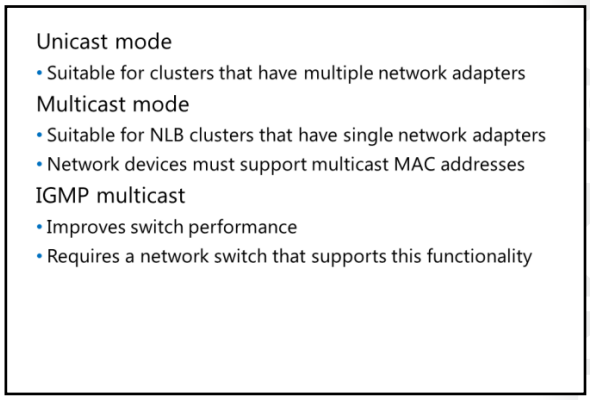

Unicast Mode

When you configure a NLB cluster to use unicast mode, all cluster hosts use the same unicast MAC address. Outgoing traffic uses a modified MAC address that is determined by the cluster host's priority setting. This prevents the switch that handles outbound traffic from having problems with all cluster hosts using the same MAC address.

When you use unicast mode with a single network adapter on each node, only computers that use the same subnet can communicate with the node using the node's assigned IP address. If you have to perform any node management tasks, (such as connecting with the Windows operating system feature Remote Desktop to apply software updates), you will need to perform these tasks from a computer that is on the same TCP/IP subnet as the node.

When you use unicast mode with two or more network adapters, one adapter will be used for dedicated cluster communication, and the other adapter or adapters can be used for management tasks. When you use unicast mode with multiple network adapters, you can perform cluster management tasks such as connecting using Remote PowerShell to add or remove roles and features.

Unicast mode can also minimize problems that occur when cluster nodes also host other non-NLB related roles or services. For example, using unicast mode means that a server that participates in a web server cluster on port 80 may also host another service such as DNS or DHCP. Although this is possible, we recommend that all cluster nodes have the same configuration.

Multicast Mode

When you configure an NLB cluster to use multicast mode, each cluster host keeps its original MAC address, but also is assigned an additional multicast MAC address. Each node in the cluster is assigned the same additional MAC multicast address. Multicast mode requires network switches and routers that support multicast MAC addresses.

Internet Group Management Protocol Multicast

Internet Group Management Protocol (IGMP) multicast mode is a special form of multicast mode that prevents the network switch from being flooded with traffic. When you deploy IGMP multicast mode, traffic is forwarded only through switch ports that participate in the NLB cluster. IGMP multicast mode requires switch hardware that supports this functionality.

Network Considerations

You can improve NLB cluster performance when using unicast mode by using separate virtual local area networks (VLANs) for cluster traffic and management traffic. Using VLANs segment traffic, you can prevent management traffic from affecting cluster traffic. When you host NLB nodes on virtual machines using Windows Server 2012 or Windows Server 2012 R2, you can also use network virtualization to segment management traffic from cluster traffic.

Lesson 3
Planning an NLB Implementation

When you are planning an NLB implementation, you must ensure that the applications that you deploy are appropriate for NLB. Not all applications are suitable for deployment on NLB clusters, and it is important for you to be able to identify which ones can benefit from this technology. You also need to know what steps you can take to secure NLB, and you should be familiar with the options that you have to scale NLB, in case the application hosted on the NLB cluster requires greater capacity.

Lesson Objectives

After completing this lesson, you will be able to:

- Explain how to design application and storage support for NLB.
- Describe the special considerations for deploying NLB clusters on virtual machines.
- Describe the considerations for securing NLB.
- Describe the considerations for scaling NLB.
- Describe the considerations for upgrading an NLB cluster to Windows Server 2012 or Windows Server 2012 R2.

Designing Applications and Storage Support for NLB

Because client traffic can be directed to any node in an NLB cluster, each node in the cluster must be able to provide a consistent experience. Therefore, when you are designing applications and storage support for NLB applications, you must ensure that you configure each node in the same way, and that each node has access to the same data.

- Each node in an NLB cluster needs to have the same configuration
- Each node needs access to the same consistent application data
- Use IIS shared configuration to ensure that web application configuration is consistent across NLB nodes
- Use CSVs to host shared application and configuration data for NLB applications

When a highly available application has multiple tiers—such as a web application that includes an SQL Server database tier—the web application tier is hosted on an NLB cluster. SQL Server, as a stateful application, is not made highly available using NLB. Instead, you use technologies such as failover clustering, mirroring, or AlwaysOn Availability Groups, to make the SQL Server database tier highly available.

All hosts in an NLB cluster should run the same applications and be configured in the same way. When you are using web applications, you can use the Internet Information Services (IIS) 8.0 shared configuration functionality to ensure that all nodes in the NLB cluster are configured in the same manner.

You can also use technologies such as file shares that are hosted on Cluster Shared Volumes (CSVs) to host application configuration information. File shares that are hosted on CSVs allow multiple hosts to have access to application data and configuration information. File shares that are hosted on CSVs are a feature of Windows Server 2012.

Considerations for Deploying an NLB Cluster on Virtual Machines

As organizations transition from physical to virtual deployments, administrators must consider several factors when determining the placement of NLB cluster nodes on Hyper-V hosts. This includes the network configuration of virtual machines, the configuration of the Hyper-V hosts, and the benefits of using the Hyper-V high availability features in conjunction with NLB.

- Configure virtual machines with multiple network adapters
- Configure one network adapter on each node member to use a shared private network switch
- Configure the NLB cluster to use unicast mode and enable MAC address spoofing on Hyper-V host
- Use the shared private network switch for cluster communication
- When NLB nodes span multiple sites, use network virtualization to separate the cluster network

Virtual Machine Placement

You should place NLB cluster nodes on separate hard disks on the Hyper-V host. That way, if a disk or disk array fails, even if one node becomes unavailable, other NLB cluster nodes that are hosted on the same Hyper-V host will remain online. As a best practice, you should configure the Hyper-V host with redundant hardware, including redundant disks, network adapters, and power supplies. This will minimize the chance that hardware failure on the Hyper-V host will lead to all nodes in an NLB cluster becoming unavailable. When you are using multiple network adapters, configure network teaming to ensure that virtual machines are able to maintain access to the network even in the event that individual network adapter hardware suffers a failure.

Where possible, deploy NLB virtual machine nodes on separate hyper-V hosts. This protects the NLB cluster from other types of server failure, such as the failure of a motherboard, or any other single point of failure. When you are planning this type of configuration, ensure that the virtual machines that participate in the NLB cluster are located on the same TCP/IP subnet.

Virtual Machine Network Configuration

Because adding additional virtual network adapters is a straightforward process, you can configure the NLB cluster to use unicast mode, and then deploy each virtual machine with multiple network adapters. You should create separate virtual switches for cluster traffic and node management traffic, because segmenting traffic can improve performance. You can also use network virtualization to partition cluster traffic from node management traffic. You can use VLAN tags as a method of partitioning cluster traffic from node management traffic.

When you are using unicast mode, ensure that you enable MAC address spoofing for the virtual network adapter on the Hyper-V host. You can do this by editing the virtual network adapter's settings on the Virtual Machine Settings dialog box, which is available through the Hyper-V Manager. Enabling MAC address spoofing allows unicast mode to configure MAC address assignment on the virtual network adapter.

NLB Cluster vs. Virtual Machine High Availability

Virtual machine high availability is the process of placing virtual machines on failover clusters. When a failover cluster node fails, the virtual machine fails over so that it is hosted on another node. Although failover clustering and NLB are both high availability technologies, they serve different purposes. Failover clustering supports stateful applications such as SQL Server, whereas NLB is suited to stateless applications such as websites. Highly available virtual machines do not allow an application to scale, because you cannot add nodes to increase capacity. However, it is possible to deploy NLB cluster nodes as highly available virtual machines. In this scenario, the NLB cluster nodes fail over to a new Hyper-V host in the event that the original Hyper-V host fails.

The degree of availability and redundancy required will fluctuate, depending on the application. A business-critical application that costs an organization millions of dollars when it is down requires an availability that differs from that of an application that causes minimal inconvenience if it is offline.

Considerations for Securing NLB

NLB clusters are almost always used to host web applications that are important to the organization. Because of this importance, you should take steps to secure NLB, both by restricting the traffic that can address the cluster, and by ensuring that appropriate permissions are applied.

- Use NLB cluster port rules to discard traffic not related to cluster applications
- Use firewall rules on cluster nodes to drop traffic not related to cluster applications or node management
- Configure applications to respond only to traffic that is addressed to the cluster
- Use SANs to create certificates that support the application name and node names
- Implement principle of least privilege to ensure that only authorized users have appropriate permissions on nodes

Configure Port Rules

When securing NLB clusters, you must first ensure that you create port rules to block traffic to all ports other than those used by applications hosted on the NLB cluster. When you do this, all incoming traffic that is not addressed specifically to applications that are running on the NLB cluster will be dropped. If you do not perform this first step, all incoming traffic that is not managed by a port rule will be forwarded to the cluster node with the lowest cluster priority value.

Configure Firewall Rules

You should also ensure that Windows Firewall with Advanced Security is configured on each NLB cluster node. When you enable NLB on a cluster node, the following firewall rules that allow NLB to function and communicate with other nodes in the cluster are created and enabled automatically:

- Network Load Balancing (DCOM-In)
- Network Load Balancing (ICMP4-ERQ-In)
- Network Load Balancing (ICMP6-ERQ-In)
- Network Load Balancing (RPCSS)
- Network Load Balancing (WinMgmt-In)
- Network Load Balancing (ICMP4-DU-In)
- Network Load Balancing (ICMP4-ER-In)
- Network Load Balancing (ICMP6-DU-In)
- Network Load Balancing (ICMP6-EU-In)

When created, these firewall rules do not include scope settings. In high security environments, you would configure an appropriate local IP address or IP address range, and a remote IP address for each of these rules. The remote IP address or address range should include the addresses that are used by other hosts in the cluster.

When you configure additional firewall rules, remember the following:

- When you are using multiple network adapters in unicast mode, configure different firewall rules for each network interface. For the interface used for management tasks, you should configure the firewall rules to allow inbound management traffic only—for example, enabling the use of remote Windows PowerShell, Windows Remote Management, and Remote Desktop for management tasks. You should configure the firewall rules on the network interface used by the cluster node, to provide an application to the cluster, and to allow access to that application. For example, allow incoming traffic on TCP ports 80 and 443 on an application that uses the HTTP and HTTPS protocols.
- When you are using multiple network adapters in multicast mode, configure firewall rules that allow access to applications that are hosted on the cluster, but block access to other ports.

Configure Applications to Respond Only to Traffic Addressed to the Cluster

You should configure applications on each node to respond only to traffic that is addressed to the cluster, and to ignore application traffic that is addressed to the individual node. For example, if you deploy a web application that is designed to respond to traffic addressed to www.adatum.com, there will be a website on each node that will accept traffic on port 80. Depending on the NLB cluster configuration, it is possible that traffic that is addressed to the node on port 80 will generate a direct response. For example, users may be able to access the A. Datum web application by entering the address http://nlb-node-3.adatum.com in a browser, instead of entering the address http://www.adatum.com. You can secure applications from this type of direct traffic by configuring them to respond only to traffic that uses the NLB cluster address. For web applications, you can do this by configuring the website to use a host header. Each application that runs on an NLB cluster will have its own unique method of allowing you to configure the application to respond only to traffic directed at the cluster, rather than at the individual cluster node.

Securing Traffic with an SSL Certificate

NLB websites must all use the same website name. When you are securing websites that you make highly available using NLB, you need to ensure that each website has an SSL certificate that matches the website name. You can use host headers on each node. In most cases, you will install the same website certificate on each node in the NLB cluster, because this is simpler than procuring separate certificates for each cluster node. In some cases, you will need to procure certificates that support subject alternative names (SANs). Certificates that support SANs allow a server to be identified by multiple names, such as the name used by the clustered application and the name of the cluster node. For example, a certificate with a SAN might support the names www.adatum.com, node1.adatum.internal, node2.adatum.internal, node3.adatum.internal, and node4.adatum.internal.

Principle of Least Privilege

Ensure that the users are only delegated permissions for tasks that they need to perform on the NLB node. Members of the local Administrators group on any single node are able to add and remove cluster nodes, even if they are not members of the local Administrators group on those nodes. Applications that run on NLB clusters should be configured so that they do not require application administrators to have local Administrator privileges on the servers that host the application. Only users whose job role requires them to be able to make remote management connections to NLB cluster nodes should be able to make those connections.

Considerations for Scaling NLB

Scaling is the process of increasing the capacity of an NLB cluster. For example, if you have a four-node NLB cluster and each node in the cluster is being heavily utilized to the point where the cluster cannot manage more traffic, you can add additional nodes. Adding nodes will spread the same load across more computers, reducing the load on each current cluster node. Capacity increases because a larger number of similarly configured computers can manage a higher workload than a smaller number of similarly configured computers.

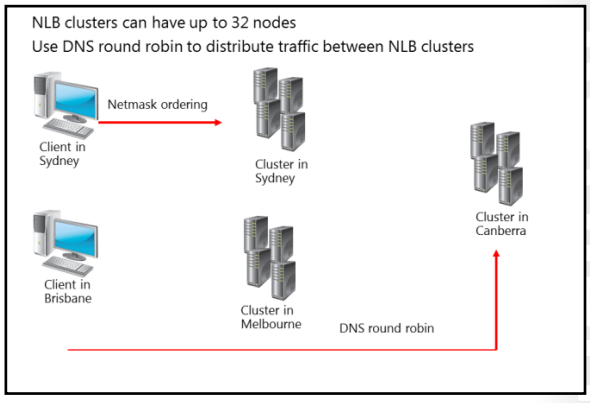

An NLB cluster supports up to 32 nodes. This means that you can scale-out a single NLB cluster so that 32 separate nodes participate in that cluster. When you consider scaling an application so that it is hosted on a 32-node NLB cluster, remember that each node in the cluster must be on the same TCP/IP subnet.

As an alternative to building single NLB clusters is to build multiple NLB clusters, you can use DNS round robin to share traffic between them. DNS round robin is a technology that allows a DNS server to provide requesting clients with different IP addresses to the same hostname, in sequential order. For example, if there are three addresses associated with a hostname, the first requesting host receives the first address, the second receives the second address, and the third receives the third address, and so forth. When you use DNS round robin with NLB, you associate the IP addresses of each cluster with the hostname that is used by the application.

Distributing traffic between NLB clusters using DNS round robin also allows you to deploy NLB clusters across multiple sites. DNS round robin can also be used in conjunction with netmask ordering. This technology ensures that clients on a subnet are provided with an IP address of a host on the same network, if one is available. For example, you might deploy three four-node NLB clusters in the cities of Sydney, Melbourne, and Canberra, and use DNS round robin to distribute traffic between them. With netmask ordering, a client in Sydney that is accessing the application in Sydney will be directed to the NLB cluster hosted in Sydney. A client that is not on the same subnet as the NLB cluster nodes, such as a client in the city of Brisbane, would be directed by DNS round robin to either the Sydney, Melbourne, or Canberra NLB cluster.

Considerations for Upgrading NLB Clusters

Upgrading NLB clusters involves moving cluster nodes from one host operating system—for example, Windows Server® 2003 or Windows Server 2008—to Windows Server 2012. Upgrading the cluster might not require performing an operating system upgrade on each node, because in some cases the original host operating system might not support a direct upgrade to Windows Server 2012. In cases where the original host operating system does not support a direct upgrade to Windows Server 2012, you can perform a migration.

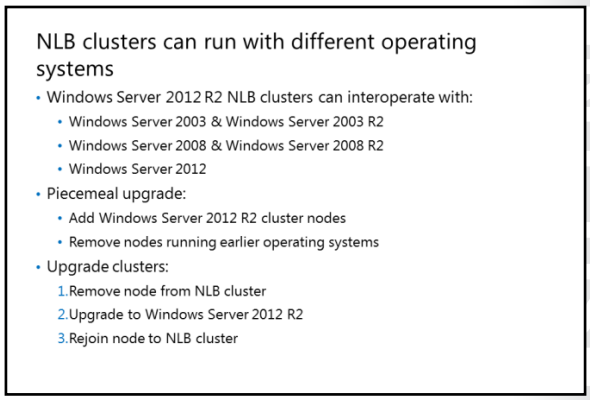

A key consideration when you upgrade NLB clusters is to remember that NLB supports having clusters that are running a mixture of operating systems. This means that you can have a cluster that runs a mixture of Windows Server 2003, Windows Server 2008, and Windows Server 2012. Keep in mind that while mixed operating system NLB clusters are supported, they are not recommended. You should configure the NLB cluster so that all hosts are running the same operating system as soon as possible.

Note: In some situations, it will not be possible to upgrade the operating system of a cluster node.

When you are performing an upgrade, you can use one of the following strategies:

- Piecemeal upgrade. During this type of upgrade, you add new Windows Server 2012 nodes to an existing cluster, and then remove the nodes that are running earlier versions of the Windows Server operating system. This type of upgrade is appropriate when the original hardware and operating system does not support a direct upgrade to Windows Server 2012.

- Rolling upgrade. During this type of upgrade, you upgrade one node in the cluster at a time. You do this by taking the node offline, performing the upgrade, and then rejoining the node back to the cluster.

To learn more about upgrading NLB clusters, go to:
http://go.microsoft.com/fwlink/?LinkId=270037

Lab: Implementing NLB

Scenario

A. Datum Corporation is an engineering and manufacturing company. The organization is based in London, England, and is quickly expanding into Australia. As the company expands, the need for scalable web applications has increased. To address this need, you will develop a pilot program to test the deployment of NLB on hosts running the Windows Server 2012 operating system.

Because you intend to automate the process of deploying Windows NLB clusters, you will use Windows PowerShell to perform many of the cluster setup and configuration tasks. You will also configure port rules and affinity, which will allow you to deploy multiple load balanced web applications on the same Windows NLB clusters.

Objectives

After completing this lab, the students will be able to:

- Implement an NLB cluster.
- Configure and manage an NLB cluster.
- Validate high availability for the NLB cluster.

Lab Setup

Estimated Time: 60 minutes

Virtual machines	**20412C-LON-DC1** **20412C-LON-SVR1** **20412C-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 **20412C-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 two through four for 20412C-LON-SVR1 and 20412C-LON-SVR2.

Exercise 1: Implementing an NLB Cluster

Scenario

You eventually want to automate the process of deploying Windows Server 2012 NLB clusters. To accomplish this, you will use Windows PowerShell to perform the majority of the NLB cluster deployment tasks.

The main tasks for this exercise are as follows:

1. Verify website functionality for stand-alone servers
2. Install NLB
3. Create a new Windows Server 2012 NLB cluster
4. Add a second host to the cluster
5. Validate the NLB cluster

▶ Task 1: Verify website functionality for stand-alone servers

1. On LON-SVR1, navigate to the **c:\inetpub\wwwroot** folder.
2. Open **iis-8.png** in Microsoft Paint, and use the Paintbrush tool and the color red to mark the IIS logo in a distinctive manner.
3. Close File Explorer.
4. Switch to LON-DC1, and then open Windows Internet Explorer.
5. Navigate to **http://LON-SVR1**, and verify that the web page is marked in a distinctive manner with the color red.
6. Navigate to **http://LON-SVR2**, and verify that the website is not marked in a distinctive manner.
7. Close Internet Explorer.

▶ Task 2: Install NLB

1. On LON-SVR1, open Windows PowerShell ISE.
2. Type the following command, and then press Enter:

```
Invoke-Command -Computername LON-SVR1,LON-SVR2 -command {Install-WindowsFeature
NLB,RSAT-NLB}
```

▶ Task 3: Create a new Windows Server 2012 NLB cluster

1. On LON-SVR1, in Windows PowerShell ISE, type the following command, and then press Enter:

```
New-NlbCluster -InterfaceName "Ethernet" -OperationMode Multicast -ClusterPrimaryIP
172.16.0.42 -ClusterName LON-NLB
```

2. In Windows PowerShell ISE, type the following command, and then press Enter:

```
Invoke-Command -Computername LON-DC1 -command {Add-DNSServerResourceRecordA      –
zonename adatum.com –name LON-NLB –Ipv4Address 172.16.0.42}
```

▶ Task 4: Add a second host to the cluster

- On LON-SVR1, in Windows PowerShell ISE, type the following command, and then press Enter:

```
Add-NlbClusterNode -InterfaceName "Ethernet" -NewNodeName "LON-SVR2" -
NewNodeInterface "Ethernet"
```

▶ Task 5: Validate the NLB cluster

1. On LON-SVR1, open the Network Load Balancing Manager console, and verify that nodes LON-SVR1 and LON-SVR2 display with the status of **Converged**.

2. View the properties of the LON-NLB cluster, and verify the following:
   - The cluster is set to use the **Multicast** operations mode.
   - There is a single port rule named **All** that starts at port **0** and ends at port **65535** for both **TCP** and **UDP** protocols, and that it uses **Single** affinity.

Results: After completing this exercise, you will have successfully implemented an NLB cluster.

Exercise 2: Configuring and Managing the NLB Cluster

Scenario

You want to deploy multiple separate websites to the NLB cluster and then differentiate these websites based on port address. To do this, you want to ensure that you are able to configure and validate port rules. You also want to experiment with affinity settings to ensure that requests are distributed evenly across hosts.

The main tasks for this exercise are as follows:

1. Configure port rules and affinity
2. Validate port rules
3. Manage host availability in the NLB cluster

▶ Task 1: Configure port rules and affinity

1. On LON-SVR2, open Windows PowerShell.
2. In Windows PowerShell, enter the following commands, and then press Enter after each command:

```
Cmd.exe
Mkdir c:\porttest
Xcopy /s c:\inetpub\wwwroot c:\porttest
Exit
New-Website -Name PortTest -PhysicalPath "C:\porttest" -Port 5678
New-NetFirewallRule -DisplayName PortTest -Protocol TCP -LocalPort 5678
```

3. Open File Explorer, and then browse to and open **c:\porttest\iis-8.png** in Microsoft Paint.
4. Use the **Blue** paintbrush to mark the IIS logo in a distinctive manner.
5. Switch to LON-DC1.
6. Open Internet Explorer, and navigate to **http://LON-SVR2:5678**.
7. Verify that the IIS Start page with the image marked with blue displays.
8. Switch to LON-SVR1.
9. On LON-SVR1, open Network Load Balancing Manager, and view the cluster properties of LON-NLB.
10. Remove the **All port** rule.
11. Add a port rule with the following properties:
    - Port range: **80 to 80**
    - Protocols: **Both**
    - Filtering mode: **Multiple Host**
    - Affinity: **None**

12. Create a new port rule with the following properties:
    - Port range: **5678 to 5678**
    - Protocols: **Both**
    - Filtering mode: **Single Host**
13. Close the **LON-NLB(172.16.0.42)**.
14. Edit the host properties of LON-SVR1.
15. Configure the **Handling Priority** value of the port rule for port **5678** as **10**.

▶ Task 2: Validate port rules

1. Switch to LON-DC1.
2. Using Internet Explorer, navigate to **http://lon-nlb**, refresh the web page 20 times, and verify that web pages with and without the distinctive red marking display.
3. On LON-DC1, navigate to address **http://LON-NLB:5678**, refresh the web page 20 times, and verify that only the web page with the distinctive blue marking displays.

▶ Task 3: Manage host availability in the NLB cluster

1. Switch to LON-SVR1.
2. On LON-SVR1, use the Network Load Balancing Manager console to suspend LON-SVR1.
3. Verify that node LON-SVR1 displays as **Suspended**, and that node LON-SVR2 displays as **Converged**.
4. Resume and then Start LON-SVR1.
5. Verify that both node LON-SVR1 and LON-SVR2 now display as **Converged**.

Results: After completing this exercise, you will have successfully configured and managed an NLB cluster.

Exercise 3: Validating High Availability for the NLB Cluster

Scenario

As part of preparing to deploy NLB in your organization's environment, you want to ensure that it is possible to perform maintenance tasks (such as reboot operations), without affecting the availability of the websites that are hosted on the cluster. To accomplish this, you will verify availability by rebooting one of the hosts while attempting to access the clustered website. You will also explore the Drainstop functionality.

The main tasks for this exercise are as follows:

1. Validate website availability when the host is unavailable
2. Configure and validate Drainstop
3. Prepare for the next module

▶ Task 1: Validate website availability when the host is unavailable

1. Restart LON-SVR1.
2. Switch to LON-DC1.

3. On LON-DC1, open Internet Explorer, and navigate to **http://LON-NLB**.
4. Refresh the website 20 times. Verify that the website is available, but that it does not display the distinctive red mark on the IIS logo until LON-SVR1 has restarted.

▶ Task 2: Configure and validate Drainstop

1. On LON-SVR1, open the Network Load Balancing Manager console, and initiate a Drainstop on LON-SVR2.
2. On LON-DC1, navigate to **http://lon-nlb**, and verify that only the welcome page with the red IIS logo displays.

▶ Task 3: Prepare for the next module

When you finish the lab, revert the virtual machines to their initial state.

1. On the host computer, start the Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20412C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps two and three for 20412C-LON-SVR1 and 20412C-LON-SVR2.

Results: After completing this exercise, you will have successfully validated high availability for the NLB cluster.

Question: How many additional nodes can you add to the LON-NLB cluster?

Question: What steps would you take to ensure that LON-SVR1 always manages requests for web traffic on port 5678, given the port rules established by the end of this exercise?

Question: What is the difference between a Stop and a Drainstop command?

Module Review and Takeaways

Review Questions

Question: You have created a four-node Windows Server 2012 NLB cluster. The cluster hosts a website that is hosted on IIS. What happens to the cluster if you shut down the World Wide Web publishing service on one of the nodes?

Question: You want to host the www.contoso.com, www.adatum.com, and www.fabrikam.com websites on a four-node NLB cluster. The cluster IP address will be a public IP address, and each fully qualified domain name is mapped in DNS to the cluster's public IP address. What steps should you take on each node to ensure that traffic is directed to the appropriate site?

Question: You have an eight-node Windows NLB cluster that hosts a web application. You want to ensure that traffic from a client that uses the cluster remains with the same node throughout their session, but that traffic from separate clients is distributed equitably across all nodes. Which option do you configure to accomplish this goal?

Real-world Issues and Scenarios

To become a true high-availability solution, use a monitoring solution with NLB that will detect application failure. This is because NLB clusters will continue to direct traffic to nodes with failed applications as long as NLB, which is independent of the application, continues to send heartbeat traffic.

Module 10

Implementing Failover Clustering

Contents:

Module Overview

Providing high availability is very important for any organization that wants to provide continuous service to its users. Failover clustering is one of the main technologies in Windows Server® 2012 that can provide high availability for various applications and services. In this module, you will learn about failover clustering, failover clustering components, and implementation techniques.

Objectives

After completing this module, you will be able to:

- Describe failover clustering.
- Implement a failover cluster.
- Configure highly available applications and services.
- Maintain a failover cluster.
- Implement multisite failover clustering.

Lesson 1
Overview of Failover Clustering

Failover clusters in Windows Server 2012 provide a high-availability solution for many server roles and applications. By implementing failover clusters, you can maintain application or service availability if one or more computers in the failover cluster fail. Before you implement failover clustering, you should be familiar with general high-availability concepts. You must understand clustering terminology and also how failover clusters work.

It also is important to be familiar with new clustering features in Windows Server 2012.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe high availability.
- Describe failover clustering improvements in Windows Server 2012.
- Describe failover clustering improvements in Windows Server® 2012 R2.
- Describe failover cluster components.
- Describe Cluster Shared Volumes (CSVs).
- Describe CSV improvements in Windows Server 2012 R2.
- Define failover and failback.
- Describe a quorum.
- Describe quorum modes in Windows Server 2012 failover clustering.
- Describe how quorum works in Windows Server 2012 R2 failover clustering.
- Describe failover cluster networks.
- Describe failover cluster storage.

What Is High Availability?

Availability refers to a level of service that applications, services, or systems provide, and it is expressed as the percentage of time that a service or system is available. Highly available systems have minimal downtime—whether planned or unplanned, and they usually need to be available on a 24-hour-a-day basis. Usually, these systems are available more than 99 percent of the time, depending on the organization's needs and budget. For example, a system that is unavailable for 8.75 hours per year would have a 99.9 percent availability rating.

- Availability is a level of service expressed as a percentage of time
- Highly-available services or systems are available more than 99 percent of the time
- High availability requirements differ based on how availability is measured
- Planned outages typically are not included when calculating availability

To improve availability, you must implement fault-tolerant mechanisms that mask or minimize how failures of the service's components and dependencies affect the system. You can achieve fault tolerance by implementing redundancy to single points of failure.

Availability requirements must be expressed so that there is no misunderstanding about the implications. Miscommunication about service-level expectations between the customer and the IT organization can result in poor business decisions, such as unsuitable investment levels and customer dissatisfaction.

The availability measurement period can also have a significant effect on the definition of availability. For example, a requirement for 99.9 percent availability over a one-year period allows for 8.75 hours of downtime, whereas a requirement for 99.9 percent availability over a rolling four-week window allows for only 40 minutes of downtime per period.

You also have to identify and negotiate planned outages, maintenance activities, service pack updates, and software updates. These are scheduled outages, and typically are not included as downtime when you calculate the system's availability. You typically calculate availability based on unplanned outages only. However, you have to determine exactly which planned outages you consider as downtime.

Failover Clustering Improvements in Windows Server 2012

Failover clustering has not changed significantly since Windows Server® 2008 R2. However, there are new features and technologies in Windows Server 2012 that help increase scalability and cluster storage availability, and provide better, easier management and faster failover.

Failover clustering improvements in Windows Server 2012	Removed and deprecated failover clustering features in Windows Server 2012
• Increased scalability • Improved CSVs • Cluster-aware updating • Active Directory integration improvements • Management improvements	• Cluster.exe command-line tool • Cluster Automation Server (MSClus) COM interface • Add-ClusterPrintServerRole cmdlet • Printer cluster

The important new features in Windows Server 2012 failover clustering include:

- Increased scalability. In Windows Server 2012, a failover cluster can have 64 physical nodes and can run 4,000 virtual machines on each cluster. This is a significant improvement over Windows Server 2008 R2 which supports only 16 physical nodes and 1,000 virtual machines per cluster. Each cluster you create is now available from the Server Manager console. Server Manager in Windows Server 2012 can discover and manage all clusters created in an Active Directory® Domain Services (AD DS) domain. If the cluster is deployed in multisite scenario, the administrator can now control which nodes in a cluster have votes for establishing quorum. Failover clustering scalability is also improved for virtual machines that are running on clusters. This will be discussed in more detail in *Module 11: Implementing Failover Clustering with Hyper-V*.
- Improved Cluster Shared Volumes (CSVs). This technology was introduced in Windows Server 2008 R2, and it became very popular for providing virtual machine storage. In Windows Server 2012, CSV volumes appear as CSV File System, and CSV supports Server Message Block (SMB) version 2.2 storage for Hyper-V® and other applications. In addition, CSV can use SMB multichannel and SMB Direct to enable traffic to stream across multiple networks in a cluster. For additional security, you can use BitLocker® Drive Encryption for CSV disks, and you can also make CSV storage visible only to a subset of nodes in a cluster. For reliability, CSV volumes can be scanned and repaired with zero offline time.
- Cluster-aware updating. Updating cluster nodes required a lot of preparation and planning in older versions of Windows Server, to minimize or avoid downtime. In addition, the procedure of updating cluster nodes was mostly manual, which required additional administrative effort. In Windows Server 2012, a new technology is introduced for this purpose. This technology is called Cluster-Aware Updating (CAU). This technology automatically updates cluster nodes with Windows Update hotfix, by keeping the cluster online, and minimizing downtime. This technology will be explained in more detail in *Lesson 4: Maintaining a Failover Cluster*.

- Active Directory integration improvements. Beginning with Windows Server 2008, failover clustering has been integrated in AD DS. In Windows Server 2012, this integration is improved. Administrators can create cluster computer objects in targeted organizational units (OUs), or by default in the same OUs as the cluster nodes. This aligns failover cluster dependencies on AD DS with the delegated domain administration model that is used in many IT organizations. In addition, failover clusters now can be deployed with access only to read-only domain controllers.
- Management improvements. Although failover clustering in Windows Server 2012 still uses almost the same management console and the same administrative techniques, there are some important management improvements. The Validation Wizard is improved; the validation speed for large failover clusters is improved; and new tests for CSVs, the Hyper-V role, and virtual machines have been added. In addition, new Windows PowerShell® cmdlets are available for managing clusters, monitoring clustered virtual machine applications, and creating highly available Internet SCSI (iSCSI) targets.

Removed and Deprecated Features

In Windows Server 2012 clustering, some features have been removed or deprecated. If you are moving from an older version of failover clustering, you should be aware of these features:

- The Cluster.exe command-line tool is deprecated; however, it can be optionally installed with failover clustering tools. Windows PowerShell cmdlets for failover clustering roles provide a functionality that is generally the same as Cluster.exe commands.
- The Cluster Automation Server (MSClus) COM interface has been deprecated, but it can be optionally installed with the failover clustering tools.
- Support for 32-bit cluster resource DLLs has been deprecated, but 32-bit DLLs can be installed optionally. Cluster resource DLLs should be updated to 64-bit.
- The Print Server role has been removed from the High Availability Wizard, and it cannot be configured in Failover Cluster Manager.
- The **Add-ClusterPrintServerRole** cmdlet has been deprecated, and it is not supported in Windows Server 2012.

Failover Clustering Improvements in Windows Server 2012 R2

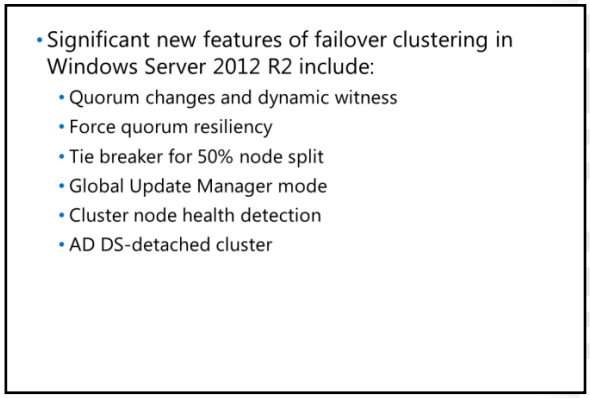

Failover clustering in Windows Server 2012 R2 has been enhanced with many new features, and existing technologies have been updated for better functionality. The quorum model has changed significantly; you now have more options and flexibility for maintaining quorum and a cluster. Also, the Failover Cluster Manager console in Windows Server 2012 R2 has a cluster dashboard where you can quickly see the health status of all managed failover clusters. In the console, next to each failover cluster that you manage, there are icons that indicate whether the cluster is running, the number and status of clustered roles, the node status, and the event status.

The most important new features in failover clustering quorum in Windows Server 2012 R2 are the following:

- Dynamic quorum. This feature enables a cluster to recalculate quorum in the event of node failure and still maintain working clustered roles, even when the number of voting nodes remaining in the cluster is less than 50 percent.
- Dynamic witness. This feature dynamically decides if the witness has a vote to maintain quorum in the cluster.
- Force quorum resiliency. This feature provides additional support and flexibility to manage *split brain syndrome* cluster scenarios. These occur when a cluster breaks into subsets of cluster nodes that are not aware of each other.
- Tie breaker for 50 percent node split. By using this feature, the cluster can adjust the running node's vote status automatically to keep the total number of votes in the cluster at an odd number.

These new quorum options and modes of work are discussed in more detail later in this lesson.

Besides updating quorum, Microsoft has made other helpful changes to failover clustering in Windows Server 2012 R2. The most important changes are the following:

Global Update Manager Mode

The Global Update Manager is responsible for updating the cluster database. In Windows Server 2012, it was not possible to configure how these updates work. Windows Server 2012 R2 enables you to configure the mode of work for the Global Update Manager.

Each time the state of a cluster changes, such as when a cluster resource is offline, all nodes in the cluster must receive notification about the event before the change is committed to the cluster database by the Global Update Manager.

In Windows Server 2012, the Global Update Manager works in *Majority* (read and write) mode. In this mode, when a change happens to a cluster, a majority of the cluster nodes must receive and process the update before it is committed to the database. When the cluster node wants to read the database, the cluster compares the latest time stamp from a majority of the running nodes and uses the data with the latest time stamp.

In Windows Server 2012 R2, the Global Update Manager can also work in the *All* (write) and *Local* (read) mode. When working in this mode, all nodes in the cluster must receive and process an update before it is committed to the database. However, when the database read request is received, the cluster will read the data from the database copy that is stored locally. Because all roles receive and process the update, the local cluster database copy can be considered a relevant source of information.

Windows Server 2012 R2 also supports a third mode for the Global Update Manager. This mode is Majority (write) and Local (read). In this mode, a majority of the cluster nodes must receive and process an update before it is committed to the database. When the database read request is received, the cluster reads the data from the database copy that is stored locally.

In Windows Server 2012 R2, the default setting for Hyper-V failover clusters is Majority (read and write). All other workloads in the clusters use All (write) and Local (read) mode. Majority (write) and Local (read) are not used by default for any workload.

Changing the working mode for the Global Update Manager improves cluster database performance and increases the performance of cluster workloads because a cluster database no longer has to perform at the speed of the slowest node.

Cluster Node Health Detection

In Windows Server 2012, the mechanism for node health detection within a cluster declares a node as down if it does not respond to heartbeats for more than five seconds. In Windows Server 2012 R2, specifically for Hyper-V failover clusters, the default threshold value is increased from five seconds to 10 seconds if nodes are in the same subnet, and to 20 seconds if nodes are in different subnets. This provides increased resiliency for temporary network failures for virtual machines that are running on a Hyper-V cluster, and this delays cluster recovery actions in cases of short network interruptions.

AD DS Detached Cluster

Failover clusters in Windows Server 2012 are integrated with AD DS, and you cannot deploy a cluster if nodes are not members of same domain. When a cluster is created, appropriate computer objects for a cluster name and a clustered role name are created in AD DS.

In Windows Server 2012 R2, you can deploy an *AD DS-detached cluster*. This is a cluster that does not have dependencies in AD DS for network names. When you deploy clusters in detached mode, the cluster network name and the network names for clustered roles are registered in a local domain name system (DNS), but corresponding computer objects for a cluster and clustered roles are not created in AD DS.

Cluster nodes still have to be joined to the same AD DS domain, but the person who creates a cluster does not need to have permission to create new objects in AD DS. Also, management of these computer objects is not needed.

When you deploy AD DS-detached clusters, side effects occur. Because computer objects are not created, you cannot use Kerberos authentication when you access cluster resources. Although Kerberos authentication is used between cluster nodes because they have their computer accounts and objects created outside the cluster, Windows NT LAN Manager (NTLM) authentication is used. Because of this, we do not recommend that you deploy AD DS-detached clusters for any scenario that requires Kerberos authentication.

To create an AD DS-detached cluster, you must run Windows Server 2012 R2 on all cluster nodes. These features cannot be configured by using the Failover Cluster Manager, so you must use Windows PowerShell.

Failover Cluster Components

A *failover cluster* is a group of independent computers that work together to increase the availability of applications and services. Physical cables and software connect the clustered servers, known as *nodes*. If one cluster nodes fails, another node begins to provide service. This process is known as *failover*. With failover, users experience a minimum of service disruptions.

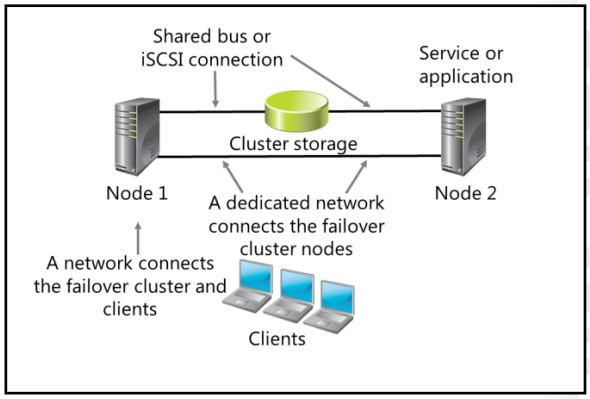

A failover clustering solution consists of several components, which include:

- Nodes. These are computers that are members of a failover cluster. These computers run cluster services, resources, and applications associated with a cluster.
- Network. This is a network across which cluster nodes can communicate with one another and with clients. There are three types of networks that can be used in a cluster. These networks are discussed in more detail in the Failover Cluster Networks topic.

- Resource. This is an entity that is hosted by a node. It is managed by the Cluster service and can be started, stopped, and moved to another node.
- Cluster storage. This is a storage system that is usually shared between cluster nodes. In some scenarios, such as clusters of servers running Microsoft® Exchange Server, shared storage is not required.
- Clients. These are computers (or users) that are using the Cluster service.
- Service or application. This is a software entity that is presented to clients and used by clients.
- Witness. This can be a file share or disk that is used to maintain quorum. Ideally, the witness should be located in a network that is both logically and physically separate from those used by the failover cluster. However, the witness must remain accessible by all cluster node members. The concepts of quorum and how the witness comes into play will be examined more closely in the following lessons.

In a failover cluster, each node in the cluster:

- Has full connectivity and communication with the other nodes in the cluster.
- Is aware when another node joins or leaves the cluster.
- Is connected to a network through which client computers can access the cluster.
- Is connected through a shared bus or iSCSI connection to shared storage.
- Is aware of the services or applications that are running locally, and the resources that are running on all other cluster nodes.

Cluster storage usually refers to logical devices—typically hard disk drives or logical unit numbers (LUNs)—to which all the cluster nodes attach, through a shared bus. This bus is separate from the bus that contains the system and boot disks. The shared disks store resources such as applications and file shares that the cluster will manage.

A failover cluster typically defines at least two data communications networks: one network enables the cluster to communicate with clients, and the second, isolated network enables the cluster node members to communicate directly with one another. If a directly-connected shared storage is not being used, then a third network segment (for iSCSI or Fibre Channel) can exist between the cluster nodes and a data storage network.

Most clustered applications and their associated resources are assigned to one cluster node at a time. The node that provides access to those cluster resources is the active node. If the nodes detect the failure of the active node for a clustered application, or if the active node is taken offline for maintenance, the clustered application is started on another cluster node. To minimize the impact of the failure, client requests are redirected immediately and transparently to the new cluster node.

What Are CSVs?

In a classic failover cluster deployment, only a single node at a time controls a LUN or a volume on the shared storage. This means that the other nodes cannot see shared storage until each node becomes an active node. CSV is a technology introduced in Windows Server 2008 R2 that enables multiple nodes to concurrently share a single LUN. Each node obtains exclusive access to individual files on the LUN instead of the entire LUN. In other words, CSVs provide a distributed file access solution so that multiple nodes in the cluster can simultaneously access the same NTFS file system.

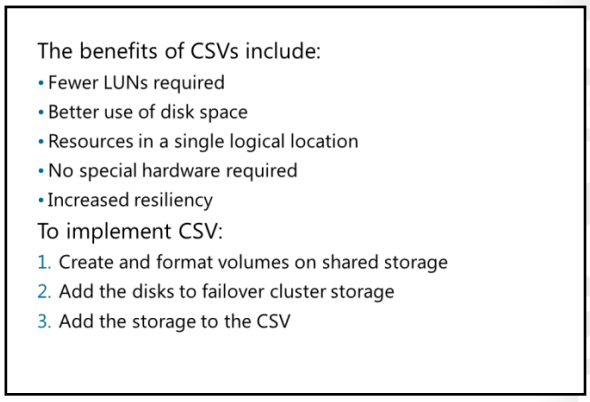

In Windows Server 2008 R2, CSVs were designed only for hosting virtual machines running on a Hyper-V server in a failover cluster. This enabled administrators to have a single LUN that hosts multiple virtual machines in a failover cluster. Multiple cluster nodes have access to the LUN, but each virtual machine runs only on one node at a time. If the node on which the virtual machine was running fails, CSV lets the virtual machine restart on a different node in the failover cluster. Additionally, this provides simplified disk management for hosting virtual machines compared to each virtual machine requiring a separate LUN.

In Windows Server 2012, CSVs have been further enhanced. You now can use CSVs for other roles, not just Hyper-V. For example, you can now configure the file server role in a failover cluster in a Scale-Out File Server scenario. The Scale-Out File Server is designed to provide scale-out file shares that are continuously available for file-based server application storage. Scale-out file shares provide the ability to share the same folder from multiple nodes of the same cluster, thereby avoiding a single point of failure that can occur from the hardware failure on one of the nodes. In this context, CSVs in Windows Server 2012 introduces support for a read cache, which can significantly improve performance in certain scenarios. In addition, a CSV File System (CSVFS) can perform CHKDSK without affecting applications with open handles on the file system.

Other important improvements in CSVs in Windows Server 2012 include:

- CSVFS benefits. In Disk Management, CSVs now appear as CSVFS. However, this is not a new file system. The underlying technology is still the NTFS file system, and CSVFS volumes are still formatted with NTFS. However, because volumes appear as CSVFS, applications can discover that they run on CSVs, which helps improve compatibility. And because of a single file namespace, all files have the same name and path on any node in a cluster.
- Multi-subnet support for CSVs. CSVs have been enhanced to integrate with SMB Multichannel to help achieve faster throughput for CSVs.
- Support for BitLocker Drive Encryption. Windows Server 2012 support BitLocker volume encryption for both traditional clustered disks and CSVs. Each node performs decryption by using the computer account for the cluster itself.
- Support for SMB 3.0 and higher storage. CSVs in Windows Server 2012 provide support for SMB 3.0 storage for Hyper-V and applications such as Microsoft® SQL Server.
- Integration with SMB Multichannel and SMB Direct. This enables CSV traffic to stream across multiple networks in the cluster and to take advantage of network adapters that support Remote Direct Memory Access.

- Integration with the Storage Spaces feature in Windows Server 2012. This can provide virtualized storage on clusters of inexpensive disks.
- Ability to scan and repair volumes. CSVs in Windows Server 2012 support the ability to scan and repair volumes with zero offline time.

Implementing CSVs

You can configure a CSV only when you create a failover cluster. After you create the failover cluster, you can enable the CSV for the cluster, and then add storage to the CSV.

Before you can add storage to the CSV, the LUN must be available as shared storage to the cluster. When you create a failover cluster, all of the shared disks configured in Server Manager are added to the cluster, and you can add them to a CSV. If you add more LUNs to the shared storage, you must first create volumes on the LUN, add the storage to the cluster, and then add the storage to the CSV.

As a best practice, you should configure CSV before you make any virtual machines highly available. However, you can convert from regular disk access to CSV after deployment. The following considerations apply:

- When you convert from regular disk access to CSV, the LUN's drive letter or mount point is removed. This means that you must re-create all virtual machines that are stored on the shared storage. If you must retain the same virtual machine settings, consider exporting the virtual machines, switching to CSV, and then importing the virtual machines in Hyper-V.
- You cannot add shared storage to CSV if it is in use. If you have a running virtual machine that is using a cluster disk, you must shut down the virtual machine, and then add the disk to CSV.

For more information on SMB, go to:
http://go.microsoft.com/fwlink/?linkID=269659

For more information on Storage Spaces, go to:
http://go.microsoft.com/fwlink/?linkID=269680

New CSV Features in Windows Server 2012 R2

In Windows Server 2012 R2, CSVs are further improved. The following topic describes the important improvements for CSV in Windows Server 2012 R2, which include:

- CSVs in Windows Server 2012 R2 provide the following enhancements and new functionalities:
  - Optimized CSV placement policies
  - Increased CSV resiliency
  - CSV cache allocation
  - CSV diagnosis
  - CSV interoperability

Optimized CSV Placement Policies

In a failover cluster for Windows Server 2012, one node in the cluster is designated as the coordinator for a CSV, and there is no automatic rebalance of this designation. The node that is the coordinator for CSV owns the physical disk resource that is associated with a LUN. All I/O operations that are specific to the file system are performed through the coordinator node. In Windows Server 2012 R2, CSV ownership now is distributed evenly between cluster nodes. This distribution is performed based on the number of CSVs that each node owns. A failover cluster service automatically performs a rebalance in scenarios when a node rejoins a cluster, when you add a new cluster, or when you restart a cluster node.

Increased CSV Resiliency

CSV in Windows Server 2012 uses SMB as a transport for I/O forwarding between nodes in a cluster. SMB uses a Server service on cluster nodes, and if this service becomes unavailable, it can result in decreased performance or the ability to access storage. Windows Server 2012 R2 implements multiple instances of Server service, which improves the resilience and scalability of inter-node SMB traffic. The default instance of Server service now accepts clients that access regular file shares, and a second CSV instance handles only inter-node CSV traffic. Also, if Server service becomes unhealthy on one cluster node, CSV ownership can be transitioned to another node automatically to ensure greater resiliency.

CSV Cache Allocation

CSV cache enables the server to use RAM memory as a cache for write-through operations, which improves performance. In Windows Server 2012, CSV cache is disabled by default, and when enabled, you can allocate up to 20 percent of total random access memory (RAM) for cache. In Windows Server 2012 R2, you can allocate up to 80 percent of memory for CSV cache, which enables you to achieve performance gains for the clustered server role. This is especially useful for Scale-Out File Server clusters. In deployments where a Hyper-V cluster is running on a Scale-Out File Server cluster, we recommend that you enable and use the CSV cache, but with greater allocation for a Scale-Out File Server deployment to achieve maximum performance of the virtual machines stored on the file servers.

 **Note:** In Windows Server 2012 R2, the name of the private property of the cluster physical disk resource has been changed from **CsvEnableBlockCache** to **EnableBlockCache**.

CSV Diagnosis

In Windows Server 2012 R2, you now can see the state of CSV on a per-node basis. For example, you can see whether I/O is direct or redirected, or whether the CSV is unavailable. If a CSV is in I/O redirected mode, you can view the reason. This information can be retrieved by using the Windows PowerShell cmdlet **Get-ClusterSharedVolumeState** with the parameters **StateInfo**, **FileSystemRedirectedIOReason**, or **BlockRedirectedIOReason**. This provides you with a better view of how CSV works across cluster nodes.

CSV Interoperability

CSVs in Windows Server 2012 R2 also support interoperability with the following technologies:

- Resilient File System (ReFS)
- Data Deduplication
- Parity storage spaces
- Tiered storage spaces
- Storage Spaces write-back caching

This added support expands the scenarios in which you can use CSVs, and enables you to take advantage of the efficiencies that are introduced in these features.

What Are Failover and Failback?

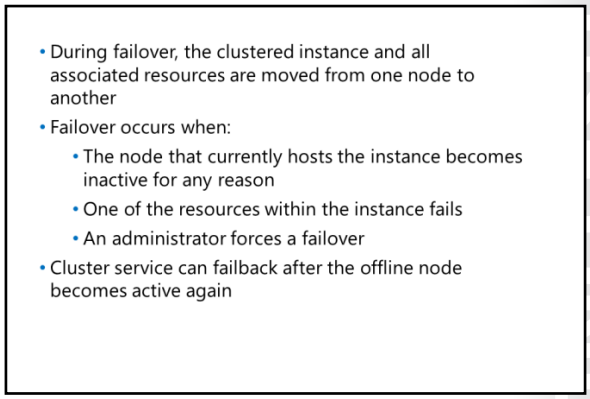

Failover transfers the responsibility of providing access to resources in a cluster from one node to another. Failover can occur when an administrator intentionally moves resources to another node for maintenance, or when unplanned downtime of one node happens because of hardware failure or other reasons. In addition, service failure on an active node can initiate failover to another node.

A failover attempt consists of the following steps:

1. The Cluster service takes all the resources in the instance offline in an order that is determined by the instance's dependency hierarchy. That is, dependent resources first, followed by the resources on which they depend. For example, if an application depends on a physical disk resource, the Cluster service takes the application offline first, which enables the application to write changes to the disk before the disk is taken offline.
2. After all the resources are offline, the Cluster service attempts to transfer the instance to the node that is listed next on the instance's list of preferred owners.
3. If the Cluster service successfully moves the instance to another node, it attempts to bring all the resources online. This time, it starts at the lowest part of the dependency hierarchy. Failover is complete when all the resources are online on the new node.

The Cluster service can *failback* instances that were originally hosted on the offline node, after the offline node becomes active again. When the Cluster service fails back an instance, it uses the same procedures that it performs during failover. That is, the Cluster service takes all the resources in the instance offline, moves the instance, and then brings all the resources in the instance back online.

What Is Quorum?

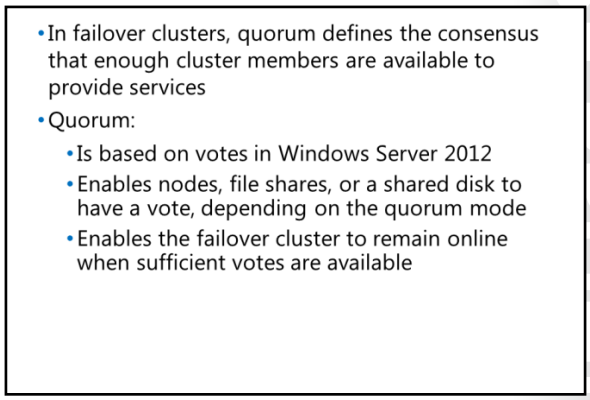

Quorum is the number of elements that must be online for a cluster to continue running. In effect, each element can cast one vote to determine whether the cluster continues to run. Each cluster node is an element that has one vote. In case there is an even number of nodes, then an additional element, which is known as a *witness*, is assigned to the cluster. The witness element can be either a disk or a file share. Each voting element contains a copy of the cluster configuration; and the Cluster service works to keep all copies synchronized at all times.

The cluster will stop providing failover protection if most of the nodes fail, or if there is a problem with communication between the cluster nodes. Without a quorum mechanism, each set of nodes could continue to operate as a failover cluster. This results in a partition within the cluster.

Quorum prevents two or more nodes from concurrently operating a failover cluster resource. If a clear majority is not achieved between the node members, then the vote of the witness becomes crucial to maintain the validity of the cluster. Concurrent operation could occur when network problems prevent

one set of nodes from communicating with another set of nodes. That is, a situation might occur in which more than one node tries to control access to a resource. If that resource is, for example, a database application, damage could result. Imagine the consequence if two or more instances of the same database are made available on the network, or if data was accessed and written to a target from more than one source at a time. If the application itself is not damaged, the data could easily become corrupted.

Because a given cluster has a specific set of nodes and a specific quorum configuration, the cluster can calculate the number of votes that are required for the cluster to continue providing failover protection. If the number of votes drops below the majority, the cluster stops running. That means that it will not provide failover protection if there is a node failure. Nodes will still listen for the presence of other nodes, in case another node appears again on the network, but the nodes will not function as a cluster until a majority consensus or quorum is achieved.

 **Note:** The full functioning of a cluster depends not just on quorum, but also on the capacity of each node to support the services and applications that fail over to that node. For example, a cluster that has five nodes could still have quorum after two nodes fail, but each remaining cluster node would continue serving clients only if it has enough capacity, such as disk space, processing power, network bandwidth, or RAM, to support the services and applications that failed over to it. An important part of the design process is to plan each node's failover capacity. A failover node must be able to run its own load and also the load of additional resources that might failover to it.

The Process of Achieving Quorum

Because a given cluster has a specific set of nodes and a specific quorum configuration, the cluster software on each node stores information about how many votes constitute a quorum for that cluster. If the number drops below the majority, the cluster stops providing services. Nodes will continue to listen for incoming connections from other nodes on port 3343, in case they appear again on the network, but the nodes will not begin to function as a cluster until quorum is achieved.

A cluster must complete several phases to achieve quorum. As a given node comes up, it determines whether there are other cluster members with which communicates. This process may be in progress on multiple nodes at the same time. After communication is established with other members, the members compare their membership views of the cluster until they agree on one view, which is based on timestamps and other information. A determination is made whether this collection of members has a quorum; or has enough members so that the total creates sufficient votes to avoid a split scenario. A *split* scenario means that another set of nodes that are in this cluster are running on a part of the network inaccessible to these nodes. Therefore, more than one node could be actively trying to provide access to the same clustered resource. If there are not enough votes to achieve quorum, the voters (the currently recognized members of the cluster) wait for more members to appear. After at least the minimum vote total is attained, the Cluster service begins to bring cluster resources and applications into service. With quorum attained, the cluster becomes fully functional.

Quorum Modes in Windows Server 2012 Failover Clustering

The same quorum modes from Windows Server 2008 are also present in Windows Server 2012. As before, a majority of votes determines whether a cluster achieves quorum. Nodes can vote, and where appropriate, either a disk in cluster storage, known as a disk witness, or a file share, known as a file share witness, can vote. There also is a quorum mode called No Majority: Disk Only, which functions like the disk-based quorum in Windows Server® 2003. Other than that mode, there is no single point of failure with the quorum modes, because only the number of votes is important and not whether a particular element is available to vote.

Quorum mode	What has the vote?	When is quorum maintained?
• Node Majority	Only nodes in the cluster have a vote	Quorum is maintained when more than half of the nodes are online
• Node and Disk Majority	The nodes in the cluster and a disk witness have a vote	Quorum is maintained when more than half of the votes are online
• Node and File Share Majority	The nodes in the cluster and a file share witness have a vote	Quorum is maintained when more than half of the votes are online
• No Majority: Disk Only	Only the quorum-shared disk has a vote	Quorum is maintained when the shared disk is online

This quorum mode is flexible. You can choose the mode best suited to your cluster.

Be aware that, most of the time, it is best to use the quorum mode selected by the cluster software. If you run the Quorum Configuration Wizard, the quorum mode that the wizard lists as recommended is the quorum mode chosen by the cluster software. We recommend that you change the quorum configuration only if you have determined that the change is appropriate for your cluster.

The four quorum modes are:

- Node Majority. Each node that is available and in communication can vote. The cluster functions only with a majority, or more than half of the votes. This model is preferred when the cluster consists of an odd number of server nodes (no witness is needed to maintain or achieve quorum).
- Node and Disk Majority. Each node plus a designated disk in the cluster storage can vote. The disk witness can vote, when it is available and in communication. The cluster functions only with a majority (more than half) of the votes. This model is based on an even number of server nodes being able to communicate with one another in the cluster, in addition to the disk witness.
- Node and File Share Majority. Each node plus a designated file share created by the administrator, which is the file share witness, can vote when they are available and in communication. The cluster functions only with a majority of the votes. This model is based on an even number of server nodes being able to communicate with one another in the cluster, in addition to the file share witness.
- No Majority: Disk Only. The cluster has quorum if one node is available and in communication with a specific disk in the cluster storage. Only the nodes that are also in communication with that disk can join the cluster.

Except for the No Majority: Disk Only mode, all quorum modes in Windows Server 2012 failover clusters are based on a simple-majority vote model. As long as a majority of the votes are available, the cluster continues to function. For example, if there are five votes in the cluster, the cluster continues to function if at least three votes are available. The source of the votes is not relevant—the vote could be a node, a disk witness, or a file share witness. The cluster will stop functioning if a majority of votes is not available.

In the No Majority: Disk Only mode, the quorum-shared disk can veto all other possible votes. In this mode, the cluster will continue to function as long as the quorum-shared disk and at least one node are available. This type of quorum also prevents more than one node from assuming the primary role.

Note: If the quorum-shared disk is not available, the cluster will stop functioning, even if all nodes are still available. In this mode, the quorum-shared disk is a single point of failure, so this mode is not recommended.

When you configure a failover cluster in Windows Server 2012, the Installation Wizard automatically selects one of two default configurations. By default, failover clustering selects:

- Node Majority, if there is an odd number of nodes in the cluster.
- Node and Disk Majority, if there is an even number of nodes in the cluster.

Modify this setting only if you determine that a change is appropriate for your cluster, and ensure that you understand the implications of making the change.

In addition to planning your quorum mode, you should also consider the capacity of the nodes in your cluster, and their ability to support the services and applications that may fail over to that node. For example, a cluster that has four nodes and a disk witness still has quorum after two nodes fail. However, if you have several applications or services deployed on the cluster, each remaining cluster node may not have the capacity to provide services.

How Quorum Works in Windows Server 2012 R2 Failover Clustering

In Windows Server 2012 R2, old quorum modes such as Node Majority, Node and Disk Majority, and Node and File Share Witness Majority, are not used anymore. Instead, Windows Server 2012 R2 introduces the concept of *Dynamic Quorum*. This feature provides the ability for a cluster to recalculate quorum in the event of a node failure and still maintain working clustered roles, even when the number of voting nodes remaining in the cluster is less than 50 percent.

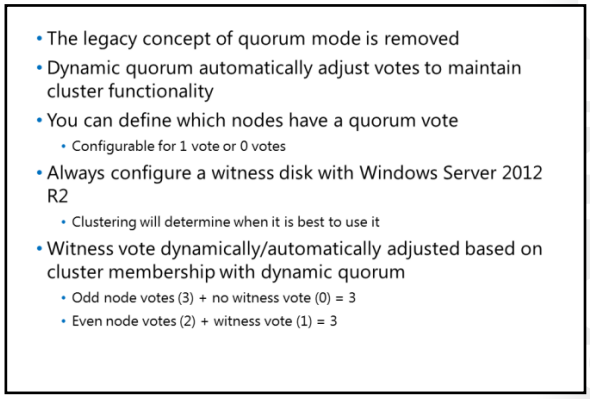

In Windows Server 2012 R2, this feature is enhanced additionally by introducing the concept of *Dynamic Witness*. When you configure a cluster in Windows Server 2012 R2, Dynamic Quorum is selected by default, but witness vote also is adjusted dynamically based on the number of voting nodes in the current cluster membership. For example, if a cluster has an odd number of votes, a quorum witness does not have a vote in the cluster. If the number of nodes is even, a quorum witness does have a vote. If a witness resource is failed or offline for some reason, the cluster will set the witness vote to a value of 0 automatically. By using this approach, the risk of a malfunctioned cluster because of a failing witness is greatly reduced. If you want to see if a witness has a vote, you can use Windows PowerShell and a new cluster property in the following cmdlet:

```
(Get-Cluster).WitnessDynamicWeight
```

A value of 0 indicates that the witness does not have a vote. A value of 1 indicates that the witness has a vote.

The cluster can now decide whether to use the witness vote based on the number of voting nodes that are available in the cluster. As an additional benefit, quorum configuration is much simpler when you create a cluster. Windows Server 2012 R2 configures quorum witness automatically when you create a cluster.

Also, when you added or evict cluster nodes, you no longer have to adjust the quorum configuration manually. The cluster now automatically determines quorum management options and quorum witness.

Force Quorum Resiliency

This feature provides additional support and flexibility to split brain syndrome cluster scenarios. This scenario happens when a cluster breaks into subsets of cluster nodes that are not aware of each other. The cluster node subset that has a majority of votes will run, while others are turned down. This scenario usually happens in multisite cluster deployments. If you want to start cluster nodes that do not have a majority, you can force quorum to start manually by using the **/fq** switch.

In Windows Server 2012 R2, in such scenarios, the cluster will detect partitions in the cluster automatically as soon as connectivity between nodes is restored. The partition that was started by forcing a quorum is considered authoritative, and other nodes rejoin the cluster. When this happens, the cluster is brought back to a single view of membership. In Windows Server 2012, partitioned nodes without quorum did not start automatically, and the administrator had to start them manually with the **/pq** switch. In Windows Server 2012 R2, both sides of the split cluster have a view of the cluster membership, and they will reconcile automatically when connectivity is restored.

Tie Breaker for 50% Node Split

Dynamic quorum in Windows Server 2012 R2 is enhanced with an additional functionality. The cluster can now adjust the running node's vote status automatically to keep the total number of votes in the cluster at an odd number. This is called *tie breaker for 50% node split*, and it works with dynamic witness functionality. You can use the dynamic witness functionality to adjust the value of a quorum witness vote. For example, if you have a cluster with an even number of nodes and a file share witness, if the file share witness fails, the cluster uses the dynamic witness functionality to remove the vote from a file share witness automatically. However, because the cluster now has even number of votes, the cluster tie breaker picks a node randomly and removes it from the quorum vote to maintain an odd number of votes. If the nodes are distributed evenly in two sites, this helps to maintain cluster functionality in one site. In previous Windows Server versions, if both sites have an equal number of nodes and a file share witness fails, both sites stop the cluster.

If you want to avoid the node being picked randomly, you can use the **LowerQuorumPriorityNodeID** property to predetermine which node has its vote removed. You can set this property by using the following Windows PowerShell command, where 1 is the example node ID for a node in the site that you consider less critical:

```
(Get-Cluster).LowerQuorumPriorityNodeID = 1
```

Failover Cluster Networks

Network and network adapters are important parts of each cluster implementation. You cannot configure a cluster without configuring the networks that the cluster will use. A network can perform one of the following roles in a cluster:

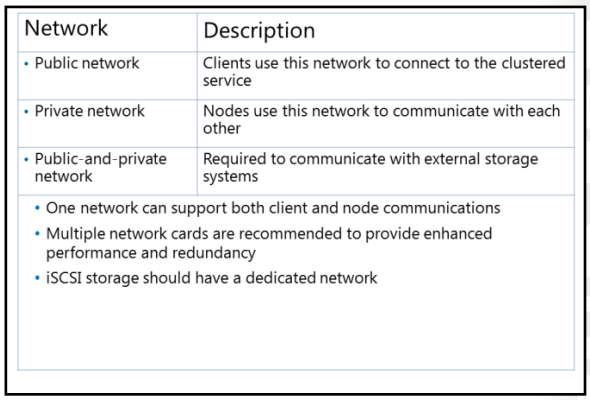

- Private network. A private network carries internal cluster communication. By using this network, cluster nodes exchange heartbeats and check for another node or nodes. The failover cluster authenticates all internal communication. However, administrators who are especially concerned about security may want to restrict internal communication to physically secure networks.
- Public network. A public network provides client systems with access to cluster application services. IP address resources are created on networks that provide clients with access to the Cluster service.
- Public-and-private network. A public-and-private network, also known as a mixed network, carries internal cluster communication and connects clients to cluster application services.

When you configure networks in failover clusters, you must also dedicate a network to connect to the shared storage. If you use iSCSI for the shared storage connection, the network will use an IP-based Ethernet communications network. However, you should not use this network for node or client communication. Sharing the iSCSI network in this manner may result in contention and latency issues for both users and the resource that that the cluster provides.

Although not a best practice, you can use the private and public networks for both client and node communications. Preferably, you should dedicate an isolated network for the private node communication. The reasoning for this is similar to using a separate Ethernet network for iSCSI: to avoid resource bottleneck and contention issues. The public network is configured to allow client connections to the failover cluster. Although the public network can provide backup for the private network, a better design practice is to define alternative networks for the primary private and public networks or at least team the network interfaces used for these networks.

The networking features in Windows Server 2012–based clusters include the following:

- The nodes transmit and receive heartbeats by using User Datagram Protocol (UDP) unicast, instead of UDP broadcast, which was used in legacy clusters. The messages are sent on port 3343.
- You can include clustered servers on different IP subnets, which reduces the complexity of setting up multisite clusters.
- The Failover Cluster Virtual Adapter is a hidden device that is added to each node when you install the failover clustering feature. The adapter is assigned a media access control (MAC) address based on the MAC address that is associated with the first enumerated physical network adapter in the node.
- Failover clusters fully support IPv6 for both node-to-node and node-to-client communication.
- You can use Dynamic Host Configuration Protocol (DHCP) to assign IP addresses, or to assign static IP addresses to all nodes in the cluster. However, if some nodes have static IP addresses and you configure others to use DHCP, the Validate a Configuration Wizard will raise an error. The cluster IP address resources are obtained based on the configuration of the network interface supporting that cluster network.

Failover Cluster Storage

Most failover clustering scenarios require shared storage to provide consistent data to a highly available service or application after failover. The following are three shared-storage options for a failover cluster:

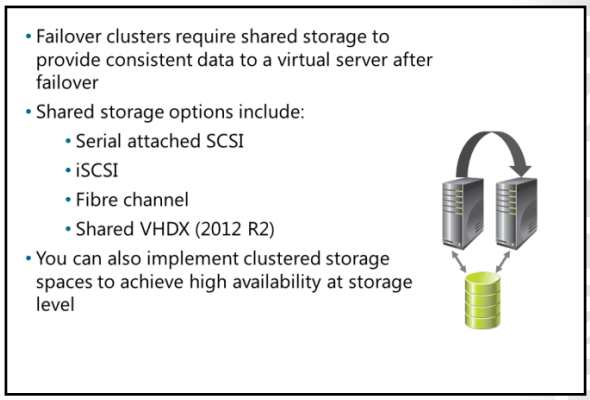

- Shared serial attached SCSI. Shared serial attached SCSI is the lowest-cost option, however, it is not very flexible for deployment because the two cluster nodes must be physically close together. In addition, the shared storage devices that are supporting serial attached SCSI have a limited number of connections for cluster nodes.
- iSCSI. iSCSI is a type of storage area network (SAN) that transmits SCSI commands over IP networks. Performance is acceptable for most scenarios when 1 gigabit per second (Gbps) or 10 Gbps Ethernet is used as the physical medium for data transmission. This type of SAN is fairly inexpensive to implement because no specialized networking hardware is required. In Windows Server 2012, you can implement iSCSI target software on any server, and present local storage over an iSCSI interface to clients.
- Fibre channel. Fibre channel SANs typically have better performance than iSCSI SANs, but are much more expensive. Specialized knowledge and hardware are required to implement a fibre channel SAN.
- Shared VHDX. In Windows Server 2012 R2, you can use a shared virtual hard disk drive as storage for virtual machine guest clustering. A shared virtual hard drive should be located on a CSV or Scale-Out File Server cluster, and it can be added to two or more virtual machines that are participating in a guest cluster, by connecting to the SCSI interface.

Note: The Microsoft iSCSI Software Target is now an integrated feature in Windows Server 2012. It can provide storage from a server over a TCP/IP network, including shared storage for applications that are hosted in a failover cluster. In addition, in Windows Server 2012, a highly available iSCSI Target Server can be configured as a clustered role by using the Failover Cluster Manager or Windows PowerShell.

In Windows Server 2012 R2, you can use failover clustering to provide high availability for the storage, in addition to using storage as a cluster component. This is done by implementing clustered storage spaces. When you implement clustered storage spaces, you help to protect your environment from risks such as physical disk failures, data access failures, data corruptions, volume unavailability, and server node failures.

Deploy Clustered Storage Spaces
http://go.microsoft.com/fwlink/?LinkID=386644

Storage Requirements

After you select a storage type, you should also be aware of the following storage requirements:

- To use the native disk support included in failover clustering, use basic disks, not dynamic disks.
- We recommend that you format the partitions with NTFS. For the disk witness, the partition must be NTFS, because file allocation table (FAT) is not supported.

- For the partition style of disk, you can use either a master boot record (MBR) or a GUID partition table (GPT).
- Because improvements in failover clusters require that the storage respond correctly to specific SCSI commands, the storage must follow the SCSI Primary Commands-3 (SPC-3) standard. In particular, the storage must support persistent reservations, as specified in the SPC-3 standard.
- The miniport driver used for the storage must work with the Microsoft Storport storage driver. Storport offers a higher performance architecture and better Fibre Channel compatibility in Windows systems.
- You must isolate storage devices, in a ratio of one cluster per device. Servers from different clusters must be unable to access the same storage devices. In most cases, a LUN that is used for one set of cluster servers should be isolated from all other servers through LUN masking or zoning.
- Consider using Multipath I/O (MPIO) software. In a highly available storage fabric, you can deploy failover clusters with multiple host bus adapters by using MPIO software. This provides the highest level of redundancy and availability. For Windows Server 2012, your multipath solution must be based on MPIO. Your hardware vendor usually supplies an MPIO device-specific module (DSM) for your hardware, although Windows Server 2012 includes one or more DSMs as part of the operating system.
- If you use a shared virtual hard disk drive, you must have a separate cluster with CSV or a file server cluster to store the virtual hard disk drive.

Lesson 2
Implementing a Failover Cluster

Failover clusters in Windows Server 2012 have specific recommended hardware and software configurations that enable Microsoft to support the cluster. Failover clusters are intended to provide a higher level of service than stand-alone servers. Therefore, cluster hardware requirements are frequently stricter than requirements for stand-alone servers.

This lesson describes how to prepare for cluster implementation, and also discusses the hardware, network, storage, infrastructure, and software requirements for Windows Server 2012 failover clusters. This lesson also outlines the steps for using the Validate a Configuration Wizard to ensure correct cluster configuration.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe how to prepare for implementing failover clustering.
- Describe hardware requirements for failover clustering.
- Describe network requirements for failover cluster implementation.
- Describe infrastructure requirements for failover clustering.
- Describe software requirements for failover cluster implementation.
- Describe how to validate and configure a failover cluster.
- Describe how to migrate and upgrade failover clusters.

Preparing for Failover Cluster Implementation

Before you implement failover clustering technology, you must identify services and applications that you want to make highly available. Failover clustering cannot be applied to all applications. In addition, you should be aware that failover clustering does not provide improved scalability by adding nodes. You can only obtain scalability by scaling up and using more powerful hardware for the individual nodes. Therefore, you should only use failover clustering when your goal is high availability, instead of scalability.

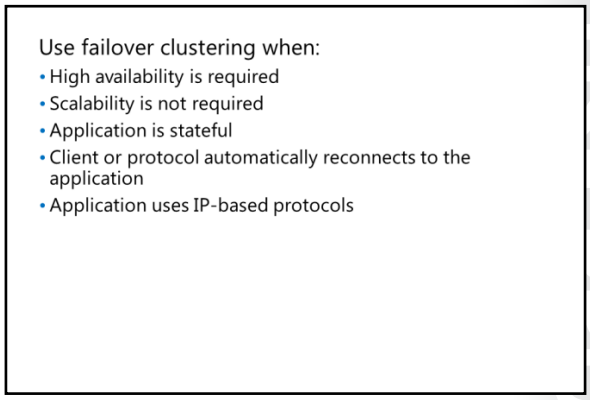

Failover clustering is best suited for stateful applications that are restricted to a single set of data. A database is one example of such an application. Data is stored in a single location, and can only be used by one database instance. You can also use failover clustering for Hyper-V virtual machines.

Failover clustering uses only IP-based protocols; therefore, it is suited only to IP-based applications. Both IP version 4 (IPv4) and IP version 6 (IPv6) are supported.

The best results for failover clustering occur when the client or network protocol can reconnect to the application automatically after failover. If the client does not reconnect automatically, then the user must restart the client application.

Consider the following guidelines when you plan node capacity in a failover cluster:

- Spread out the highly available applications from a failed node. When all nodes in a failover cluster are active, the highly available services or applications from a failed node should be spread out among the remaining nodes to prevent a single node from being overloaded.
- Ensure that each node has sufficient idle capacity to service the highly available services or applications that are allocated to it when another node fails. This idle capacity should be a sufficient buffer to avoid nodes running at near capacity after a failure event. Failure to adequately plan resource utilization can result in decrease in performance following node failure.
- Use hardware with comparable capacity for all nodes in a cluster. This simplifies the planning process for failover because the failover load will be evenly distributed among the surviving nodes.
- Use standby servers to simplify capacity planning. When a passive node is included in the cluster, then all highly available services or applications from a failed node can be failed over to the passive node. This avoids the need for complex capacity planning. If this configuration is selected, it is important that the standby server has sufficient capacity to run the load from more than one node failure.

You also should examine all cluster configuration components to identify single points of failure. You can remedy many single points of failure with simple solutions, such as adding storage controllers to separate and stripe disks, teaming network adapters, or using multipath software. These solutions reduce the probability that a failure of a single device causing a failure in the cluster. Typically, server class computer hardware has options for multiple power supplies for power redundancy, and for creating redundant array of independent disks (RAID) sets for disk data redundancy.

Hardware Requirements for Failover Cluster Implementation

It is very important to make good decisions when you select hardware for cluster nodes. Failover clusters must satisfy the following criteria to meet availability and support requirements:

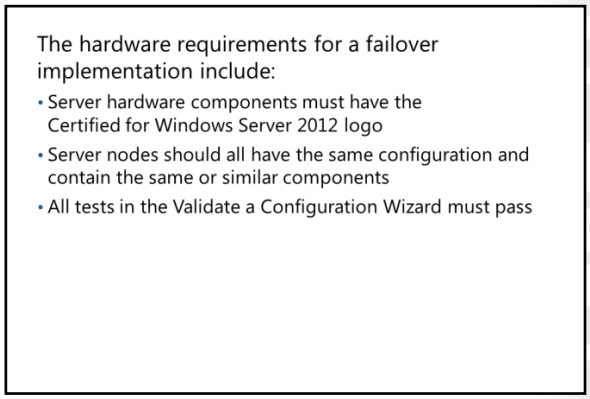

- All hardware that you select for a failover cluster should meet the "Certified for Windows Server 2012" logo requirements. Hardware that has this logo was independently tested to meet the highest technical bar for reliability, availability, stability, security, and platform compatibility. This means that official support options exist in case malfunctions arise.
- You should install the same or similar hardware on each failover cluster node. For example, if you choose a specific model of network adapter, you should install this adapter on each cluster node.
- If you are using serial attached SCSI or Fibre Channel storage connections, the mass-storage device controllers that are dedicated to the cluster storage should be identical in all clustered servers. The controllers should also use the same firmware version.
- If you use iSCSI storage connections, each clustered server should have one or more network adapters or host bus adapters dedicated to the cluster storage. The network that you use for iSCSI storage connections should not be used for network communication. In all clustered servers, the network adapters that you use to connect to the iSCSI storage target should be identical, and we recommend that you use Gigabit Ethernet or more.

- After you configure the servers with the hardware, all tests provided in the Validate a Configuration Wizard must pass before the cluster is considered a configuration that is supported by Microsoft.

Network Requirements for Failover Cluster Implementation

Failover cluster network components must have the Certified for Windows Server 2012 logo, and also must pass the tests in the Validate a Configuration Wizard. Additional requirements include the following:

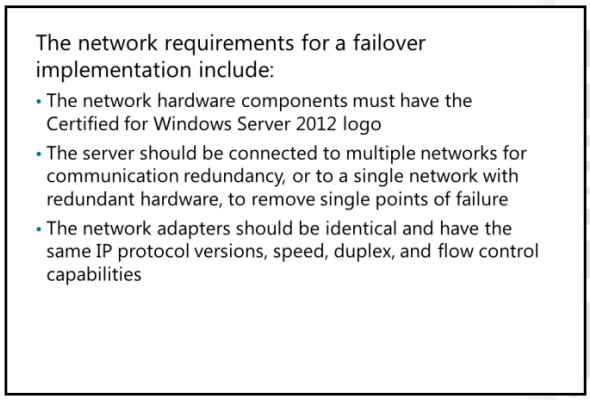

- The network adapters in each node should be identical and have the same IP protocol version, speed, duplex, and flow control capabilities that are available.
- The networks and network equipment to which you connect the nodes should be redundant, so that even a single failure allows for the nodes to continue communicating with one another. You can use network adapter teaming to provide single network redundancy. We recommend multiple networks to provide multiple paths between nodes for inter-node communication; otherwise, a warning will generate during the validation process.
- The network adapters in a cluster network must have the same IP address assignment method, which means either that they all use static IP addresses or that they all use DHCP.

Note: If you connect cluster nodes with a single network, the network passes the redundancy requirement in the Validate a Configuration Wizard. However, the report from the wizard includes a warning that the network should not have single points of failure.

Infrastructure Requirements for Failover Cluster

Failover clusters depend on infrastructure services. Each server node must be in the same Active Directory domain; and if you use DNS, the nodes should use the same DNS servers for name resolution.

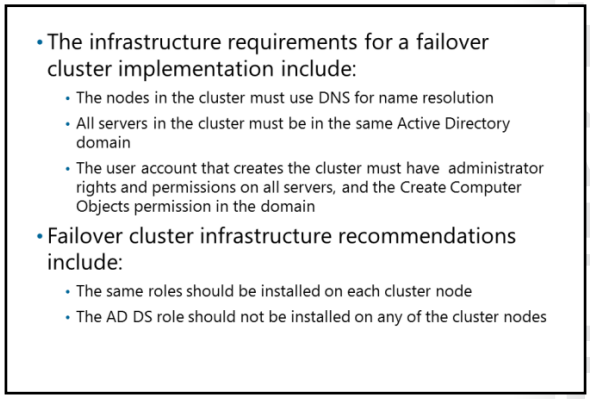

We recommend that you install the same Windows Server 2012 features and roles on each node. Inconsistent configuration on cluster nodes can cause instability and performance issues. In addition, you should not install the AD DS role on any of the cluster nodes because AD DS has its own fault-tolerance mechanism. If you install the AD DS role on one of the nodes, you must install it on all nodes.

You must have the following network infrastructure for a failover cluster:

- Network settings and IP addresses. When you use identical network adapters for a network, also use identical communication settings on those adapters such as speed, duplex mode, flow control, and media type. Also compare the settings between the network adapter and the switch to which it connects, and ensure that no settings are in conflict. Otherwise, network congestion or frame loss might occur that could adversely affect how the cluster nodes communicate among themselves, with clients, or with storage systems.
- Unique subnets. If you have private networks that are not routed to the rest of the network infrastructure, ensure that each of these private networks uses a unique subnet. This is necessary even if you give each network adapter a unique IP address. For example, if you have a cluster node in a central office that uses one physical network, and another node in a branch office that uses a separate physical network; do not specify 10.0.0.0/24 for both networks, even if you give each adapter a unique IP address. This avoids routing loops and other network communications problems if, for example, the segments are accidentally configured into the same collision domain because of incorrect virtual local area network (VLAN) assignments.
- DNS. The servers in the cluster typically use DNS for name resolution. DNS dynamic update protocol is a supported configuration.
- Domain role. All servers in the cluster must be in the same Active Directory domain. As a best practice, all clustered servers should have the same domain role, either member server or domain controller. The recommended role is member server because AD DS inherently includes its own failover protection mechanism.
- Account for administering the cluster. When you first create a cluster or add servers to it, you must be logged on to the domain with an account that has administrator rights and permissions on all servers in that cluster. The account does not have to be a Domain Admins account, but can be a Domain Users account that is in the Administrators group on each clustered server. In addition, if the account is not a Domain Admins account, the account, or the group in which the account is a member, must be given the Create Computer Objects permission in the domain.

In Windows Server 2012, there is no cluster service account. Instead, the Cluster service runs automatically in a special context that provides the specific permissions and credentials that are necessary for the service, which is similar to the local system context, but with reduced credentials. When a failover cluster is created and a corresponding computer object is created in AD DS, that object is configured to prevent accidental deletion. In addition, the cluster Network Name resource has additional health check logic, which periodically checks the health and properties of the computer object that represents the Network Name resource.

Software Requirements for Failover Cluster Implementation

Failover clusters require that each cluster node must run the same edition of Windows Server 2012. The edition can be either Windows Server 2012 Standard or Windows Server 2012 Datacenter. The same applies for Windows Server 2012 R2. The nodes should also have the same software updates and service packs. Depending on the role that will be clustered, a Server Core installation may also meet the software requirements. However, you cannot install Server Core and the full editions in the same cluster.

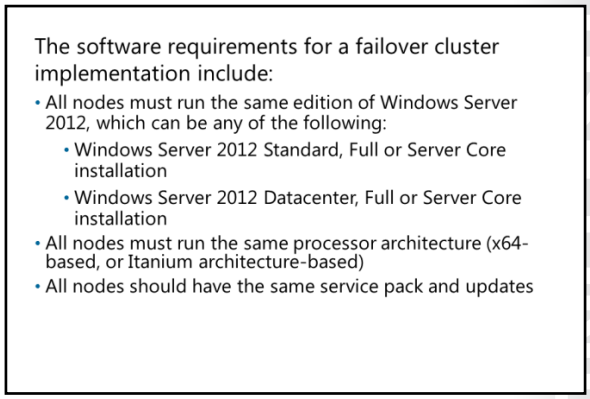

It is also very important that the same version of service packs or any operating system updates exist on all nodes that are parts of a cluster.

Note: Windows Server 2012 provides CAU technology that can help you maintain updates on cluster nodes. This feature will be discussed in more detail in *Lesson 4: Maintaining a Failover Cluster*.

Each node must run the same processor architecture. This means that each node must have the same processor family, which might be the Intel Xeon processor family with Extended Memory 64 Technology, the AMD Opteron AMD64 family, or the Intel Itanium-based processor family.

Demonstration: Validating and Configuring a Failover Cluster

The Validate a Configuration Wizard runs tests that confirm if the hardware and hardware settings are compatible with failover clustering. By using the wizard, you can run the complete set of configuration tests or a subset of the tests. We recommend that you run the tests on servers and storage devices before you configure the failover cluster, and again after any major changes are made to the cluster. You can access the test results in the %*Windir*%\cluster\Reports directory.

Demonstration Steps

1. Start the Failover Cluster Manager on the **LON-SVR3** virtual machine.
2. Start the Validate Configuration Wizard. Add LON-SVR3 and LON-SVR4 as cluster nodes.
3. Review the report.
4. Create a new cluster. Add LON-SVR3 and LON-SVR4 as cluster nodes.
5. Name the cluster as **Cluster1**.
6. Use **172.16.0.125** as **IP address**.

Migrating and Upgrading Failover Clusters

In certain scenarios, such as replacing cluster nodes or upgrading to a newer version of a Windows operating system, you will need to migrate clustered roles or services from one cluster to another. In Windows Server 2012, it is possible to migrate clustered roles and cluster configuration from clusters running Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008. You can migrate these roles and configurations in one of two ways:

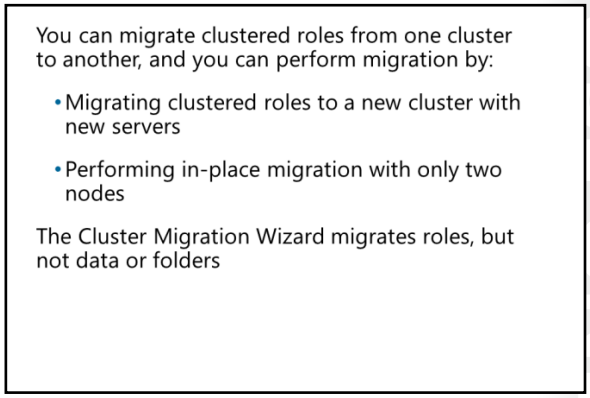

- Migrate from an existing cluster to a new cluster that is running Windows Server 2012. In this scenario, you have two new cluster nodes running Windows Server 2012, and you then perform migration from an existing cluster with nodes running Windows Server 2008 or newer.
- Perform an in-place migration on a two-node cluster. This is a more complex scenario, where you want to migrate a cluster to a newer version of the Windows operating system. In this scenario, you do not have additional computers for new cluster nodes. For example, you may want to upgrade a cluster that is currently running on Windows Server 2008 R2 to a cluster that is running Windows Server 2012. To achieve this, you must first remove resources from one node, and then evict that node from a cluster. Next, you perform a clean installation of Windows Server 2012 on that server. After Windows Server 2012 is installed, you create a one-node failover cluster, migrate the clustered services and applications from the old cluster node to that failover cluster, and then remove the old node from cluster. The last step is to install Windows Server 2012 on another cluster node, together with failover cluster feature, and add the server to the failover cluster. Then you run validation tests to confirm that the overall configuration works correctly.

The Cluster Migration Wizard is a tool that enables you to migrate clustered roles. Because the Cluster Migration Wizard does not copy data from one storage location to another, you must copy or move data or folders (including shared folder settings) during a migration. In addition, the Cluster Migration Wizard does not migrate mount-point information (information about hard disk drives that do not use drive letters and are mounted in a folder on another hard disk drive). However, it can migrate physical disk resource settings to and from disks that use mount points.

Lesson 3

Configuring Highly Available Applications and Services on a Failover Cluster

After you have configured a failover clustering infrastructure, you should configure specific role or service to be highly available. Not all roles can be clustered. Therefore, you should first identify the resource that you want to put in a cluster, and check whether it is supported. In this lesson, you will learn about configuring roles and applications in failover clusters, and about configuring failover cluster settings.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe and identify cluster resources and services.
- Describe the process for clustering server roles.
- Configure a cluster role.
- Describe failover cluster management tasks.
- Describe how to manage cluster nodes.
- Describe how to configure application failover settings.

Identifying Cluster Resources and Services

A clustered service that contains an IP address resource and a Network Name resource (and other resources) is published to a client on the network under a unique server name. Because this group of resources is displayed as a single logical server to clients, it is called a cluster instance.

- Clustered services:
  - Are services or applications that are made highly available by installing them on a failover cluster
  - Are active on one node, but can be moved to another node
- Resources:
  - Are the components that make up a clustered service
  - Are moved to another node when one node fails
  - Can only run on one node at a time
  - Include components such as shared disks, names, and IP addresses

Users access applications or services on an instance in the same manner they would if the applications or services were on a non-clustered server. Usually, applications or users do not know that they are connecting to a cluster; nor do they know the node to which they are connected.

Resources are physical or logical entities, such as a file share, disk, or IP address that the failover cluster manages. Resources may provide a service to clients or may be an important part of the cluster. Resources are the most basic and smallest configurable unit. At any time, a resource can run only on a single node in a cluster, and it is online on a node when it provides its service to that specific node.

Server Cluster Resources

A cluster resource is any physical or logical component that has the following characteristics:

- It can be brought online and taken offline.
- It can be managed in a server cluster.
- It can be hosted (owned) by only one node at a time.

To manage resources, the Cluster service communicates to a resource DLL through a resource monitor. When the Cluster service makes a request of a resource, the resource monitor calls the appropriate entry-point function in the resource DLL to check and control the resource state.

Dependent Resources

A dependent resource is one that requires another resource to operate. For example, a network name must be associated with an IP address. Because of this requirement, a network name resource depends on an IP address resource. Dependent resources are taken offline before the resources upon which they depend are taken offline; similarly, they are brought online after the resources on which they depend are brought online. A resource can specify one or more resources on which it is dependent. Resource dependencies also determine bindings. For example, clients will be bound to the particular IP address on which a network name resource depends.

When you create resource dependencies, consider the fact that, although some dependencies are strictly required, others are not required but are recommended. For example, a file share that is not a Distributed File System (DFS) root has no required dependencies. However, if the disk resource that holds the file share fails, the file share will be inaccessible to users. Therefore, it is logical to make the file share dependent on the disk resource.

A resource can also specify a list of nodes on which it can run. Possible nodes and dependencies are important considerations when administrators organize resources into groups.

The Process for Clustering Server Roles

Failover clustering supports the clustering of several Windows Server roles, such as File Services, DHCP, and Hyper-V. To implement clustering for a server role, or for external applications such as SQL Server or Exchange Server, perform the following procedure:

1. Install the failover clustering feature
2. Verify the configuration and create a cluster
3. Install the role on all cluster nodes, using Server Manager
4. Create a clustered application by using the Failover Cluster Management snap-in
5. Configure the application
6. Test the failover

1. Install the failover clustering feature. Use Server Manager or Deployment Image Servicing and Management (DISM) to install the failover clustering feature on all computers that will be cluster members.
2. Verify configuration, and create a cluster with the appropriate nodes. Use the Failover Cluster Management snap-in to first validate a configuration, and then create a cluster with selected nodes.
3. Install the role on all cluster nodes. Use Server Manager to install the server role that you want to use in the cluster.
4. Create a clustered application by using the Failover Cluster Management snap-in.
5. Configure the application. Configure options on the application that is used in the cluster.
6. Test failover. Use the Failover Cluster Management snap-in to test failover by intentionally moving the service from one node to another.

After the cluster is created, you can monitor its status by using the Failover Cluster Management console, and manage available options.

Demonstration: Clustering a File Server Role

In this demonstration, you will see how to cluster a file server role.

Demonstration Steps

1. Open the Failover Cluster Manager and verify that three cluster disks are available.
2. Start the Configure Role Wizard and Configure the **File Server** as clustered role.
3. For the Client Access Point, use the name **AdatumFS** and the IP address of **172.16.0.130**.
4. Select **Cluster Disk 2** as the storage for the File Server role.

Failover Cluster Management Tasks

You can perform several failover cluster management tasks. These tasks range from adding and removing cluster nodes to modifying the quorum settings. Some of the most frequently used configuration tasks include:

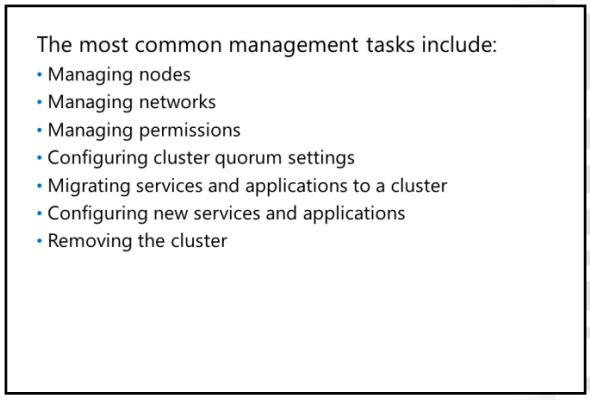

- Managing cluster nodes. For each node in a cluster, you can stop cluster service temporarily, pause it, initiate remote desktop to the node, or evict node from the cluster. You also can choose to drain nodes in the cluster, for example, if you want to perform maintenance or install updates. This functionality is part of the infrastructure that enables CAU for patching nodes in a cluster.
- Managing cluster networks. You can add or remove cluster networks, and you can configure networks that will be dedicated just for inter-cluster communication.
- Managing permissions. By managing permissions, you delegate rights to administer a cluster.
- Configuring cluster quorum settings. By configuring quorum settings, you determine how quorum is achieved as well as who can vote in a cluster.
- Migrating services and applications to a cluster. You can implement existing services to the cluster and make them highly available.
- Configuring new services and applications to work in a cluster. You can implement new services to the cluster.
- Removing a cluster.

You can perform most of these administrative tasks by using the Failover Cluster Management console.

Managing Cluster Nodes

Cluster nodes are mandatory for each cluster. After you create a cluster and put it in to production, you might have to manage cluster nodes occasionally.

There are three aspects of managing cluster nodes:

- You can add a node to an established failover cluster by selecting Add Node in the Failover Cluster Management Actions pane. The Add Node Wizard prompts you for information about the additional node.

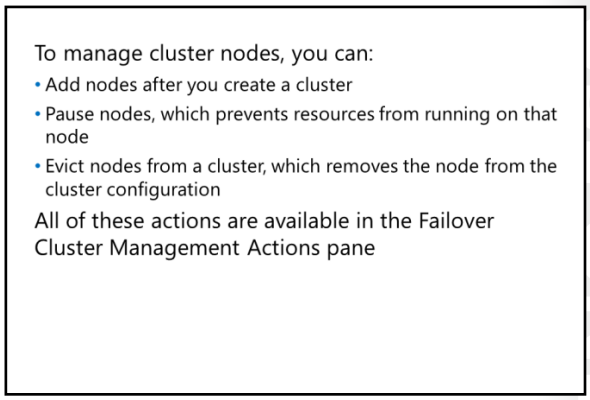

- You can pause a node to prevent resources from being failed over or moved to the node. You typically pause a node when a node is undergoing maintenance or troubleshooting.
- You can evict a node, which is an irreversible process for a cluster node. After you evict the node, it must be re-added to the cluster. You evict nodes when a node is damaged beyond repair or is no longer needed in the cluster. If you evict a damaged node, you can repair or rebuild it, and then add it back to the cluster by using the Add Node Wizard.

You can manage cluster nodes by using the Failover Cluster Management console.

Configuring Application Failover Settings

You can adjust the failover settings, including preferred owners and failback settings, to control how the cluster responds when the application or service fails. You can configure these settings on the property sheet for the clustered service or application found on the General tab or on the Failover tab. The following table provides examples that show how these settings work.

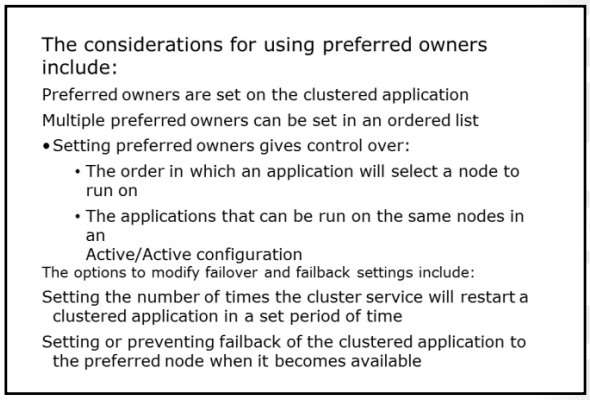

Setting	Result
Example 1: General tab, Preferred owner: Node1 Failover tab, Failback setting: Allow failback (immediately)	If the service or application fails over from Node1 to Node2, when Node1 is again available, the service or application will fail back to Node1.
Example 2: Failover tab, Maximum failures in the specified period: 2 Failover tab, Period (hours): 6	In a six-hour period, if the application or service fails no more than two times, it will be restarted or failed over every time. If the application or service fails a third time in the six-hour period, it will be left in the failed state. The default value for the maximum number of failures is n-1, where n is the number of nodes. You can change the value, but we recommend that you set a fairly low value so that if multiple node failures occur, the application or service will not be moved between nodes indefinitely.

Lesson 4
Maintaining a Failover Cluster

When a cluster infrastructure is up and running, it is very important to establish monitoring to prevent possible failures. In addition, it is important to have backup and restore procedures for cluster configuration. In Windows Server 2012, there is a new technology that enables you to update cluster nodes without downtime. In this lesson, you will learn about monitoring, backup, and restore, and about updating cluster nodes.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe how to monitor failover clusters.
- Describe how to back up and restore a failover cluster configuration.
- Describe how to monitor and troubleshoot failover clusters.
- Describe CAU.
- Describe how to configure CAU.

Monitoring Failover Clusters

Many tools are available to help you monitor failover clusters. You can use standard Windows Server tools, such as the Event Viewer and the Performance and Reliability Monitor snap-in, to review cluster event logs, and performance metrics. You can also use Tracerpt.exe to export data for analysis. In addition, you can use the MHTML-formatted cluster configuration reports and the Validate a Configuration Wizard to troubleshoot problems with the cluster configuration and hardware changes.

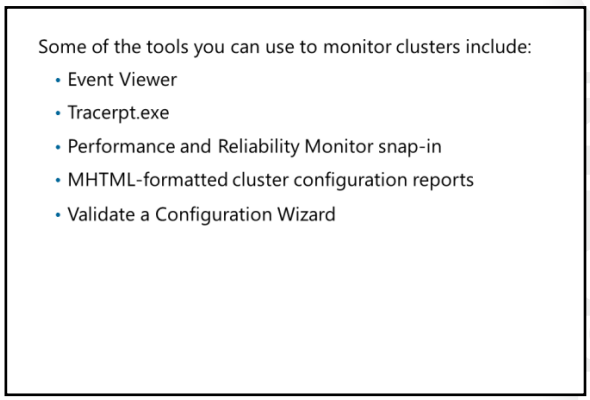

Event Viewer

When problems arise in the cluster, use the Event Viewer to view events with a Critical, Error, or Warning severity level. In addition, informational-level events are logged to the Failover Clustering Operations log, which can be found in the Event Viewer in the Applications and Services Logs\Microsoft\Windows folder. Informational-level events are usually common cluster operations, such as cluster nodes leaving and joining the cluster, or resources going offline or coming online.

In earlier versions of Windows Server, event logs were replicated to each node in the cluster. This simplified cluster troubleshooting, because you could review all event logs on a single cluster node. Windows Server 2012 does not replicate the event logs between nodes. However, the Failover Cluster Management snap-in has a Cluster Events option that enables you to view and filter events across all cluster nodes. This feature is helpful in correlating events across cluster nodes.

The Failover Cluster Management snap-in also provides a Recent Cluster Events option that queries all the Error and Warning events from all the cluster nodes in the last 24 hours.

You can access additional logs, such as the Debug and Analytic logs, in the Event Viewer. To display these logs, modify the view on the top menu by selecting the Show Analytic and Debug Logs options.

Windows Event Tracing

Windows event tracing is a kernel component that is available early after startup and late into shutdown. It is designed to enable fast tracing and delivery of events to trace files and to consumers. Because it is designed to be fast, it enables only basic in-process filtering of events based on event attributes.

The event trace log contains a comprehensive accounting of the failover cluster actions. To view the data, use Tracerpt.exe to access the information in the event trace log.

Tracerpt.exe will parse the event trace logs only on the node on which it is run. All of the individual logs are collected in a central location. To transform the XML file into a text file or an HTML file that can be opened in Internet Explorer®, you can parse the XML-based file by using the Microsoft XSL parsing command-prompt utility msxsl.exe, and an XSL style sheet.

Performance and Reliability Monitor Snap-In

The Performance and Reliability Monitor snap-in enables you to do the following:

- Trend application performance on each node. To determine how an application is performing, you can view and trend specific information on system resources that are being used on each node.
- Trend application failures and stability on each node. You can pinpoint when application failures occur, and match the application failures with other events on the node.
- Modify trace log settings. You can start, stop, and adjust trace logs, including their size and location.

Backing Up and Restoring Failover Cluster Configuration

Cluster configuration can be a time-consuming process with many details, so backup of cluster configuration is very important. You can perform backup and restore of cluster configuration with Windows Server Backup or a third-party backup tool.

When you back up the cluster configuration, be aware of the following:

- You must test your backup and recovery process before you put a cluster in to production.

When backing up failover clusters, keep in mind that:

- Windows Server Backup is an optional Windows Server feature
- Backup and restore operations involve the VSS
- Third-party tools are also available to perform backups and restores
- You must perform system-state backups

Two types of restore are:

- A non-authoritative restore completely restores a single node in the cluster
- An authoritative restore restores the entire cluster configuration to a point in time

- You must first add the Windows Server Backup feature, if you decide to use it. You can do this by using the Server Manager Add Roles and Features Wizard. Windows Server Backup uses Volume Shadow Service (VSS) to perform a backup.

Windows Server Backup is the built-in backup and recovery feature available with Windows Server 2012. To complete a successful backup, consider the following:

- For a backup to succeed in a failover cluster, the cluster must be running and must have quorum. In other words, enough nodes must be running and communicating—perhaps with a witness disk or witness file share, depending on the quorum configuration—that the cluster has achieved quorum.
- You must back up all clustered applications. For example, if you cluster an SQL Server database, you must have a backup plan for the databases and configuration outside the cluster configuration.
- If the application data must be backed up, the disks on which you store the data must be made available to the backup software. You can achieve this by running the backup software from the cluster node that owns the disk resource, or by running a backup against the clustered resource over

the network. When you have CSVs enabled in your cluster, you need to run it from any node which is a member of the CSV cluster.

- The cluster service keeps track of which cluster configuration is the most recent, and it replicates that configuration to all cluster nodes. If the cluster has a witness disk, the cluster service also replicates the configuration to the witness disk.

Restoring a Cluster

The two types of restore are:

- Non-authoritative restore. Use a non-authoritative restore when a single node in the cluster is damaged or rebuilt, and the rest of the cluster is operating correctly. Perform a non-authoritative restore by restoring the system recovery (system state) information to the damaged node. When you restart that node, it joins the cluster and receives the latest cluster configuration automatically.
- Authoritative restore. Use an authoritative restore when the cluster configuration must be rolled back to a previous time. For example, you would use an authoritative restore if an administrator accidentally removed clustered resources or modified other cluster settings. Perform the authoritative restore by stopping the cluster resource on each node, and then performing a system recovery (system state) on a single node by using the command-line Windows Server Backup interface. After the restored node restarts the cluster service, the remaining cluster nodes can also start the cluster service.

Maintaining and Troubleshooting Failover Clusters

Cluster validation functionality in Windows Server 2012 failover clustering helps prevent misconfigurations and non-working clusters. However, in some situations, you may still have to perform maintenance or cluster troubleshooting.

Some common maintenance tasks that can help prevent problems with cluster configuration include the following:

Failover cluster troubleshooting techniques include:

- Reviewing events in logs, such as: cluster, hardware and storage
- Using the Validate a Configuration Wizard
- Defining a process for troubleshooting failover clusters
- Reviewing storage configuration
- Checking for group and resource failures

- Use the Validate a Configuration Wizard to highlight configuration issues that might cause cluster problems.
- Review cluster events and trace logs to identify application or hardware issues that might cause an unstable cluster.
- Review hardware events and logs to help pinpoint specific hardware components that might cause an unstable cluster.
- Review SAN components, switches, adapters, and storage controllers to help identify any potential problems.

When you troubleshoot failover clusters, do the following:

- Identify the perceived problem by collecting and documenting the symptoms of the problem.
- Identify the scope of the problem so that you can understand what is being affected by the problem, and the impact of that effect on the application and the clients.
- Collect information so that you can accurately understand and pinpoint the possible problem. After you identify a list of possible problems, you can prioritize them by probability, or by the impact of a repair. If you cannot pinpoint the problem, you should attempt to re-create the problem.
- Create a schedule for repairing the problem. For example, if the problem only affects a small subset of users, you can delay the repair to an off-peak time so that you can schedule downtime.
- Complete and test each repair one at a time so that you can identify the fix.

To troubleshoot SAN issues, start by checking physical connections and by checking each of the hardware component logs. Additionally, run the Validate a Configuration Wizard to verify that the current cluster configuration is still supportable.

Note: When you run the Validate a Configuration Wizard, ensure that the storage tests that you select can be run on an online failover cluster. Several of the storage tests cause loss of service on the clustered disk when the tests are run.

Troubleshooting Group and Resource Failures

To troubleshoot group and resource failures:

- Use the Dependency Viewer in the Failover Cluster Management snap-in to identify dependent resources.
- Check the Event Viewer and trace logs for errors from the dependent resources.
- Determine whether the problem only happens on a specific node or nodes, by trying to re-create the problem on different nodes

What Is CAU?

Applying operating system updates to nodes in a cluster requires special attention. If you want to provide zero downtime for a clustered role, you must manually update cluster nodes one after the other, and you must manually move resources from the node being updated to another node. This procedure can be very time consuming. In Windows Server 2012, Microsoft has implemented a new feature for automatic updating of cluster nodes called Cluster-Aware Updating (CAU).

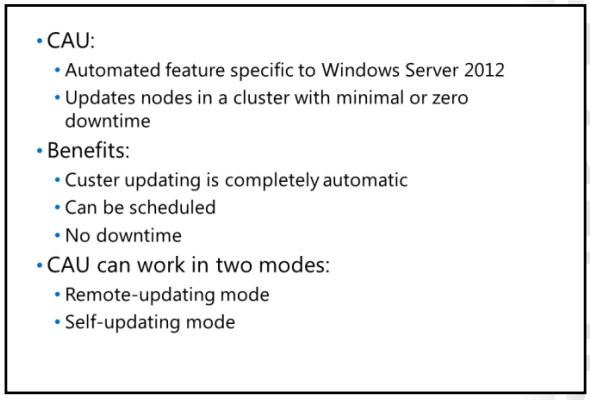

CAU is a feature that enables administrators to automatically update cluster nodes with little or no loss in availability during the update process. During an update procedure, CAU transparently takes each cluster node offline, installs the updates and any dependent updates, and then performs a restart if necessary. CAU then brings the node back online, and moves to update the next node in a cluster.

For many clustered roles, this automatic update process triggers a planned failover, and it can cause a transient service interruption for connected clients. However, for continuously available workloads in

Windows Server 2012, such as Hyper-V with the live migration feature or file server with SMB Transparent Failover, CAU can orchestrate cluster updates with no effect on the service availability.

CAU is based on orchestrating a process of cluster-node updating.

CAU can orchestrate the complete cluster updating operation in two modes:

- Remote-updating mode. In this mode, a computer that is running Windows Server 2012 or Windows 8 is called and configured as an orchestrator. To configure a computer as a CAU orchestrator, you must install failover clustering administrative tools on it. The orchestrator computer is not a member of the cluster that is updated during the procedure. From the orchestrator computer, the administrator triggers on-demand updating by using a default or custom Updating Run profile. Remote-updating mode is useful for monitoring real-time progress during the Updating Run, and for clusters that are running on Server Core installations of Windows Server 2012.
- Self-updating mode. In this mode, the CAU clustered role is configured as a workload on the failover cluster that is to be updated, and an associated update schedule is defined. In this scenario, CAU does not have a dedicated orchestrator computer. The cluster updates itself at scheduled times by using a default or custom Updating Run profile. During the Updating Run, the CAU orchestrator process starts on the node that currently owns the CAU clustered role, and the process sequentially performs updates on each cluster node. In the self-updating mode, CAU can update the failover cluster by using a fully automated, end-to-end updating process. An administrator can also trigger updates on demand in this mode, or use the remote-updating approach if desired. In the self-updating mode, an administrator can access summary information about an Updating Run in progress by connecting to the cluster and running the **Get-CauRun** Windows PowerShell cmdlet.

To use CAU, you must install the failover clustering feature in Windows Server 2012 and create a failover cluster. The components that support CAU functionality are automatically installed on each cluster node.

You must also install the CAU tools on the orchestrator node or any cluster node; these tools are included in the failover clustering tools and also are part of the Remote Server Administration Tools (RSAT). The CAU tools consist of the CAU UI and the CAU Windows PowerShell cmdlets. The failover clustering tools and CAU tools are installed by default on each cluster node when you install the failover clustering feature. You can also install these tools on a local or a remote computer that is running Windows Server 2012 or Windows 8 and that has network connectivity to the failover cluster.

Demonstration: Configuring CAU

In this demonstration, you will see how to configure CAU.

Demonstration Steps

1. Make sure that the cluster is configured and running on LON-SVR3 and LON-SVR4.
2. Add the **Failover Clustering Feature** to LON-DC1.
3. Run **Cluster-Aware Updating** on LON-DC1, and configure it to connect to **Cluster1**.
4. Preview updates that are available for nodes LON-SVR3 and LON-SVR4.
5. Review available options for the Updating Run Profile.
6. Apply available updates to **Cluster1** from LON-DC1.
7. After updates are applied, configure **Cluster self-updating options** on LON-SVR3.

Lesson 5
Implementing a Multisite Failover Cluster

In certain scenarios, you must deploy cluster nodes on different sites. Usually, you do this when you build disaster-recovery solutions. In this lesson, you will learn about deploying multisite failover clusters.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe a multisite cluster.
- Describe prerequisites for implementing multisite failover clusters.
- Describe synchronous and asynchronous replication.
- Describe how to select a quorum mode for multisite clusters.
- Describe the process for configuring multisite failover clusters.
- Describe the challenges for implementing multisite clusters.
- Describe considerations for multisite failover and failback.

What Is a Multisite Cluster?

A multisite cluster provides highly available services in more than one location. Multisite clusters can solve several specific problems. However, they also present specific challenges.

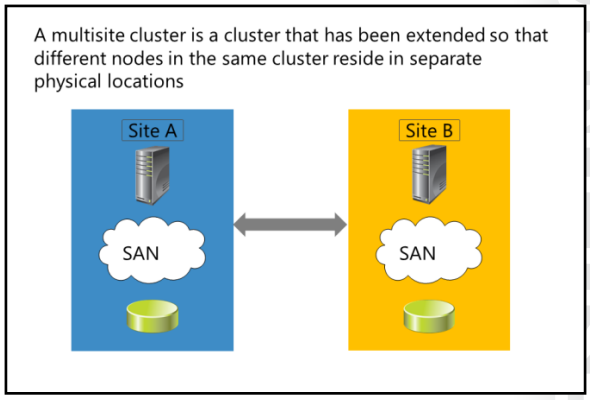

In a multisite cluster, each site usually has a separate storage system with replication between the sites. Multisite cluster storage replication enables each site to be independent, and provides fast access to the local disk. With separate storage systems, you cannot share a single disk between sites.

A multisite cluster has three main advantages in a failover site, compared to a remote server:

- When a site fails, a multisite cluster can automatically fail over the clustered service or application to another site.
- Because the cluster configuration is automatically replicated to each cluster node in a multisite cluster, there is less administrative overhead than with a cold standby server, which requires you to manually replicate changes.
- The automated processes in a multisite cluster reduce the possibility of human error, which is present in manual processes.

Because of the increased cost and complexity of a multisite failover cluster, it might not be an ideal solution for every application or business. When you are considering whether to deploy a multisite cluster, you should evaluate the importance of the applications to the business, the type of applications, and any alternative solutions. Some applications can provide multisite redundancy easily with log shipping or other processes, and can still achieve sufficient availability with only a modest increase in cost and complexity.

The complexity of a multisite cluster requires better architectural and hardware planning than is required for a single-site cluster. It also requires you to develop business processes to routinely test the cluster functionality.

Prerequisites for Implementing a Multisite Failover Cluster

Prerequisites for implementation of a multisite cluster are different from those for single-site cluster implementation. It is important to understand what you must prepare before you start to implement a multisite cluster.

To implement a multisite failover cluster, you must provide the following:

- ✓ Additional hardware to ensure enough nodes on each site
- ✓ Same operating systems and service packs on each node
- ✓ At least one low-latency and reliable network connection between sites
- ✓ Storage replication mechanism
- ✓ Infrastructure services on each site

Before you implement a multisite failover cluster, you must ensure the following:

- You must have enough nodes and votes on each site so that the cluster can be online even if one site is down. This setup requires additional hardware, and can require significant financial costs.
- All nodes must have the same operating system and service pack version.
- You must provide at least one low-latency and reliable network connection between sites. This is important for cluster heartbeats. By default, regardless of subnet configuration, heartbeat frequency, also known as subnet delay, is once every second (1,000 milliseconds). The range for heartbeat frequency is once every 250 to 2,000 milliseconds on a common subnet and 250 to 4,000 milliseconds across subnets. By default, when a node misses a series of five heartbeats, another node will initiate failover. The range for this value, also known as subnet threshold, is three through 10 heartbeats.
- You must provide a storage replication mechanism. Failover clustering does not provide a storage replication mechanism, so you must provide another solution. This also requires that you have multiple storage solutions, including one for each cluster you create.
- You must ensure that all other necessary services for cluster, such as AD DS and DNS, are also available on a second site.

You must ensure that client connections can be redirected to a new cluster node when failover happens.

Synchronous and Asynchronous Replication

Usually, it is not possible for a geographically dispersed failover cluster to use shared storage between physical locations. Geographically dispersed failover clusters must synchronize data between locations by using specialized hardware. Multisite data replication can be either synchronous or asynchronous, as follows:

- In synchronous replication, the host receives a "write complete" response from the primary storage after the data is written successfully to both storage locations
- In asynchronous replication, the host receives a "write complete" response from the primary storage after the data is written successfully on the primary storage

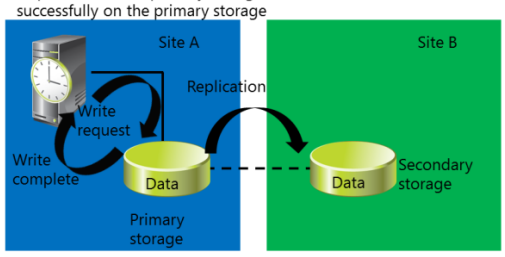

- When you use synchronous replication, the host receives a "write complete" response from the primary storage after the data is written successfully on both storage systems. If the data is not written successfully to both

storage systems, the application must attempt to write to the disk again. With synchronous replication, both storage systems are identical.

- When you use asynchronous replication, the node receives a write complete response from the storage after the data is written successfully on the primary storage. The data is written to the secondary storage on a different schedule, depending on the hardware or software vendor's implementation. Asynchronous replication can be storage based, host based, or even application based. However, not all forms of asynchronous replication are sufficient for a multisite cluster. For example, a Distributed File System (DFS) Replication provides file-level asynchronous replication. However, it does not support multisite failover clustering replication. This is because DFS Replication replicates smaller documents that are not held open continuously. Therefore, it was not designed for high-speed, open-file replication.

When to Use Synchronous or Asynchronous Replication

Use synchronous replication when data loss cannot be tolerated. Synchronous replication solutions require low-disk write latency, because the application waits for both storage solutions to acknowledge the data writes. The requirement for low-latency disk writes also limits the distance between the storage systems, because increased distance can cause higher latency. If the disk latency is high, the performance and even the stability of the application can be affected.

Asynchronous replication overcomes latency and distance limitations by acknowledging local disk writes only, and by reproducing the disk write on the remote storage system in a separate transaction. Because asynchronous replication writes to the remote storage system after it writes to the local storage system, the possibility of data loss during a failure is increased.

Selecting a Quorum Mode for Multisite Clusters

For a geographically dispersed cluster, you cannot use quorum configurations that require a shared disk, because geographically dispersed clusters do not use shared disks. Both the Node and Disk Majority, and the No Majority: Disk Only quorum modes require a shared witness disk to provide a vote for determining quorum. You should only use these two quorum modes if the hardware vendor specifically recommends and supports them.

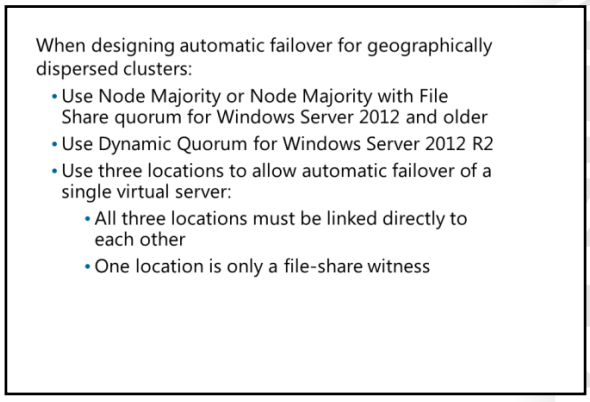

To use the Node and Disk Majority and No Majority: Disk Only modes in a multisite cluster, the shared disk requires that:

- You preserve the semantics of the SCSI commands across the sites, even if a complete communication failure occurs between sites.
- You replicate the witness disk in real-time synchronous mode across all sites.

Because multisite clusters can have WAN failures, in addition to node and local network failures, Node Majority and Node and File Share Majority are better solutions for multisite clusters. If there is a WAN failure that causes the primary and secondary sites to lose communication, a majority must still be available to continue operations.

If there is an odd number of nodes, use the Node Majority quorum. If there is an even number of nodes, which is typical in a geographically dispersed cluster, you can use the Node Majority with File Share quorum.

If you are using Node Majority and the sites lose communication, you need a mechanism to determine which nodes stay up, and which nodes drop out of the cluster membership. The second site requires another vote to obtain quorum after a failure. To obtain another vote for quorum, you must join another node to the cluster, or create a file share witness.

The Node and File Share Majority mode can help maintain quorum without adding another node to the cluster. To provide for a single-site failure and enable automatic failover, the file share witness might have to exist at a third site. In a multisite cluster, a single server can host the file share witness. However, you must create a separate file share for each cluster.

Note: If you use Windows Server 2012 R2 as host operating system in multisite cluster environment, you should use Dynamic Quorum, as discussed earlier in this module.

You must use three locations to enable automatic failover of a highly available service or application. Locate one node in the primary location that runs the highly available service or application. Locate a second node in a disaster-recovery site, and locate the third node for the file share witness in another location.

There must be direct network connectivity among all three locations. In this manner, if one site becomes unavailable, the two remaining sites can still communicate and have enough nodes for a quorum.

Note: In Windows Server 2008 R2, administrators could configure the quorum to include nodes. However, if the quorum configuration included nodes, all nodes were treated equally according to their votes. In Windows Server 2012, cluster quorum settings can be adjusted so that when the cluster determines whether it has quorum, some nodes have a vote and some do not. This adjustment can be useful when solutions are implemented across multiple sites.

When you use Windows Server 2012 R2 as the operating system for cluster nodes in a multisite cluster, you can also leverage Force Quorum Resiliency technology. This technology, as discussed earlier in this module, can be particularly useful in scenarios when sites that have cluster nodes lose connectivity.

Process for Configuring a Multisite Failover Cluster

Configuration of multisite cluster is somewhat different than configuring a single-site cluster. Multisite clusters are more complex to configure and maintain, and require more administrative effort to support. High-level steps to configure a multisite cluster are as follows:

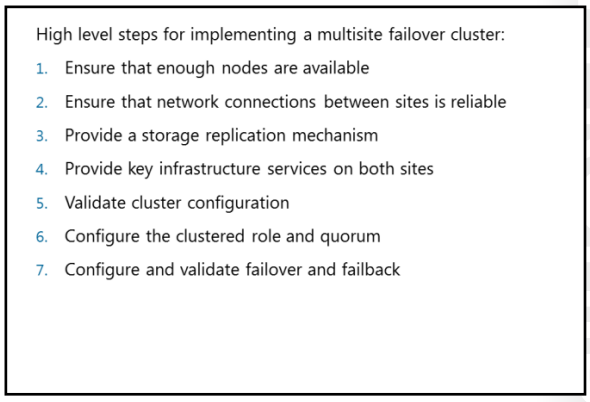

1. Ensure that you have enough cluster nodes on each site. In addition, ensure that cluster nodes have similar hardware configurations, and have the same version of operating system and service pack.
2. Ensure that networking between sites is operational, and that network latency is acceptable for configuring the cluster. You can validate this by using the Validate Configuration Wizard in Failover Cluster Manager.
3. Ensure that you have deployed a reliable storage replication mechanism between sites. Also, choose the type of replication for use.

4. Ensure that key infrastructure services such as AD DS, DNS, and DHCP are present on each site.
5. Run the Validate a Configuration Wizard on all of the cluster nodes to determine if your configuration is acceptable for creating a cluster.
6. Determine the role that you will configure in a cluster.
7. Determine the cluster quorum mode that you will use.
8. Create a clustered role.
9. Configure failover and failback settings.
10. Validate failover and failback.

You should be aware that multisite clusters require more administrative effort during failover and failback. While single-site cluster failover and failback is mostly automatic, with multisite clusters this is not the case.

Challenges with Implementing a Multisite Cluster

Implementation of multisite clusters is more complex than implementation of single-site clusters, and can also present several challenges to the administrator. The most important challenges you experience when you implement multisite clusters are related to storage and network.

Challenge	Description
Requires a separate or third-party data replication solution	• Hardware (block level) storage-based replication • Software (file system level) host-based replication • Application-based replication, such as Exchange 2007 Cluster Continuous Replication
Can be either synchronous or asynchronous replication	• Synchronous. No acknowledgement of data changes made in Site A until the data is successfully written to Site B • Asynchronous. Data changes made in Site A will eventually be written to the storage in Site B
• Inter-node communications are time sensitive; you might need to configure these thresholds to meet the higher WAN latency • DNS replication might impact client reconnect times when failover is based on hostname • Active Directory replication latency might affect application data availability • Some applications might require all of the nodes to be in the same Active Directory site	

In a multisite cluster, there is no shared storage that the cluster node uses. This means that nodes on each site must have their own storage instance. On the other hand, failover clustering does not include any built-in functionality to replicate data between sites. There are three options for replicating data: block level hardware-based replication, software-based file replication installed on the host, or application-based replication.

Multisite data replication can be either synchronous or asynchronous. Synchronous replication does not acknowledge data changes that are made in, for example, Site A, until the data is successfully written to Site B. With asynchronous replication, data changes that are made in Site A are eventually written to Site B.

When you deploy a multisite cluster and run the Validate a Configuration Wizard, the disk tests will not find any shared storage, and therefore will not run. However, you can still create a cluster. If you follow the hardware manufacturer's recommendations for Windows Server failover clustering hardware, Microsoft will support the solution.

Windows Server 2012 enables cluster nodes to exist on different IP subnets, which enables a clustered application or service to change its IP address based on the IP subnet. DNS updates the clustered application's DNS record so that clients can locate the IP address change. Because clients rely on DNS to find a service or application after a failover, you might have to adjust the DNS records' Time to Live (TTL), and the speed at which DNS data is replicated. In addition, when cluster nodes are in multiple sites, network latency might require you to modify the inter-node communication or heartbeat delay and time-out thresholds.

Multisite Failover and Failback Considerations

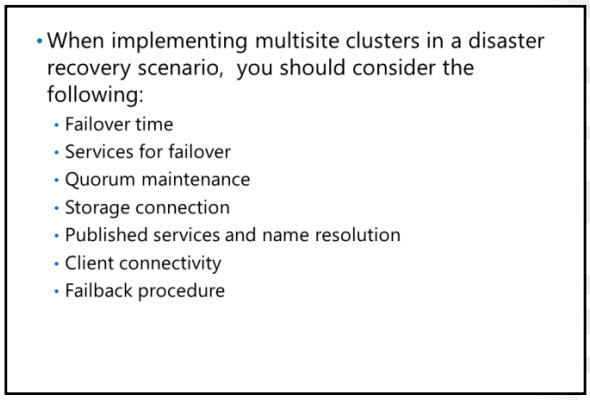

When you establish multisite failover clustering structure, it is very important that you define a procedure and the tasks that should be performed in the case of a site disaster, including tasks for failback.

In most cases, failover of critical services to another site is not automatic, but rather a manual or semi-manual procedure. When you define your failover process, you should consider the following factors:

- Failover time. You should decide how long you to wait before you pronounce a disaster and start the failover process to another site.
- Services for failover. You should clearly define the critical services, such as AD DS, DNS, and DHCP, that must be available, and that should failover to another site. It is not enough to have a cluster designed to failover to another site. Failover clustering requires that you have AD DS services up and running on a second site. You cannot make all necessary services highly available by using failover clustering, so you have to consider other technologies to achieve the result. For example, for AD DS and DNS, you can deploy additional domain controllers that also run DNS service on a second site.
- Quorum maintenance. It is very important to design quorum model in a way that each site has enough votes for maintaining the cluster functionality. If that is not possible, you can use options such as forcing a quorum or Dynamic Quorum, in Windows Server 2012 R2, to establish a quorum in case of disaster.
- Storage connection. A multisite cluster usually requires that you have storage available at each site. Because of this you should carefully design storage replication and the procedure for how to failover to secondary storage in case of a disaster.
- Published services and name resolution. If you have services published to your internal or external users, such as email, or a web page, failover to another site in some cases requires a name or IP address change. If that is the case, you should have a procedure of changing DNS records in an internal or public DNS. To reduce the downtime, we recommended that you reduce TTL on the one that contains critical the DNS records.
- Client connectivity. A failover plan must also include a design for client connectivity in case of disaster. This includes both internal and external clients. If your primary site fails, you should have a way for your clients to connect to a second site.
- Failback procedure. Once the primary site comes back online, you should plan and implement a failback process. Failback is as important as a failover, because if you perform it incorrectly, you might cause data loss and services downtime. Therefore, it is very important to clearly define steps in how to perform failback to a primary site, without data loss or corruption. The failback process is very rarely automated, and it usually happens in a very controlled environment.

Establishing a multisite cluster involves much more than just defining the cluster, cluster role, and quorum options. When you design a multisite cluster, you should consider the broader picture of failover as a part of a disaster-recovery strategy. Windows Server 2012 R2 has several technologies that can help failover and failback, but you should also consider other technologies be included in your infrastructure. In addition, each failover and failback procedure greatly depends on a service or services implemented in a cluster.

Lab: Implementing Failover Clustering

Scenario

As A. Datum's business grows, it is becoming increasingly important that many of the applications and services on the network are available at all times. A. Datum has many services and applications that must be available to internal and external users who work in different time zones around the world. Many of these applications cannot be made redundant by using Network Load Balancing (NLB). Therefore, you have to use a different technology to make these applications highly available.

As one of the senior network administrators at A. Datum, you are responsible for implementing failover clustering on the Windows Server 2012 R2 servers to provide high availability for network services and applications. You are also responsible for planning the failover cluster configuration, and for deploying applications and services on the failover cluster.

Objectives

After completing this lab, you will be able to:

- Configure a failover cluster.
- Deploy and configure a highly available file server.
- Validate the deployment of the highly available file server.
- Configure cluster-aware updating on the failover cluster.

Lab Setup

Estimated Time: 60 minutes

Virtual machines	20412C-LON-DC1 20412C-LON-SVR1 20412C-LON-SVR3 20412C-LON-SVR4
User name	**Adatum\Administrator**
Password	**Pa$$w0rd**

Lab Setup

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

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In the Hyper-V Manager, click **20412C-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 two through four for 20412C-LON-SVR1, 20412C-LON-SVR3, and 20412C-LON-SVR4.

Exercise 1: Configuring a Failover Cluster

Scenario

A. Datum has important applications and services that the company wants to make highly available. Some of these services cannot be made redundant by using NLB. Therefore, you decided to implement failover clustering. Because iSCSI storage is set up, you decide to use the iSCSI storage for failover clustering. First, you will implement the core components for failover clustering and validate the cluster, and then create the failover cluster.

The main tasks for this exercise are as follows:

1. Connect cluster nodes to the iSCSI targets
2. Install the failover clustering feature
3. Validate the servers for failover clustering
4. Create the failover cluster
5. Configuring CSV

▶ Task 1: Connect cluster nodes to the iSCSI targets

1. On LON-SVR3, start the **iSCSI Initiator**, and configure **Discover Portal** with the IP address **172.16.0.21**.
2. Connect to the discovered target in the **Targets** list.
3. Repeat steps one and two on LON-SVR4.
4. Open **Disk Management** on LON-SVR3.
5. Bring online and initialize the three new disks.
6. Make a simple volume on each disk, and format it with NTFS.
7. On LON-SVR4, open **Disk Management**, and bring online and initialize the three new disks.

▶ Task 2: Install the failover clustering feature

1. On LON-SVR3, install the Failover Clustering feature by using the Server Manager.
2. On LON-SVR4, install the Failover Clustering feature by using the Server Manager.

▶ Task 3: Validate the servers for failover clustering

1. On LON-SVR3, open the Failover Cluster Manager console.
2. Start the **Validate a Configuration Wizard**.
3. Use LON-SVR3 and LON-SVR4 as nodes for test.
4. Run all tests.
5. Review report.

▶ Task 4: Create the failover cluster

1. On LON-SVR3, in the Failover Cluster Manager, start the **Create Cluster Wizard**.
2. Use LON-SVR3 and LON-SVR4 as cluster nodes.
3. Specify **Cluster1** as the **Cluster name**.
4. Specify the **IP address** as **172.16.0.125**.

▶ Task 5: Configuring CSV

1. On LON-SVR3, in the Failover Cluster Manager console, navigate to **Storage->Disks**.
2. Locate the disk that is assigned to **Available Storage**. (If possible, use Cluster Disk 2).
3. Add this to **Cluster Shared Volumes**.

Results: After this exercise, you will have installed and configured the failover clustering feature.

Exercise 2: Deploying and Configuring a Highly Available File Server

Scenario

In A. Datum, File Services is one of the important services that must be highly available. After you have created a cluster infrastructure, you decide to configure a highly available file server and implement settings for failover and failback.

The main tasks for this exercise are as follows:

1. Add the File Server application to the failover cluster
2. Add a shared folder to a highly available file server
3. Configure failover and failback settings

▶ Task 1: Add the File Server application to the failover cluster

1. Add the File Server role service to LON-SVR4 (LON-SVR3 already has File Server Role service installed), by using the Server Manager console.
2. On LON-SVR3, open the Failover Cluster Manager console.
3. In the **Storage** node, click **Disks**, and verify that three cluster disks are online.
4. Add **File Server** as a cluster role. Select the **File Server for general use** option.
5. Specify **AdatumFS** as **Client Access Name**.
6. Specify **172.16.0.130** as the **IP address** for the cluster role.
7. Select **Cluster Disk 3** as the storage disk for AdatumFS role.
8. Complete the wizard.

▶ Task 2: Add a shared folder to a highly available file server

1. On LON-SVR4, open the Failover Cluster Manager.
2. Start the New Share Wizard, and add a new shared folder to the AdatumFS cluster role.
3. Specify the File share profile as **SMB Share – Quick**.
4. Accept the default values on the **Select the server and the path for this share** page.
5. Name the shared folder **Docs**.
6. Accept the default values on the **Configure share settings** and **Specify permissions to control access** pages.
7. Create the share at the end of wizard.

▶ Task 3: Configure failover and failback settings

1. On LON-SVR4, in the Failover Cluster Manager, open the **Properties** for the **AdatumFS** cluster role.
2. Enable failback between 4 and 5 hours.
3. Select both **LON-SVR3** and **LON-SVR4** as the preferred owners.
4. Move **LON-SVR4** to be first in the Preferred Owners list.

Results: After this exercise, you will have configured a highly available file server.

Exercise 3: Validate the Deployment of the Highly Available File Server

Scenario

In the process of implementing failover cluster, you want to perform failover and failback tests. Also, you want to change the disk witness in quorum.

The main tasks for this exercise are as follows:

1. Validate the highly available file server deployment
2. Validate the failover and quorum configuration for the file server role

▶ Task 1: Validate the highly available file server deployment

1. On LON-DC1, open File Explorer, and attempt to access the **\\AdatumFS** location. Make sure that you can access the **Docs** folder.
2. Create a test text document inside this folder.
3. On LON-SVR3, in the Failover Cluster Manager, move **AdatumFS** to the second node.
4. On LON-DC1, in File Explorer, verify that you can still access **\\AdatumFS** location.

▶ Task 2: Validate the failover and quorum configuration for the file server role

1. On LON-SVR3, determine the current owner of the **AdatumFS** role.
2. Stop the Cluster service on the node that is the current owner of the **AdatumFS** role.
3. Verify that **AdatumFS** has moved to another node, and that the **\\AdatumFS** location is still available, by trying to access it from LON-DC1.
4. Start the Cluster service on the node in which you stopped it in step two.
5. Browse to the Disks node, and take the disk marked as **Disk Witness in Quorum** offline.
6. Verify that **AdatumFS** is still available, by trying to access it from LON-DC1.
7. Bring the disk witness online.
8. Open Cluster Quorum Settings.
9. Choose to perform advanced configuration
10. Change the witness disk to Cluster Disk 3. Do not make any other changes.

Results: After this exercise, you will have tested the failover scenarios.

Exercise 4: Configuring CAU on the Failover Cluster

Scenario

In previous Windows Server versions, implementing updates to servers with critical service was causing unwanted downtime. To enable seamless and zero-downtime cluster updating, you want to implement the CAU feature and test updates for cluster nodes.

The main tasks for this exercise are as follows:

1. Configure CAU

2. Update the failover cluster and configure self-updating

3. Prepare for the next module

▶ Task 1: Configure CAU

1. On LON-DC1, install the failover clustering feature by using the Server Manager console.
2. On LON-SVR3, open the **Windows Firewall with Advanced Security** window, and verify that following two inbound rules are enabled:
   - **Inbound Rule for Remote Shutdown (RPC-EP-In)**
   - **Inbound Rule for Remote Shutdown (TCP-In)**
3. Repeat step two on LON-SVR4.
4. On LON-DC1, from Server Manager, open Cluster-Aware Updating.
5. Connect to **Cluster1**.
6. Preview the updates available for nodes in **Cluster1**. (Note: An Internet connection is required for this step.)

▶ Task 2: Update the failover cluster and configure self-updating

1. On LON-DC1, start the update process for **Cluster1**, by selecting **Apply updates to this cluster**.
2. Accept the default values in the update wizard.
3. Wait until the update process is completed. The process is finished when both nodes have a **Succeeded** value in the **Last Run status** column.
4. On LON-SVR3, open **Cluster- Aware Updating**, and then connect to **Cluster1**.
5. Select the **Configure cluster self-updating options** option.
6. Choose to add the CAU clustered role with the self-updating mode enabled in this cluster.
7. Configure self-updating to be performed **weekly**, on **Sundays** at **4:00 AM.**
8. Close the wizard.

▶ Task 3: Prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. On the **Virtual Machines** list, right-click **20412C-LON-DC1**, and then click **Revert**.

3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps two and three for 20412C-LON-SVR1, 20412C-LON-SVR3, and 20412C-LON-SVR4.

Results: After this exercise, you will have configured CAU.

Question: What information do you have to collect as you plan a failover cluster implementation and choose the quorum mode?

Question: After running the Validate a Configuration Wizard, how can you resolve the network communication single point of failure?

Question: In what situations might it be important to enable failback of a clustered application only during a specific time?

Module Review and Takeaways

Review Questions

Question: Why is using a disk-only quorum configuration generally not a good idea?

Question: What is the purpose of CAU?

Question: What is the main difference between synchronous and asynchronous replication in a multisite cluster scenario?

Question: What is an enhanced feature in multisite clusters in Windows Server 2012?

Real-world Issues and Scenarios

Your organization is considering the use of a geographically dispersed cluster that includes an alternative data center. Your organization has only a single physical location, together with an alternative data center. Can you provide an automatic failover in this configuration?

Answer: No, you cannot provide an automatic failover in this configuration. To provide an automatic failover, you must have at least three sites.

Tools

The tools for implementing failover clustering include:

Tool	Use for	Where to find it
Failover Cluster Manager console	Cluster management	Administrative Tools
Cluster-Aware Updating console	Cluster update management	Administrative Tools
Windows PowerShell	Cmdlet-based management	Administrative Tools
Server Manager	General server mangement	Taskbar
iSCSI initiator	Establishing a connection with an iSCSI target	Administrative Tools
Disk Management	Disk and volume managment	Computer Management

Best Practice:

- Try to avoid using a quorum model that depends just on the disk for Hyper-V high availability or a Scale-Out File Server.
- Perform regular backups of the cluster configuration.
- Ensure that in case one node fails, other nodes can handle the load.
- Carefully plan multisite clusters.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Cluster Validation Wizard reports an error	Review the report that Cluster Validation Wizard provides and determine the problem.
Create Cluster Wizard reports that not all nodes support the desired clustered role	Review installed roles and features on cluster nodes. A clustered role must be installed on each cluster node.
You cannot create a print server cluster	This is not supported in Windows Server 2012. You should use other technologies, such as configuring print server in the virtual machine that is highly available, to provide a highly available print server.

Module 11

Implementing Failover Clustering with Hyper-V

Contents:

Module Overview

One benefit of implementing server virtualization is the opportunity to provide high availability, both for applications or services that have built-in high availability functionality, and for applications or services that do not provide high availability in any other way. With the Windows Server® 2012 Hyper-V® technology, failover clustering, and Microsoft® System Center 2012 Virtual Machine Manager (VMM), you can configure high availability by using several different options.

In this module, you will learn about how to implement failover clustering in a Hyper-V scenario to achieve high availability for a virtual environment. You will also learn about basic features of VMM.

Objectives

After completing this module, you will be able to:

- Describe how Hyper-V integrates with failover clustering.
- Implement Hyper-V virtual machines on failover clusters.
- Implement Hyper-V virtual machine movement.
- Manage a Hyper-V virtual environment by using VMM.

Lesson 1
Overview of Integrating Hyper-V with Failover Clustering

Failover clustering is a Windows Server 2012 feature that enables you to make applications or services highly available. To make virtual machines highly available in a Hyper-V environment, you should implement failover clustering on the Hyper-V host computers.

This lesson summarizes the high availability options for Hyper-V-based virtual machines, and then focuses on how failover clustering works, and how to design and implement failover clustering for Hyper-V.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the options for making virtual machines highly available.
- Describe how failover clustering works with Hyper-V nodes.
- Describe the new features of failover clustering for Hyper-V in Windows Server 2012.
- Describe the new features of failover clustering for Hyper-V in Windows Server 2012 R2.
- Describe the best practices for implementing high availability in a virtual environment.

Options for Making Virtual Machines Highly Available

Most organizations have some applications that are business critical and must be highly available. To make an application highly available, you must deploy it in an environment that provides redundancy for all components that the application requires. For virtual machines to be highly available, you can choose between several options. You can implement virtual machines as a clustered role (host clustering); you can implement clustering inside virtual machines (guest clustering); or you can use Network Load Balancing (NLB) inside virtual machines.

High availability options	Description
Host clustering	• Virtual machines are highly available • Does not require virtual machine operating system or application to be cluster aware
Guest clustering	• Virtual machines are failover cluster nodes • Virtual machine applications must be cluster aware • Requires iSCSI or virtual fiber channel interface for shared storage connections
NLB	• Virtual machines are NLB cluster nodes • Use for web-based applications

Host Clustering

Host clustering enables you to configure a failover cluster by using the Hyper-V host servers. When you configure host clustering for Hyper-V, you configure the virtual machine as a highly available resource. Failover protection is implemented at the host-server level. This means that the guest operating system and applications that are running within the virtual machine do not have to be cluster-aware. However, the virtual machine is still highly available.

Some examples of non-cluster aware applications are a print server (in Windows Server 2012 and newer), or a proprietary network-based application such as an accounting application. If the host node that controls the virtual machine unexpectedly becomes unavailable, the secondary host node takes control and restarts or resumes the virtual machine as quickly as possible. You can also move the virtual machine from one node in the cluster to another in a controlled manner. For example, you could move the virtual machine from one node to another while patching the host management operating system.

The applications or services that are running in the virtual machine do not have to be compatible with failover clustering, and they do not have to be aware that the virtual machine is clustered. Because the

failover is at the virtual machine level, there are no dependencies on software that is installed in the virtual machine.

Guest Clustering

Guest failover clustering is configured similarly to physical-server failover clustering, except that the cluster nodes are virtual machines. In this scenario, you create two or more virtual machines, and enable failover clustering within the guest operating system. The application or service is then enabled for high availability between the virtual machines. Because failover clustering is implemented within each virtual machine node's guest operating system, you can locate the virtual machines on a single host. This can be a quick and cost-effective configuration in a test or staging environment.

For production environments, however, you can better protect the application or service if you deploy the virtual machines on separate failover clustering enabled Hyper-V host computers. With failover clustering implemented at both the host and virtual machine levels, the resource can restart regardless of whether the node that fails is a virtual machine or a host. This configuration is also known as a *Guest Cluster Across Hosts*. It is considered an optimal high-availability configuration for virtual machines running mission-critical applications in a production environment.

You should consider several factors when you implement guest clustering:

- The application or service must be failover cluster-aware. This includes any of the Windows Server 2012 services that are cluster-aware, and any applications, such as clustered Microsoft SQL Server® and Microsoft® Exchange Server.
- Hyper-V virtual machines can use Fibre Channel-based connections to shared storage (this is specific only to Microsoft Hyper-V Server 2012 and newer), or you can implement Internet Small Computer System Interface (iSCSI) connections from the virtual machines to the shared storage. In Windows Server 2012 R2, you can also use the shared virtual hard disk feature to provide shared storage for virtual machines.

You should deploy multiple network adapters on the host computers and the virtual machines. Ideally, you should dedicate a network connection to the iSCSI connection, if you use this method to connect to storage, to the private network between the hosts, and to the network connection that the client computers use.

NLB

NLB works with virtual machines in the same manner that it works with physical hosts. It distributes IP traffic to multiple instances of a TCP/IP service, such as a web server that is running on a host within the NLB cluster. NLB transparently distributes client requests among the hosts, and it enables the clients to access the cluster by using a virtual host name or a virtual IP address. From the client computer's perspective, the cluster appears to be a single server that answers these client requests. As enterprise traffic increases, you can add another server to the cluster.

Therefore, NLB is an appropriate solution for resources that do not have to accommodate exclusive read or write requests. Examples of NLB-appropriate applications include web-based front ends to database applications or Exchange Server Client Access Servers.

When you configure an NLB cluster, you must install and configure the application on all virtual machines that will participate in the NLB cluster. After you configure the application, you install the NLB feature in Windows Server 2012 within each virtual machine's guest operating system (not on the Hyper-V hosts), and then configure an NLB cluster for the application. Older versions of Windows Server also support NLB, so that the guest operating system is not limited to only Windows Server 2012; however, you should use the same operating system versions within one NLB cluster. Similar to a Guest Cluster Across Hosts, the NLB resource typically benefits from overall increased I/O performance when the virtual machine nodes are located on different Hyper-V hosts.

Note: As with older versions of Windows Server, NLB and failover clustering should not be implemented within the same operating system because the two technologies conflict with each other.

Question: Do you use any high availability solution for virtual machines in your environment?

How Does a Failover Cluster Work with Hyper-V Nodes?

When you implement failover clustering and configure virtual machines as highly available resources, the failover cluster treats the virtual machines like any other application or service. For example, if there is a host failure, failover clustering will act to restore access to the virtual machine as quickly as possible on another host in the cluster. Only one node at a time runs the virtual machine. However, you can also move the virtual machine to any other node in the same cluster as a part of a planned migration.

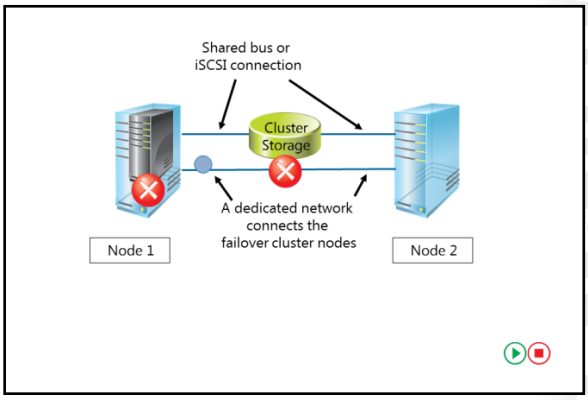

The failover process transfers the responsibility of providing access to resources in a cluster from one node to another. A *Planned failover*, also known as *switchover*, can occur when an administrator intentionally moves resources to another node for maintenance or other reasons, or when unplanned downtime of one node occurs because of a hardware failure or other reasons.

The failover process consists of the following steps:

1. The node where the virtual machine is running owns the clustered instance of the virtual machine, controls access to the shared bus or iSCSI connection to the cluster storage, and has ownership of any disks, or Logical Unit Numbers (LUNs), assigned to the virtual machine. All of the nodes in the cluster use a private network to send regular signals, known as *heartbeat signals*, to one another. The heartbeat indicates that a node is functioning and communicating on the network. The default heartbeat configuration specifies that each node send a heartbeat over TCP/UDP port 3343 each second (or 1,000 milliseconds).
2. Failover initiates when the node hosting the virtual machine does not send regular heartbeat signals over the network to the other nodes. By default, this is five consecutively missed heartbeats (or 5,000 milliseconds elapsed). Failover might occur because of a node failure or network failure. When heartbeat signals stop arriving from the failed node, one of the other nodes in the cluster begins taking over the resources that the virtual machines use.

   You define the one or more nodes that could take over by configuring the Preferred and Possible Owners properties. The Preferred Owner specifies the hierarchy of ownership if there is more than one possible failover node for a resource. By default, all nodes are members of Possible Owners. Therefore, removing a node as a Possible Owner absolutely excludes it from taking over the resource in a failure situation. Suppose that a failover cluster is implemented by using four nodes. However, only two nodes are configured as Possible Owners. In a failover event, the resource might still be taken over by the third node if neither of the Preferred Owners is online. Although the fourth node is not configured as a Preferred Owner, as long as it remains a member of Possible Owners, the failover cluster uses it to restore access to the resource, if necessary. Resources are brought online in order of dependency. For example, if the virtual machine references an iSCSI LUN, access to the appropriate

host bus adapters (HBAs), network(s), and LUNs will be stored in that order. Failover is complete when all the resources are online on the new node. For clients that are interacting with the resource, there is a short service interruption, which most users might not notice.

3. You also can configure the cluster service to fail back to the offline node after it again becomes active. When the cluster service fails back, it uses the same procedures that it performs during failover. This means that the cluster service takes all of the resources associated with that instance offline, moves the instance, and then brings all of the resources in the instance back online.

What Is New in Failover Clustering for Hyper-V in Windows Server 2012?

In Windows Server 2012, failover clustering has greatly improved with respect to Hyper-V clusters. Some of the most important improvements include:

The Failover Clustering improvements for Hyper-V in Windows Server 2012 include:

- Support for up to 8,000 virtual machines per cluster
- Multi-select virtual machines for Live Migration
- Virtual machine priority attribute
- CSV improvements
- Virtual machine application monitoring
- Storage of virtual machines on highly available SMB file share

- Failover clustering now supports up to 64 nodes and 8,000 virtual machines per cluster (and 1024 virtual machines per node), and the improved Failover Cluster Manager snap-in simplifies managing multiple virtual machines.
- Administrators now can perform multi-select actions to queue live migrations of multiple virtual machines, instead of doing it one by one, as was required in older versions of Windows Server.
- Administrators can configure virtual machine priority attributes to control the order in which virtual machines are started. Priority is also used to ensure that lower-priority virtual machines automatically release resources if they are needed by higher-priority virtual machines.
- The Cluster Shared Volume (CSV) feature, which simplifies the configuration and operation of virtual machines, is improved to allow more security and better performance. It now supports scalable file-based server-application storage, increased backup and restore, and single, consistent file namespace. In addition, you now can protect CSVs by using BitLocker® Drive Encryption and configuring them to make storage visible to only a subset of nodes.
- Virtual machine application monitoring is enhanced. You now can monitor services running on clustered virtual machines. In clusters running Windows Server 2012, administrators can configure monitoring of services on clustered virtual machines that are also running Windows Server 2012. This functionality extends the high-level monitoring of virtual machines that is implemented in Windows Server 2008 R2 failover clusters.
- It is now possible to store virtual machines on server message block (SMB) file shares in a file server cluster. This is a new way to provide high availability for virtual machines. Instead of making a cluster between Hyper-V Server nodes, you now can have Hyper-V nodes out of cluster but with virtual machine files on a highly available file share. To enable this feature, you should deploy a file-server cluster in a Scale-Out File Server mode. Scale-Out File Servers also can use CSVs for storage.

What Is New in Failover Clustering for Hyper-V in Windows Server 2012 R2?

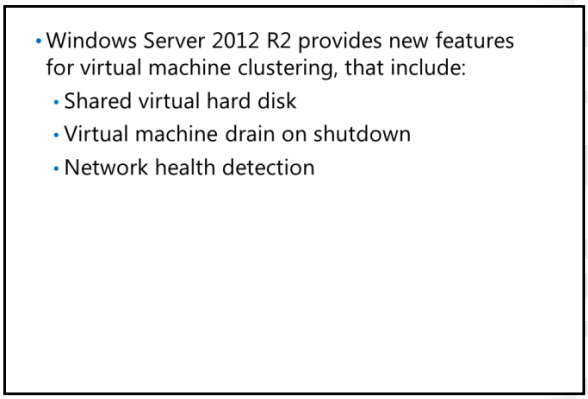

Microsoft has enhanced the existing functionalities in Windows Server 2012 R2, and has also provided additional functionalities for virtual machine clustering.

These features provide you with the ability to implement failover clustering with less administrative time, and to manage and monitor cluster resources more effectively.

The new features for virtual machine clustering in Windows Server 2012 R2 are:

- Shared virtual hard disk. When creating a guest cluster, you can now use a VHDX virtual hard disk to provide shared storage for cluster nodes. By using this, it is no longer required to have shared storage on a Fibre Channel or to have an iSCSI interface available to virtual machines.
- Virtual machine drain on shutdown. This feature provides an additional safety mechanism in scenarios when one cluster node shuts down. In Windows Server 2012 R2, if such a scenario occurs, virtual machines are automatically live migrated (instead of placed in a saved state, such as in a Quick Migration) to another cluster node.
- Network health detection. This feature helps in scenarios when virtual machines lose a connection to the physical or external network. If this happens to highly available virtual machines, failover clustering will automatically migrate affected virtual machines to another cluster node.

When you plan for high availability for virtual machines in Windows Server 2012 R2, you should be aware of these features so that you can build a stable environment with fewer downtime periods. These features are discussed in more detail in the next lesson.

Question: Do you think that these new features will be useful for your environment? If yes, which one(s)?

Best Practices for Implementing High Availability in a Virtual Environment

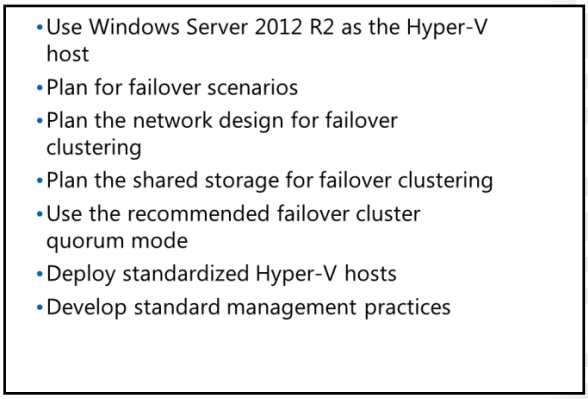

After you determine which applications will be deployed on highly available failover clusters, you plan and deploy the failover clustering environment. Apply the following recommendations when you implement the failover cluster:

- Use Windows Server 2012 R2 as the Hyper-V host. Windows Server 2012 R2 provides enhancements such as new version of Hyper-V, improved CSVs, virtual machine migrations, and other features that improve flexibility and performance when you implement host failover clustering.
- Plan for failover scenarios. When you design the hardware requirements for the Hyper-V hosts, ensure that you include the hardware capacity required when hosts fail. For example, if you deploy a six-

node cluster, you must determine the number of host failures that you want to accommodate. If you decide that the cluster must sustain the failure of two nodes, then the four remaining nodes must have the capacity to run all of the virtual machines in the cluster.

- Plan the network design for failover clustering. To optimize the failover cluster performance and failover, you should dedicate a fast network connection for internode communication. As with older versions, this network should be logically and physically separate from the network segment(s) used for clients to communicate with the cluster. You can also use this network connection to transfer virtual machine memory during a Live Migration. If you are using iSCSI for any virtual machines, ensure that you also dedicate a network connection to the iSCSI network connection.
- Plan the shared storage for failover clustering. When you implement failover clustering for Hyper-V, the shared storage must be highly available. If the shared storage fails, the virtual machines will all fail, even if the physical nodes are functional. To ensure the storage availability, plan for redundant connections to the shared storage, and redundant array of independent disks (RAID) on the storage device. If you decide to use a shared virtual hard disk (specific to Windows Server 2012 R2 Hyper-V), ensure that the shared disk is located on a highly available resource, such as a Scale-Out File Server.
- Use the recommended failover cluster quorum mode. If you deploy a cluster with an even number of nodes, and shared storage is available to the cluster, the Failover Cluster Manager automatically selects Node and Disk Majority quorum mode. If you deploy a cluster with an odd number of nodes, the Failover Cluster Manager selects the Node Majority quorum mode. You should not modify the default configuration unless you understand the implications of doing this. Consider using Dynamic Quorum if you are using Windows Server 2012 R2.
- Deploy standardized Hyper-V hosts. To simplify the deployment and management of the failover cluster and Hyper-V nodes, develop a standard server hardware and software platform for all nodes.
- Develop standard management practices. When you deploy multiple virtual machines in a failover cluster, you increase the risk that a single mistake may shut down a large part of the server deployment. For example, if an administrator accidentally configures the failover cluster incorrectly, and the cluster fails, all virtual machines in the cluster will be offline. To avoid this, develop and thoroughly test standardized instructions for all administrative tasks.

Lesson 2

Implementing Hyper-V Virtual Machines on Failover Clusters

Implementation of highly available virtual machines is somewhat different than implementing other roles in a failover cluster. Failover clustering in Windows Server 2012 provides many features for Hyper-V clustering, in addition to tools for virtual machine high availability management. In this lesson, you will learn about how to implement highly available virtual machines.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the components of a Hyper-V cluster.
- Describe the prerequisites for Hyper-V failover cluster implementation.
- Implement Hyper-V virtual machines on a failover cluster.
- Configure CSV.
- Explain how to configure a shared virtual hard disk.
- Explain how to use Scale-Out File Servers over SMB 3.0 for virtual machine storage.
- Describe the considerations for implementing Hyper-V virtual machines in a cluster.
- Explain how to maintain and monitor virtual machines in clusters.

Components of Hyper-V Clusters

Hyper-V as a role has some specific requirements for cluster components. To form a Hyper-V cluster, you must have at least two physical nodes. Whereas other clustered roles such as Dynamic Host Configuration Protocol (DHCP) or file server allow nodes to be virtual machines, Hyper-V nodes must be composed of physical hosts. You cannot run Hyper-V within a virtual machine on a Hyper-V host.

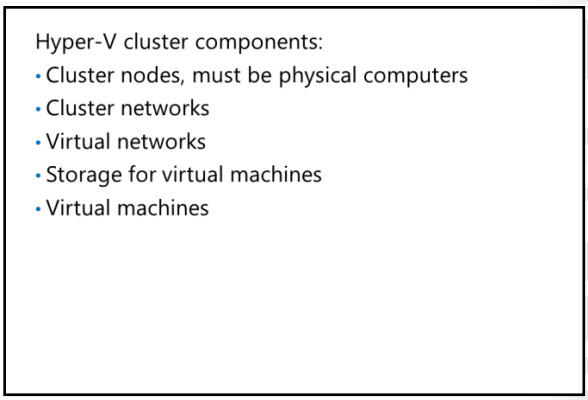

In addition to having nodes, you also must have physical and virtual networks. Failover clustering requires a cluster network interface for internal cluster communication, and also a network interface for clients. You can also implement a storage network separately, depending of type of storage being used. As a reminder, specific to the Hyper-V role, you should also consider virtual networks for clustered virtual machines. It is very important to create the same virtual networks on all physical hosts that participate in one cluster. Failing to do this causes a virtual machine to lose network connectivity when it is moved from one host to another.

Storage is an important component of virtual machine clustering. You can use any type of storage that is supported by Windows Server 2012 failover clustering and Windows Server 2012 R2 failover clustering. We recommended that you configure storage as a CSV. This is discussed in a following topic.

When you use host clustering, virtual machines are also components of a Hyper-V cluster. In the Failover Cluster Manager, you can create new highly available virtual machines, or you can make existing virtual machines highly available. In both cases, the virtual machine storage location must be on shared storage that can be accessible to both nodes. You might not want to make all virtual machines highly available. In the Failover Cluster Manager, you can select the virtual machines that are part of a cluster configuration.

Prerequisites for Implementing Hyper-V Clusters

To deploy a Hyper-V cluster, you must ensure that you meet the hardware, software, account, and network-infrastructure requirements that the following sections detail.

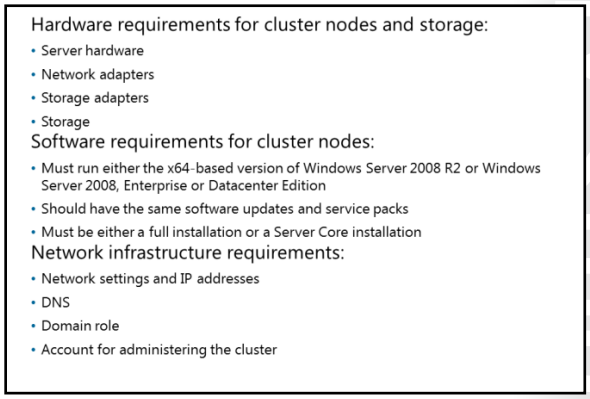

Hardware Requirements for Failover Clustering with Hyper-V

You must have the following hardware for a two-node failover cluster:

- Server hardware. Windows Server 2012 Hyper-V requires an x64-based processor, hardware-assisted virtualization, and hardware-enforced Data Execution Prevention (DEP). As a best practice, the servers should have very similar hardware. If you are using Windows Server 2008, the processors on the servers must be the same version. If you are using Windows Server 2012 R2, Windows Server 2012, or Windows Server 2008 R2, the processors must use the same architecture.

Note: Microsoft supports a failover cluster solution only if all the hardware features are marked as "Certified for Windows Server". In addition, the complete configuration (servers, network, and storage) must pass all tests in the Validate This Configuration Wizard, which is included in the Failover Cluster Manager snap-in.

- Network adapters. The network adapter hardware, like other features in the failover cluster solution, must be marked as "Certified for Windows Server." To provide network redundancy, you can connect cluster nodes to multiple networks, or you can connect the nodes to one network that uses teamed network adapters, redundant switches, redundant routers, or similar hardware to remove single points of failure. We recommend that you configure multiple physical network adapters on the host computer that you configure as a cluster node. One network adapter should be connected to the private network that the inter-host communications uses.
- Storage adapters. If you use serial attached SCSI or Fibre Channel, the mass-storage device controllers in all clustered servers should be identical and should use the same firmware version. If you use iSCSI, each clustered server should have one or more network adapters that are dedicated to the cluster storage. The network adapters that you use to connect to the iSCSI storage target should be identical, and you should use Gigabit Ethernet or a faster network adapter.
- Storage. You must use shared storage that is compatible with Windows Server 2012 R2 or Windows Server 2012. If you deploy a failover cluster that uses a *witness disk*, the storage must contain at least two separate volumes (LUNs). One volume functions as the witness disk, and additional volumes

contain the virtual machine files that are shared between the cluster nodes. Storage considerations and recommendations include the following:

- Use basic disks, not dynamic disks. Format the disks with the NTFS file system.
- Use either master boot record (MBR) or GUID partition table (GPT).
- If you use a storage area network (SAN), the miniport driver that the storage uses must work with the Microsoft Storport storage driver.
- Consider using Microsoft Multipath I/O (MPIO) software. If your SAN uses a highly available network design with redundant components, you can deploy failover clusters with multiple HBAs by using MPIO software. This provides the highest level of redundancy and availability. For Windows Server 2012 and Windows Server 2008 R2, your multipath solution must be based on MPIO.
- For environments without direct access to SAN or iSCSI storage, consider using shared virtual hard disks. Also consider using SMB 3.0 file shares as storage.

Software Requirements for Using Hyper-V and Failover Clustering

The following are the software requirements for using Hyper-V and failover clustering:

- All of the servers in a failover cluster must run the x64-based version of Windows Server 2012 Standard or Datacenter Edition. The nodes in a single failover cluster cannot run different versions, because that configuration is not supported.
- All of the servers should have the same software updates and service packs.
- All of the servers should have same drivers.

Network Infrastructure Requirements

The following network infrastructure is required for a failover cluster and an administrative account with the following domain permissions:

- Network settings and IP addresses. Use identical communication settings on all network adapters, including the speed, duplex mode, flow control, and media-type settings. Ensure that all network hardware supports the same settings.
- Private networks. If you use private networks that are not routed to your entire network infrastructure for communication between cluster nodes, ensure that each of these private networks uses a unique subnet.
- DNS. The servers in the cluster must use Domain Name System (DNS) for name resolution. You should use the DNS dynamic update protocol. It is recommended that all cluster nodes use same DNS servers for name resolution.
- Domain role. All servers in the cluster must be in the same Active Directory® domain. As a best practice, all clustered servers should have the same domain role (either member server or domain controller). The recommended role is member server.
- Account for administering the cluster. When you first create a cluster or add servers to it, you must be logged on to the domain with an account that has administrator rights and permissions on all of the cluster's servers. In addition, if the account is not a Domain Admins account, the account must have the Create Computer Objects permission in the domain.

Implementing Failover Clustering for Hyper-V Virtual Machines

To implement failover clustering for Hyper-V, you must complete the following high-level steps:

1. Install and configure Windows Server 2012
2. Configure shared storage
3. Install the Hyper-V and failover clustering features
4. Validate the cluster configuration
5. Create the cluster
6. Create a virtual machine on one of the cluster nodes
7. Make the virtual machine highly available, for existing virtual machine
8. Test virtual machine failover

1. Install and configure the required versions of Windows Server 2012. After you complete the installation, configure the network settings, join the computers to an Active Directory domain, and configure the connection to the shared storage.
2. Configure the shared storage. You must use Disk Manager to create disk partitions on the shared storage.
3. Install the Hyper-V and failover clustering features on the host servers. You can use Server Manager in Microsoft Management Console (MMC) or Windows PowerShell® to do this.
4. Validate the cluster configuration. The Validate This Cluster Wizard checks all of the prerequisite components that are required to create a cluster, and provides warnings or errors if any components do not meet the cluster requirements. Before you continue, resolve any issues that the Validate This Cluster Wizard identifies.

Note: Although it is possible to create a cluster without running cluster validation, we strongly recommended that you run the Validate This Cluster Wizard and resolve all issues before you create a cluster and put it into production.

5. Create the cluster. When the components pass the Validate This Cluster Wizard, you can create a cluster. When you configure the cluster, assign a cluster name and an IP address. A computer account for the cluster name is created in Active Directory domain, and the IP address is registered in DNS. In Windows Server 2012 R2, you can also create an Active Directory-detached cluster.

Note: You can enable Clustered Shared Storage for the cluster only after you create the cluster and add eligible storage to it. If you want to use CSV, you should configure CSV before you move to the next step.

6. Create a virtual machine on one of the cluster nodes. When you create the virtual machine, ensure that all files associated with the virtual machine, including both the Virtual Hard Disk (VHD or VHDX) and virtual machine configuration files, are stored on the shared storage. You can create and manage virtual machines in either Hyper-V Manager or Failover Cluster Manager. We recommended that you use the Failover Cluster Manager console for creating virtual machines. When you create a virtual machine using Failover Cluster Manager, the virtual machine is automatically made highly available.
7. Make the virtual machine highly available only for existing virtual machines. If you created a virtual machine before you implemented failover clustering, you should manually make it highly available. To make the virtual machine highly available, in the Failover Cluster Manager, select to make a new service or application highly available. Failover Cluster Manager then presents a list of services and applications that can be made highly available. When you select the option to make virtual machines highly available, you can select the virtual machine that you created on shared storage.

Note: When you make a virtual machine highly available, you see a list of all virtual machines hosted on all cluster nodes, including virtual machines that are not stored on the shared storage. If you make a virtual machine that is not located on shared storage highly available, you will receive a warning, but Hyper-V adds the virtual machine to the services and applications list. However, when you try to migrate the virtual machine to a different host, the migration fails.

8. Test virtual machine failover. After you make the virtual machine highly available, you can migrate the computer to another node in the cluster. If you are running Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2, you can select to perform a Quick Migration or a Live Migration.

Configuring CSV

CSVs in a Windows Server 2012 failover cluster allow multiple nodes in the cluster to simultaneously have read-write access to the same disk that is provisioned as an NTFS volume and added as storage to the cluster. When you use CSVs, clustered roles can fail over from one node to another more quickly without requiring a change in drive ownership, or dismounting and remounting a volume. CSVs also help in simplifying the management of a potentially large number of LUNs in a failover cluster.

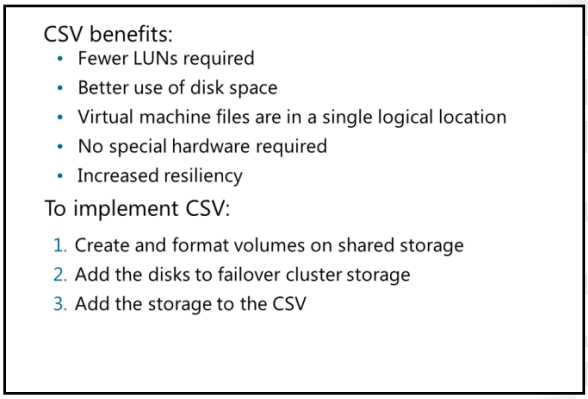

CSVs provide a general-purpose, clustered file system in Windows Server 2012, which is layered above NTFS. They are not restricted to specific clustered workloads, but currently, they are only supported for Hyper-V clusters and Scale-out File Server clusters.

Although CSVs provide additional flexibility and reduce downtime, it is not required to configure and use CSV when you implement high availability for virtual machines in Hyper-V. You can also cluster Hyper-V by using the traditional approach. However, we recommend that you use CSV because of the following advantages:

- Reduced LUNs for the disks. You can use CSV to reduce the number of LUNs that your virtual machines require. When you configure a CSV, you can store multiple virtual machines on a single LUN, and multiple host computers can access the same LUN concurrently.
- Better use of disk space. Instead of placing each .vhd or .vhdx file on a separate disk with empty space so that the .vhd/vhdx file can expand, you can oversubscribe disk space by storing multiple .vhd/.vhdx files on the same LUN.
- Single location for virtual machine files. You can track the paths of .vhd or .vhdx files and other files that virtual machines use. Instead of using drive letters or Globally Unique Identifiers (GUIDs) to identify disks, you can specify the path names. When you implement CSV, all added storage appears in the \ClusterStorage folder. The \ClusterStorage folder is created on the cluster node's system folder, and you cannot move it. This means that all Hyper-V hosts that are members of the cluster must use the same drive letter as their system drive, or the virtual machine failovers fail.
- No specific hardware requirements. There are no specific hardware requirements to implement CSV. You can implement CSV on any supported disk configuration, and on either Fibre Channel or iSCSI SANs.

- Increased resiliency. CSV increases resiliency because the cluster can respond correctly even if connectivity between one node and the SAN is interrupted, or part of a network is down. The cluster reroutes the CSV traffic through an intact part of the SAN or network.

Implementing CSV

After you create the failover cluster, you can enable CSV for the cluster, and then add storage to the CSV.

Before you can add storage to the CSV, the LUN must be available as shared storage to the cluster. When you create a failover cluster, all of the shared disks configured in Server Manager are added to the cluster, and you can add them to a CSV. You also have the option to add storage to the cluster, after the cluster is created. If you add more LUNs to the shared storage, you must first create volumes on the LUN, add the storage to the cluster, and then add the storage to the CSV.

As a best practice, you should configure CSV before you make any virtual machines highly available. However, you can convert from regular disk access to CSV after deployment. The following considerations apply:

- The LUN's drive letter or mount point is removed when you convert from regular disk access to CSV. This means that you must re-create all virtual machines that are stored on the shared storage. If you must keep the same virtual machine settings, consider exporting the virtual machines, switching to CSV, and then importing the virtual machines in Hyper-V.
- You cannot add shared storage to CSV if it is in use. If you have a running virtual machine that is using a cluster disk, you must shut down the virtual machine, and then add the disk to CSV.

Configuring a Shared Virtual Hard Disk

In previous versions of Windows Server, to implement guest clustering, you had to expose shared storage to the virtual machine. You could connect to the shared storage by using a virtual Fibre Channel interface or by using iSCSI. In some scenarios, it is a complicated task to perform, if you do not have the support of appropriate drivers for virtual Fibre Channel, or if you do not have iSCSI support on the storage. Also, in some scenarios, for example, when a virtual machine is hosted at a hosting provider, administrators do not want to expose a storage layer to the virtual machine users or tenant administrators.

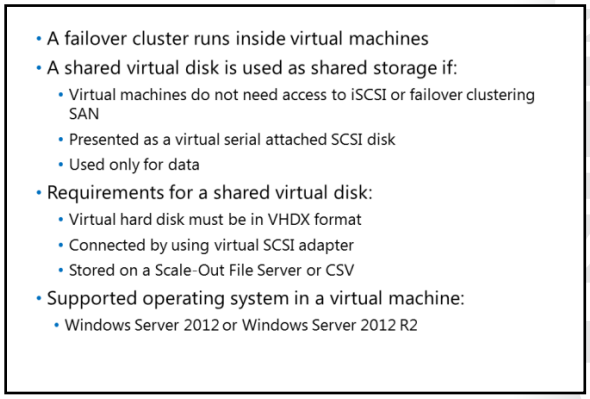

To address these issues, Windows Server 2012 R2 provides an additional layer of abstraction for virtual machine cluster storage. It is now possible to share a virtual hard disk (in .vhdx format only) between two or more virtual machines, and use that virtual hard disk as shared storage when building guest clusters. You can use the shared virtual hard disk as a witness disk or as a data disk in a cluster.

How Does a Shared Virtual Hard Disk Work?

You add shared virtual hard disks as SCSI drives in the virtual machine settings. The disks appear as virtual serial attached SCSI disks in the virtual machine. You can add a shared virtual hard disk to any virtual machine running on a Windows Server 2012 R2 Hyper-V platform. By using this technology, guest clustering configuration is simplified because you have several options for providing shared storage for guest clusters. These options include shared virtual hard disk, Fibre Channel, SMB, storage spaces, and iSCSI storage. You can use shared virtual disks to provide storage for solutions such as SQL Server databases and file server clusters.

How to Configure Shared Virtual Hard Disks?

Shared virtual disks are used only in guest cluster scenarios. To configure a guest failover cluster that uses shared virtual hard disks, you require the following:

- At least a two-node Hyper-V failover host cluster.
- All servers must be running Windows Server 2012 R2.
- All servers must belong to the same Active Directory domain.
- Configured shared storage resources must be available—for example, CSVs on block storage (such as clustered storage spaces), or a Scale-Out File Server cluster (running Windows Server 2012 R2) with SMB 3.0 for file-based storage.
- Sufficient memory, disk, and processor capacity within the failover cluster is necessary to support multiple virtual machines that are implemented as guest failover clusters.

For the guest operating systems, you can use both Windows Server 2012 R2 and Windows Server 2012. However, if you use Windows Server 2012 in virtual machines that use shared virtual hard disks, you must install Hyper-V integration services from Windows Server 2012 R2. Both Generation 1 and Generation 2 virtual machines are supported.

When you decide to implement shared virtual hard disks as storage for guest clusters, you must first decide where to store the shared virtual hard disk. You can deploy the shared virtual hard disk at the following locations:

- CSV location. In this scenario, all virtual machine files, including the shared .vhdx files, are stored on a CSV that is configured as shared storage for a Hyper-V failover cluster.
- Scale-Out File Server SMB 3.0 share. This scenario uses an SMB file-based storage as the location for the shared .vhdx files. You must deploy a Scale-Out File Server and create an SMB file share as the storage location. You also need a separate Hyper-V failover cluster.

Note: You cannot deploy a shared virtual hard disk on an ordinary file share or on a local hard disk on the host machine. You must deploy the shared virtual hard disk on a highly available location.

You can configure a shared virtual hard disk by using Hyper-V Manager graphical user interface (GUI), or by using Windows PowerShell. After you prepare your environment, and create a virtual hard disk in .vhdx format in an appropriate location, open virtual machine settings in Hyper-V Manager, and add a new SCSI disk drive. When you add a new drive, you must specify the location of your shared virtual hard disk. Before accepting changes in the virtual machine settings interface, you must mark this drive as shared in the advanced properties of the SCSI disk. Then, repeat this procedure on all virtual machines that will use this shared virtual disk drive.

To share a virtual hard disk by using Windows PowerShell, you should use the **Add-VMHardDiskDrive** cmdlet with the ***–ShareVirtualDisk*** parameter. This command must run under administrator privileges on the Hyper-V host, for each virtual machine that will use the shared .vhdx file.

For example, the following command adds a shared virtual hard disk (Data1.vhdx) stored on volume 1 of CSV to a virtual machine that is named VM1.

```
Add-VMHardDiskDrive -VMName VM1 -Path C:\ClusterStorage\Volume1\Data1.vhdx -
ShareVirtualDisk
```

Also, the following command adds a shared virtual hard disk (Witness.vhdx) that is stored on an SMB file share (\\Server1\Share1) to a virtual machine that is named VM2.

```
Add-VMHardDiskDrive -VMName VM2 -Path \\Server1\Share1\Witness.vhdx -ShareVirtualDisk
```

Comparing Shared Virtual Disk and Other Shared Storage Technologies

The following table shows a comparison among shared virtual disks, virtual Fibre Channel, and iSCSI when used for virtual machine shared storage:

Capability	Shared VHDX	Virtual Fibre Channel	ISCSI in virtual machine
Supported storage	Storage spaces, serial attached SCSI, Fibre Channel, iSCSI, SMB	Fibre Channel SAN	iSCSI SAN
Storage presented in the virtual machine as	Virtual serial attached SCSI	Virtual Fibre Channel LUN	iSCSI LUN
Data flows through the Hyper-V switch	No	No	Yes
Storage is configured at the Hyper-V host level	Yes	Yes	No
Provides low latency and a low central processing unit (CPU) use	Yes, Remote Direct Memory Access (RDMA) or Fibre Channel	Yes, Fibre Channel	No
Requires specific hardware	No	Yes	No
Requires switch to be reconfigured when virtual machine is migrated	No	Yes	No
Exposes storage architecture	No	Yes	Yes

Question: What is the main benefit of using shared hard virtual disks?

Using Scale-Out File Servers Over SMB 3.0 for Virtual Machines

In Windows Server 2012, it is possible to use one more technique to make virtual machines highly available. Instead of using host or guest clustering, virtual machine files can now be stored on a highly available SMB 3.0 file share. By using this approach, high availability is achieved not by clustering Hyper-V nodes, but by file servers that host virtual machine files on their file shares. With this new capability, Hyper-V can store all virtual machine files, including configuration, virtual hard disk (VHD) files, and checkpoints, on highly available SMB file shares.

- In Windows Server 2012 or newer, you can store virtual machine files on an SMB 3.0 file share
- File servers should run Windows Server 2012 R2
- File server cluster should run in Scale-Out File Server mode
- Hyper-V Manager can be used to create or move virtual machine files to an SMB file share

What is a Scale-Out File Server?

A Scale-Out File Server, introduced in Windows Server 2012, provides continuously available storage for file-based server applications. You create the Scale-Out File Server by implementing a file server cluster on the CSV. It is different from the file server clusters that were implemented in previous versions of Windows Server in several ways. An ordinary file-server cluster serves the clients only by using one node at a time; however, a Scale-Out File Server cluster can engage all nodes at the same time. This is achieved by taking advantage of the new Windows Server failover clustering features and the new capabilities in the new version of Windows file server protocol, SMB 3.0. Therefore, by adding nodes to the Scale-Out File Server cluster, performance of the entire cluster increases. As a result, it is now possible to store resources such as databases or virtual machine hard disks on the folder shares that are hosted on the scale-out file server cluster.

The key benefits of using a Scale-Out File Server cluster are:

- Active-active clustering. When all other failover clusters work in an active-passive mode, a Scale-Out File Server cluster works so that all nodes can accept and serve SMB client requests. In Windows Server 2012 R2, SMB 3.0 is upgraded to SMB 3.0.2. This version improves scalability and manageability for Scale-Out File Servers. SMB client connections, in Windows Server 2012 R2, are tracked per file share (instead of per server), and clients are then redirected to the cluster node with the best access to the volume used by the file share.
- Increased bandwidth. In previous versions of Windows Server, bandwidth of the file server cluster was constrained to the bandwidth of a single cluster node. Because of the active-active mode in the Scale-Out File Server cluster, you can have much higher bandwidth, which can be further increased by adding cluster nodes.
- CSV cache. Because the Scale-Out File Server clusters use CSVs, they also benefit by the use of CSV Cache. CSV Cache is a feature that you can use to allocate random access memory (RAM) as a write-through cache. The CSV Cache provides caching of read-only unbuffered I/O. This can improve performance for applications such as Hyper-V, which conducts unbuffered I/O when it accesses a VHD file. With Windows Server 2012, you can allocate up to 20 percent—and with Windows Server 2012, 80 percent—of the total physical RAM for CSV write-through cache, which will be consumed from non-paged pool memory.
- Simpler management. When using a Scale-Out File Server cluster, you can add CSV storage and shares at any time after the cluster is created.

To implement this technology, the following requirements must be met:

- One or more computers running Windows Server 2012 with the Hyper-V role must be installed.
- One or more computers running Windows Server 2012 with the File and Storage Services role must be installed.
- A common Active Directory infrastructure. The servers that are running Active Directory® Domain Services (AD DS) do not need to run Windows Server 2012.

Before you implement virtual machines on an SMB file share, you should set up a file server cluster. To do this, you should have at least two cluster nodes with file services and failover clustering installed. In the Failover Clustering console, you should create a Scale-Out File Server cluster. After you configure the cluster, you deploy the new SMB file share for applications. This share is used to store virtual machine files. When the share is created, you can use the Hyper-V Manager console to deploy new virtual machines on the SMB file share, or you can migrate existing virtual machines to the SMB file share by using the Storage Migration method.

Question: Have you considered storing virtual machines on the SMB share? Why or why not?

Considerations for Implementing Hyper-V Clusters

By implementing host failover clustering, you can make virtual machines highly available. However, implementing host failover clustering also adds significant cost and complexity to a Hyper-V deployment. You must invest in additional server hardware to provide redundancy, and you should implement or have access to a shared storage infrastructure.

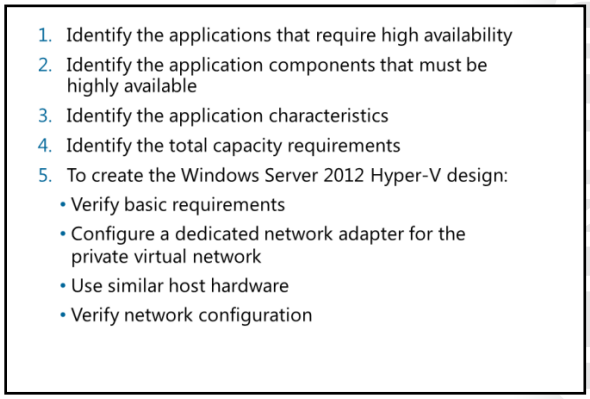

Use the following recommendations to ensure that the failover clustering strategy meets the organization's requirements:

- Identify the applications or services that require high availability. If you were to ask the people who use the organization's applications about their preferences, most of them would probably say that they want all applications to be highly available. However, unless you have the option of making all virtual machines highly available, you must develop priorities for which applications will be made highly available.
- Identify the components that must be highly available to make the applications highly available. In some cases, the application might run on a single server. If so, making that server highly available is all that you have to do. Other applications may require that several servers, and other components, such as storage or the network, be highly available.
- Identify the application characteristics. You must understand several things about the application:
    - Is virtualizing the server that is running the application an option? Some applications are not supported or recommended in a virtual environment.
    - What options are available for making the application highly available? You can make some applications highly available through options other than host clustering. If other options are available, evaluate the benefits and disadvantages of each option.

  - What are the performance requirements for each application? Collect performance information on the servers currently running the applications to gain an understanding of the hardware requirements that must be met when you virtualize the server.

- What capacity is required to make the Hyper-V virtual machines highly available? As soon as you identify all of the applications that must be highly available by using host clustering, you can start to design the actual Hyper-V deployment. By identifying the performance requirements, and the network and storage requirements for applications, you can define the hardware that you must implement in a highly available environment.

Live Migration is one of the most important aspects of Hyper-V clustering. When you implement Live Migration, consider the following:

- Verify basic requirements. The basic requirements for Live Migration are that all hosts must be part of a Windows Server 2012 or Windows Server 2012 R2 failover cluster, and that host processors must be from the same manufacturer. All hosts in the cluster must have access to shared storage.
- Configure a dedicated network adapter for the private virtual network. When you implement failover clustering, you should configure a private network for the cluster heartbeat traffic. You use this network to transfer the virtual machine memory during a failover. To optimize this configuration, configure a network adapter for this network that has a capacity of one gigabit per second (Gbps) or faster.

 **Note:** You must enable the Client for Microsoft Networks and File and Printer Sharing for Microsoft Networks components for the network adapter that you want to use for the private network.

- Use similar host hardware. All failover cluster nodes must use the same hardware for connecting to shared storage, and all cluster nodes must have processors from the same manufacturer. Although you can enable failover for virtual machines on a host with different processor versions by configuring processor compatibility settings, the failover experience and performance is more consistent if all servers have very similar hardware.
- Verify network configuration. All nodes in the failover cluster must connect through the same IP subnet, so that the virtual machine can keep the same IP address after Live Migration. In addition, the IP addresses assigned to the private network on all nodes must be on the same logical subnet. This means that multisite clusters must use a stretched virtual local area network (VLAN), which is a subnet that spans a wide area network (WAN) connection.

Demonstration: Implementing Virtual Machines on Clusters (optional)

In this demonstration, you will see how to implement highly available virtual machines with failover clustering.

Demonstration Steps

1. Ensure that LON-HOST1 is the owner of the ClusterVMs disk. If it is not, move the ClusterVMs disk to LON-HOST1.
2. On LON-HOST1, open File Explorer, browse to **E:\Program Files\Microsoft Learning\20412\Drives\20412C-LON-CORE\Virtual Hard Disks**, and then copy the **20412C-LON-CORE.vhd** virtual hard disk file to the **C:\ClusterStorage\Volume1** location. (Note: The drive letter may be different based upon the number of drives on the physical host machine)
3. In the Failover Cluster Manager, click the **Roles** node, and then start the New Virtual Machine Wizard.

4. Select LON-HOST1 as the cluster node.
5. Name the computer as **TestClusterVM**, and select it to be Generation 1 VM.
6. Store the file at **C:\ClusterStorage\Volume1**.
7. Assign 1536 MB of RAM to the TestClusterVM.
8. Connect the machine to the existing virtual hard disk drive **20412C-LON-CORE.vhd**, located at C:\ClusterStorage\Volume1.
9. Open Settings for TestClusterVM.
10. Enable the option for migration to computers with a different processor version.
11. From the Roles node, start the virtual machine.
12. On LON-HOST2, in Failover Cluster Manager, start **Live Migration** failover of **TestClusterVM** from LON-HOST1 to LON-HOST2.
13. Connect to **TestClusterVM**, and ensure that you can operate it.

Maintaining and Monitoring Virtual Machines in Clusters

In Windows Server 2012 failover clustering, you can implement the following technologies for virtual machine maintenance and monitoring:

- Service and virtual machine health monitoring
- Network health detection, for Windows Server 2012 R2 only
- Virtual machine drain on shutdown, for Windows Server 2012 R2 only

Failover clusters provide high availability for the roles that are configured in the cluster, that monitor the roles, and that take action when there is an issue with role availability. A virtual machine is one of the cluster roles, and when this role does not respond to a heartbeat, the failover cluster can restart or failover the role to a different cluster node. However, in versions prior to Windows Server 2012, the failover cluster was not able to monitor applications that were running inside a virtual machine. For example, if you used a virtual machine as a print server, the failover cluster was not able to detect if the Print Spooler service in the virtual machine had stopped. It would not take any action, although the print server did not work, because the virtual machine was still responding to a heartbeat.

Failover clustering in Windows Server 2012 has the ability to monitor and detect application health for applications and services that run inside a virtual machine. If a service in a virtual machine stops responding or an event is added to the System, Application, or Security logs, the failover cluster can take actions such as restarting the virtual machine or failing it over to a different node to restore the service. The only requirement is that both the failover cluster node and virtual machine must be running Windows Server 2012 or a newer operating system, and also have integration services installed.

You can configure virtual machine monitoring using the Failover Cluster Manager or Windows PowerShell. By default, a failover cluster is configured to monitor virtual machine health. To enable heartbeat monitoring, you must install integration services on the virtual machine. You can verify the monitoring configuration on the Settings tab of the virtual machine resource properties. To enable monitoring of any specific services that are running on the virtual machine, you must right-click the virtual machine cluster role, click More actions, and then click Configure Monitoring. Here, you can select services to monitor inside the virtual machine. The failover cluster will take action only if a service stops responding, and if, in the Services Control Manager, you have configured the service with the Take No Actions recovery setting.

Windows Server 2012 R2 can also monitor failure of virtual machine storage and loss of network connectivity, with a technology called *network health detection*. Storage failure detection can detect the failure of a virtual machine boot disk or any other virtual hard disk that the virtual machine uses. If failure happens, the failover cluster moves and restarts the virtual machine on a different node.

You can also configure a virtual network adapter to connect to a protected network. If network connectivity to such a network is lost because of reasons such as a physical switch failure or a disconnected network cable, the failover cluster will move the virtual machine to a different node to restore network connectivity.

Windows Server 2012 R2 also enhances virtual machine availability in scenarios when one Hyper-V node shuts down before being placed in maintenance mode, and before draining any clustered roles from it. In Windows Server 2012, shutting down the cluster node before draining it results in virtual machines being put into a saved state, and then moved to other nodes and resumed. This caused an interruption to the availability of the virtual machines. In Windows Server 2012 R2, if such a scenario occurs, the cluster automatically live migrates all running virtual machines before the Hyper-V node shuts down.

 **Note:** We still recommend that you drain clustered roles and place the node in maintenance mode before you perform a shutdown operation.

Configuration of this functionality, called *virtual machine drain on shutdown*, is not accessible through the Failover Cluster Manager. To configure it, you must use Windows PowerShell, and configure the **DrainOnShutdown** cluster property. It is enabled by default, and the value of this property is set to 1. If you want to check the value, you should run Windows PowerShell as Administrator, and execute the following command:

```
(Get-Cluster).DrainOnShutdown
```

Question: What are some alternative technologies that you can use for virtual machine and network monitoring?

Lesson 3
Implementing Hyper-V Virtual Machine Movement

Moving virtual machines from one location to another is a fairly common procedure in the administration of Windows Server 2012 Hyper-V environments. Most of the moving techniques in previous versions of Windows Server required downtime. Windows Server 2012 introduces new technologies that enable seamless virtual machine movement. In this lesson, you will learn about virtual machine movement and migration options.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the migration options for virtual machines.
- Explain how virtual machine and Storage Migration work.
- Describe benefits of using Offloaded Data Transfer (ODX) capable storage with Hyper-V.
- Explain how Live Migration works.
- Explain how Hyper-V Replica works.
- Describe new features of Hyper-V Replica in Windows Server 2012 R2.
- Implement Hyper-V Replica.

Virtual Machine Migration Options

There are several scenarios in which you would want to migrate a virtual machine from one location to another. For example, you might want to move a virtual machine's virtual hard disk from one physical drive to another on the same host. Another example is moving a virtual machine from one node in a cluster to another, or just moving a computer from one host server to another host server without the hosts being members of a cluster. Compared with Windows Server 2008 R2, Windows Server 2012 provides significant enhancements to, and simplified procedures for this process.

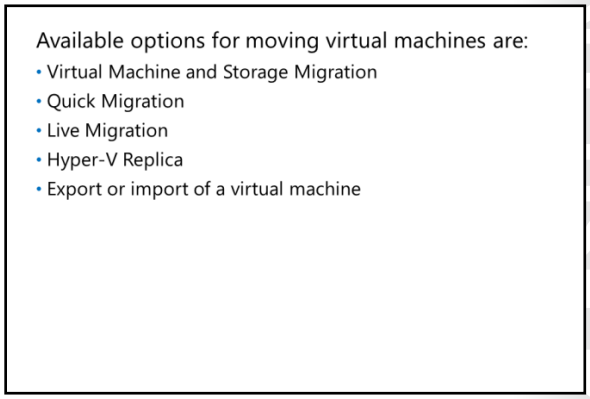

In Windows Server 2012, you can perform migration of virtual machines by using these methods:

- Virtual Machine and Storage Migration. With this method, you move a powered-on virtual machine from one location to another or from one host to another by using the Move Virtual Machine Wizard in Hyper-V Manager. Virtual Machine and Storage Migration does not require failover clustering or any other high availability technology.
- Quick Migration. This method is also available in Windows Server 2008. It requires failover clustering to be installed and configured.
- Live Migration. This improvement over Quick Migration is also available in Windows Server 2008 R2. It enables you to migrate a virtual machine from one host to another without experiencing downtime.

- Hyper-V Replica. This new feature in Windows Server 2012 enables you to replicate a virtual machine to another host or into the cloud, instead of moving the virtual machine, and to synchronize all virtual machine changes from the primary host to the host that holds the replica.
- Exporting and importing virtual machines. This is an established method of moving virtual machines without using a cluster. You export a virtual machine on one host, and then physically move exported files to another host by performing an import operation. This is a very time-consuming operation. It requires that a virtual machine be turned off during export and import. In Windows Server 2012, this migration method is improved. You can import a virtual machine to a Hyper-V host without exporting it before import. Windows Server 2012 Hyper-V is now capable of configuring all necessary settings during the import operation.

  Question: When would you export and import a virtual machine instead of migrating it?

How Does Virtual Machine and Storage Migration Work?

There are many cases in which an administrator might want to move the virtual machine files to another location. For example, if the disk where a virtual machine hard disk resides runs out of space, you must move the virtual machine to another drive or volume. Moving a virtual machine to another host is a very common procedure.

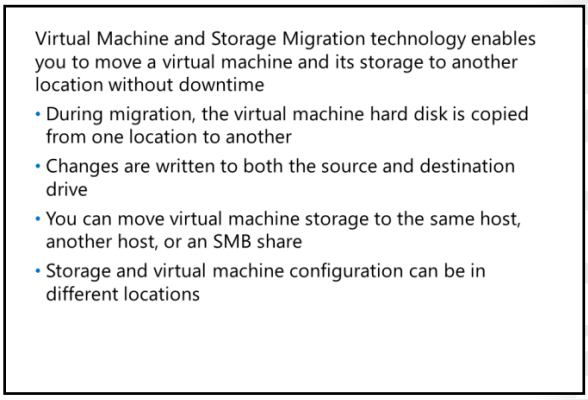

In older versions of Windows Server, such as Windows Server 2008 or Windows Server 2008 R2, moving a virtual machine resulted in downtime because the virtual machine had to be turned off. If you moved a virtual machine between two hosts, you also had to perform export and import operations for that specific machine. Export operations can be time consuming, depending on the size of the virtual machine hard disks.

In Windows Server 2012, Virtual Machine and Storage Migration enables you to move a virtual machine to another location on the same host or on another host computer without turning off the virtual machine.

To copy a virtual hard disk, an administrator starts live storage migration by using the Hyper-V console or Windows PowerShell, and completes the wizard, or specifies parameters in Windows PowerShell. A new virtual hard disk is created on the destination location, and the copy process starts. During the copy process, the virtual machine is fully functional. However, all changes that occur during copying are written to both the source and destination location. Read operations are performed only from the source location.

As soon as the disk copy process is complete, Hyper-V switches virtual machines to run on the destination virtual hard disk. In addition, if the virtual machine is moved to another host, the computer configuration is copied, and the virtual machine is associated with another host. If a failure were to occur on the destination side, there is always a fail-back option to run on the source directory. After the virtual machine is successfully migrated to and associated with a new location, the process deletes the source VHD/VHDX files and virtual machine configuration.

The time that is required to move a virtual machine depends on the source and destination location, the speed of hard disks or storage, and the size of the virtual hard disks. The moving process is accelerated if the source and destination locations are on storage, and the storage supports ODX.

When you move a virtual machine's VHDs/VHDXs and configuration files to another location, a wizard presents three available options:

- Move all the virtual machine's data to a single location: You specify one single destination location, such as disk file, configuration, checkpoint, or smart paging.
- Move the virtual machine's data to a different location: You specify individual locations for each virtual machine item.
- Move only the virtual machine's virtual hard disk: You move only the virtual hard disk file.

Using ODX Capable Storage for Virtual Machines

Windows Server 2012 provides support for ODX technology, which provides the ability to use intelligent storage arrays to directly transfer data within or between compatible storage devices as instructed by the host computer. This helps to minimize latency and maximize data bandwidth, and also greatly reduces resource usage on the host machine side, such as a CPU and a network. If you have an ODX-capable storage device or devices, Windows Server 2012 and Windows Server 2012 R2 will offload data transfer from the host and use ODX while copying or moving files by using File Explorer or command-line tools.

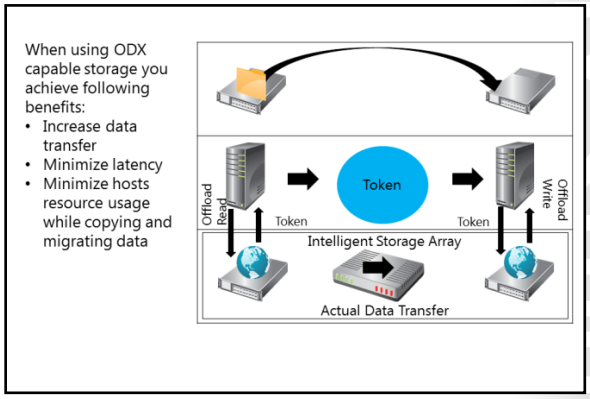

In the context of Hyper-V, ODX enables you to rapidly import and export virtual machines that are stored on storage systems supporting ODX. Moving virtual machines in a shared-nothing live migration scenario can also be significantly faster if ODX capable storage is used. In general, ODX can help with large data transfers of any type.

In a classical file transfer during copy or move operations, data is read from the storage through the source server; it is then transferred across the network to the destination server and is written back to storage but through the destination server. This procedure includes significant engagement of host-machine resources, and also engages substantial network resources.

When you use ODX capable storage, a token-based mechanism reads and writes data within one storage device or between ODX-compatible storage devices. Unlike traditional scenarios where data is copied through the source and destination host, when ODX is used, only a small token is exchanged between the source and destination computer. This token is a point-in-time representation of the data that is copied or moved. In this case, for example, when you migrate a virtual machine between storage locations, a token representing the virtual machine file is copied, and that removes the need to copy the underlying data through the servers. Copying or moving is done directly within a storage device or between storage devices, without taking almost any resource from the servers.

To use ODX-capable storage, the following procedure is required:

1. A user copies or moves a file by using File Explorer, command-line tools, or as part of a virtual machine migration.
2. Windows Server 2012 automatically translates this transfer request into an ODX, if supported by the storage device, and it receives a token that represents the data.
3. The token is copied between the source server and destination server.

4. The token is delivered to the storage array.
5. The storage array internally performs the copy or move and provides status information to the user.

Storage array that supports ODX must be connected to iSCSI, Fibre Channel, Fibre Channel over Ethernet, or a serial attached SCSI interface. On the volumes where you want to use an ODX file transfer, you cannot use Data Deduplication or BitLocker® Drive Encryption, or any other file encryption. Also, Storage Spaces and dynamic volumes are not supported.

Note: ODX file transfer is not supported by all applications that can perform copy or move operations. Currently, you can use ODX with Hyper-V management tools, File Explorer, command-line copy utilities, and cmdlets in Windows PowerShell.

How Does Live Migration Work?

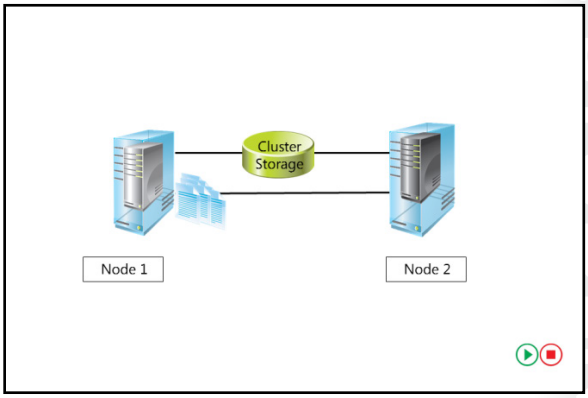

Windows Server 2012 Hyper-V allows you to move virtual machines between physical Hyper-V nodes without the need to shut down the virtual machines. This process is called Live Migration, and it can be performed in a cluster or non-cluster environment. When used within a failover cluster, Live Migration enables you to move running virtual machines from one failover cluster node to another node. If used without a cluster, Live Migration performs as a Storage Migration (described in the previous topic) and it is called *shared-nothing Live Migration*. With Live Migration, users that are connected to the virtual machine should not experience any server outage.

Note: Although you can also perform Live Migration of virtual machines using Virtual Machine and Storage Migration as described in previous topic, you should be aware that Live Migration is based on a different technology, failover clustering. Unlike the Storage Migration scenario, Live Migration is performed only if a virtual machine is highly available.

You can start a Live Migration with one of the following:

- The Failover Cluster Management console.
- The VMM Administrator console, if you use VMM to manage your physical hosts.
- Windows Management Instrumentation (WMI) or a Windows PowerShell script.

Note: Live Migration enables you to reduce the perceived outage of a virtual machine significantly during a planned failover. During a planned failover, you start the failover manually. Live Migration does not apply during an unplanned failover, such as when the node that hosts the virtual machine fails.

The Live Migration Process

The Live Migration process consists of four steps:

1. Migration setup. When the administrator starts the failover of the virtual machine, the source node creates a TCP connection with the target physical host. This connection is used to transfer the virtual machine configuration data to the target physical host. Live Migration creates a temporary virtual machine on the target physical host, and allocates memory to the destination virtual machine. The migration preparation also checks to determine whether a virtual machine can be migrated.
2. Guest-memory transfer. The guest memory is transferred iteratively to the target host while the virtual machine is still running on the source host. Hyper-V on the source physical host monitors the pages in the working set. As the system modifies memory pages, it tracks and marks them as being modified. During this phase, the migrating virtual machine continues to run. Hyper-V iterates the memory copy process several times, and every time that a smaller number of modified pages are copied to the destination physical computer. A final memory-copy process copies the remaining modified memory pages to the destination physical host. Copying stops as soon as the number of dirty pages drops below a threshold, or after 10 iterations are complete.
3. State transfer. To actually migrate the virtual machine to the target host, Hyper-V stops the source partition, transfers the state of the virtual machine, including the remaining dirty memory pages, to the target host, and then restores the virtual machine on the target host. The virtual machine has to be paused during the final state transfer.
4. Cleanup. The cleanup stage finishes the migration by tearing down the virtual machine on the source host, terminating the worker threads, and signaling the completion of the migration.

Note: In Windows Server 2012 R2, you can perform a Live Migration of virtual machines by using SMB 3.0 as a transport. This means that you can take advantage of key SMB features, such as traffic compression, SMB Direct (RDMA), and SMB Multichannel, which provide high-speed migration with low CPU utilization.

How Does Hyper-V Replica Work?

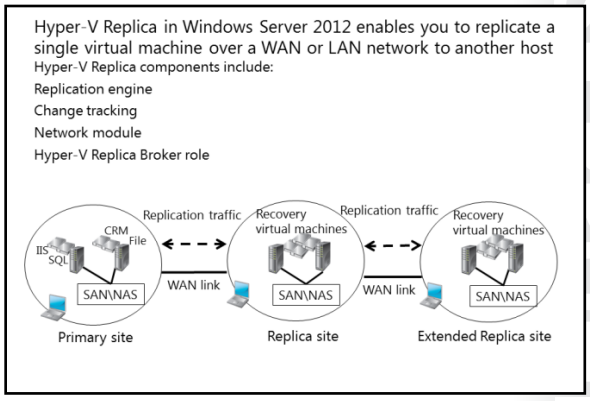

In some cases, you might want to have a spare copy of one virtual machine that you can run if the original virtual machine fails. By implementing high availability, you have one instance of a virtual machine. High availability does not prevent corruption of software that is running inside the virtual machine. One way to address the issue of corruption is to manually copy the virtual machine periodically. You also can back up the virtual machine and its storage. Although this solution achieves the desired result, it is resource intensive and time consuming. In addition, because backups are performed periodically, you never have the exact same copy as the running virtual machine.

To resolve this problem, and to enable administrators to have an up-to-date copy of a single virtual machine, Windows Server 2012 implements *Hyper-V Replica* technology. This technology enables virtual machines running at a primary site (or a location or host) to be efficiently replicated to a secondary site (location or host) across a WAN or a local area network (LAN) link. Hyper-V Replica enables you to have two instances of a single virtual machine residing on different hosts, one as the primary (live) copy and the

other as a replica (offline) copy. These copies are synchronized on a regular interval, which is configurable in the Windows Server 2012 R2 version, and you can fail over at any time.

In the event of a failure at a primary site, caused by natural disaster, a power outage, or a server failure, an administrator can use Hyper-V Manager to execute a failover of production workloads to replica servers at a secondary location within minutes, thus incurring minimal downtime. Hyper-V Replica enables an administrator to restore virtualized workloads to a specific point in time depending on the Recovery History configuration settings for the virtual machine.

Hyper-V Replica technology consists of several components:

- Replication engine. This component is the core of Hyper-V Replica. It manages the replication configuration details and handles initial replication, delta replication, failover, and test-failover operations. It also tracks virtual machine and storage mobility events and takes appropriate actions as required. For example, the replication engine pauses replication events until migration events complete, and then resumes where these events left off.
- Change tracking. This component tracks changes that are happening on the primary copy of the virtual machine. It is designed to make the scenario work regardless of where the virtual machine VHD file(s) resides.
- Network module. This module provides a secure and efficient way to transfer virtual machine replicas between the primary host and the replica host. Data compression is enabled by default. This communication is also secure, because it relies on HTTPS and certification-based authentication.
- Hyper-V Replica Broker role. This is a new role implemented in Windows Server 2012. It is configured in failover clustering, and it enables you to have Hyper-V Replica functionality even when the virtual machine being replicated is highly available and can move from one cluster node to another. The Hyper-V Replica Broker redirects all virtual machine-specific events to the appropriate node in the Replica cluster. The Broker queries the cluster database to determine which node should handle which events. This ensures that all events are redirected to the correct node in the cluster, in the event that a Quick Migration, Live Migration, or Storage Migration process was executed.

When you plan hardware configurations on the sites, you do not have to use the same server or storage hardware. It is important, however, to ensure that sufficient hardware resources are available to run the Hyper-V Replica virtual machine.

Configuring Hyper-V Replica

Before you implement Hyper-V Replica technology, ensure that these prerequisites are met:

- The server hardware supports the Hyper-V role on Windows Server 2012.
- Sufficient storage exists on both the primary and replica servers to host the files that are used by replicated virtual machines.
- Network connectivity exists between the locations that host the primary and replica servers. This can be a WAN or LAN link.
- Firewall rules are correctly configured to enable replication between the primary and replica sites (default traffic is going over TCP port 80 or 443).
- An X.509v3 certificate exists to support Mutual Authentication with certificates, if desired.

You do not have to install Hyper-V Replica separately because it is not a Windows Server role or feature. Hyper-V Replica is implemented as part of the Hyper-V role. It can be used on Hyper-V servers that are stand-alone, or on servers that are part of a failover cluster, in which case you should configure Hyper-V Replica Broker. Unlike failover clustering, a Hyper-V role is not dependent on AD DS. You can use it with Hyper-V servers that are stand-alone, or that are members of different Active Directory domains, except when servers that participate in Hyper-V replica are part of the same failover cluster.

To enable Hyper-V Replica technology, first configure Hyper-V server settings. In the Replication Configuration group of options, enable the Hyper-V server as a replica server, select the authentication and port options and configure the authorization options. You can choose to enable replication from any server that successfully authenticates, which is convenient in scenarios where all servers are part of same domain, or you can type fully qualified domain names (FQDNs) of servers that you accept as replica servers. In addition, you must configure the location for replica files. These settings should be configured on each server that will serve as a replica server.

After you configure options on the server level, enable replication on a virtual machine. During this configuration, you must specify both the replica server name and the options for connection. You can select which virtual hard disk drives you replicate, in cases when a virtual machine has more than one VHD, and you can also configure the Recovery History and the initial replication method. Specific to Windows Server 2012 R2, you can also configure replication interval; for example, for 30 seconds, five minutes (which is a default in Windows Server 2012), or 15 minutes. After you have configured these options, you can start replication. After you make the initial replica, in Windows Server 2012 R2, you can also make an extended replica to a third physical or cloud-based instance running Hyper-V. The extended replica site is built from the first replica site, not from the primary virtual machine. It is possible to configure the different replication intervals for replica and extended replica instances of a virtual machine.

New Features of Hyper-V Replica in Windows Server 2012 R2

In Windows Server 2012 R2, the Hyper-V Replica feature is improved with the following enhancements:

Hyper-V Replica in Windows Server 2012 R2 is enhanced with the following features:

- The ability to change the replication frequency:
  - The available intervals are 30 seconds, 5 minutes, and 15 minutes
- Extended replication:
  - You can extend Hyper-V Replica to include a third host

- The ability to change the replication frequency. In previous versions of Windows Server, Hyper-V Replica was to set to a five-minute replication interval, and you could not change this value. In Windows Server 2012 R2, you can set the replication interval to 30 seconds, five minutes, or 15 minutes. This means that you can configure your replication traffic based on your real environment. However, keep in mind that replica with a higher latency, for example 15 minutes, will generate more traffic when it happens.
- Extended replication. In Windows Server 2012, it is possible to have only one replica of an existing virtual machine. Windows Server 2012 R2 provides you the ability to replicate a single virtual machine to a third server. This means that you can replicate a running virtual machine to two independent servers. However, the replication does not happen from one server to two other servers. The server that is running an active copy of the virtual machine replicates to the replica server, and the replica server then replicates to the extended replica server. You create a second replica by running the Extend Replication wizard on a passive copy. In this wizard, you can set the same options that you chose when you configured the first replica.

Administrators can now benefit from these features as they help to optimize the usage of Hyper-V Replica and increase the availability of critical virtual machines.

Note: Hyper-V Replica now allows administrators to use a Windows Azure™ instance as a replica repository. This enables administrators to leverage Windows Azure, rather than having to build out a Disaster Recovery site, or manage backup tapes off-site. To use Windows Azure for

this purpose, you must have a valid subscription. Note that this services might not be available in all world regions.

Question: Do you see extended replication as a benefit for your environment?

Demonstration: Implementing Hyper-V Replica (optional)

In this demonstration, you will see how to implement Hyper-V Replica.

Demonstration Steps

1. On LON-HOST1 and LON-HOST2, configure each server to be a Hyper-V Replica server.
2. Use Kerberos (HTTP) for authentication.
3. Enable replication from any authenticated server.
4. Create and use the folder **E:\VMReplica** as a default location to store replica files.
5. Enable the firewall rule named **Hyper-V Replica HTTP Listener (TCP-In)** on both hosts.
6. On LON-HOST1, enable replication for the **20412C-LON-CORE** virtual machine:
   - Use Kerberos (HTTP).
   - Select to have only latest recovery point available.
   - Set the replication frequency to 30 seconds.
   - Start replication immediately.
7. Wait for initial replication to finish, and ensure that the **20412C-LON-CORE** virtual machine has appeared in the Hyper-V Manager console on LON-HOST2.
8. On LON-HOST2, view the replication health for **20412C-LON-CORE**.
9. On LON-HOST1, perform a planned failover to LON-HOST2. Verify that 20412C-LON-CORE is running on LON-HOST2.
10. On LON-HOST1, remove the replication for 20412C-LON-CORE.
11. On LON-HOST2, shut down 20412C-LON-CORE.

Lesson 4
Managing Hyper-V Virtual Environments by Using VMM

VMM is a part of the System Center 2012 product offerings. It is a successor to Virtual Machine Manager 2008 R2. Its main purposes are to extend management functionality for Hyper-V hosts and virtual machines, and to provide deployment and provisioning for virtual machines and services. In this lesson, you will learn the basics of VMM.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe VMM.
- Describe the new features in VMM 2012 R2.
- Describe the prerequisites for installing VMM.
- Describe the private cloud infrastructure components in VMM.
- Describe how VMM manages hosts and host groups.
- Describe how to deploy virtual machines with VMM.
- Describe VMM Services and Service Templates.
- Describe the considerations for deploying a highly available VMM server.

What Is VMM?

VMM is a management solution for a virtualized data center. VMM enables you to create and deploy virtual machines and services to private clouds by configuring and managing your virtualization host, networking, and storage resources.

VMM is a component of System Center 2012 that discovers, captures, and aggregates knowledge of the virtualization infrastructure. VMM also manages policies, processes, and best practices with automations by discovering, capturing, and aggregating knowledge of the virtualization infrastructure.

VMM provides centralized administration and management of your virtual environment and private clouds
VMM components are:
- VMM server
- Database
- Management console
- Library
- Command shell
- Self-Service Portal

VMM 2012 succeeds VMM 2008 R2, and is a key component in enabling private cloud infrastructure, which helps transition enterprise IT from an infrastructure-focused deployment model into a service-oriented, user-centric environment.

VMM architecture consists of several interrelated components. These components are:

- VMM server. The VMM server is the computer on which the VMM service runs. The VMM server processes commands and controls communications with the VMM database, the library server, and the virtual machine hosts. The VMM server is the hub of a VMM deployment through which all other VMM components interact and communicate. The VMM server also connects to an SQL Server database (VMM database) that stores all VMM configuration information.

- Database. VMM uses an SQL Server database to store the information that you view in the Virtual Machine Management Console, such as managed virtual machines, virtual machine hosts, virtual machine libraries, jobs, and other virtual machine-related data.
- Management console. The management console is a program that you use to connect to a VMM management server, to view and manage physical and virtual resources, including virtual machine hosts, virtual machines, services, and library resources.
- Library. The VMM Library is a catalog of resources (for example, virtual hard disks, templates, and profiles), that are used to deploy virtual machines and services. A library server also hosts shared folders that store file-based resources. The VMM management server is always the default library server, but you can add additional library servers later.
- Command shell. Windows PowerShell is the command-line interface you use to execute cmdlets that perform all available VMM functions. You can use these VMM-specific cmdlets to manage all the actions in a VMM environment.
- Self-Service Portal. The Self-Service Portal is a website that users who are assigned to a self-service user role can use to deploy and manage their own virtual machines.

What Is New in VMM 2012 R2?

VMM 2012 R2 includes several enhancements to the VMM product. There are enterprise-class performance enhancements that add significant advances over a previous version of VMM. The latest version of VMM includes simplified provisioning and migration abilities, support for clouds and cloud infrastructure, and enhanced ability for business units to manage the resources individually with multitenant cloud infrastructure improvements. Additionally, System Center has been extended to allow further provisioning of on-premise virtual machines and resources into the Windows Azure cloud infrastructure.

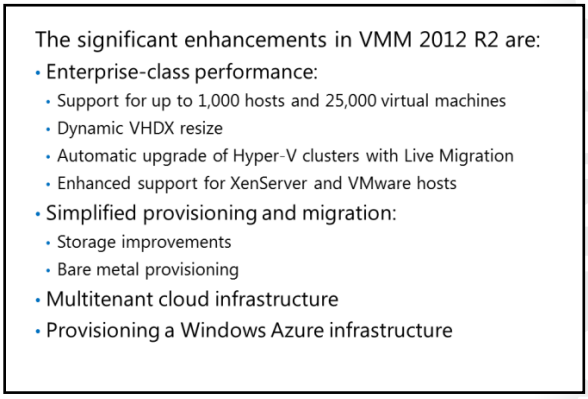

Enterprise-class Performance

System Center 2012 R2 supports an enterprise-class scale and performance for Windows Server-based environments. The VMM component of System Center 2012 R2 is key to enabling the virtualization and advanced management platform for virtualization. In this version, VMM server can support up to 1,000 hosts and 25,000 virtual machines.

Another important VMM enhancement is the Dynamic VHDX resize feature, which you can use to grow a SCSI virtual disk without any downtime. VMM support for an automated Hyper-V cluster upgrades without downtime and reduces the time, effort, cost, and downtime required to upgrade from Windows Server 2012 to Windows Server 2012 R2. You can automatically upgrade Hyper-V clusters by using Live Migration with VMM2012 R2.

There are many new and enhanced private cloud management capabilities in this new release. VMM enables Dynamic Memory changes as well as checkpoints of running virtual machines without downtime. Additionally, VMM 2012 R2 includes enhanced support for deploying VMM services to XenServer and VMware ESX hosts. This provides consistent management of Hyper-V, XenServer, and ESX-based virtual machines through the VMM console. You can manage ESX and XenServer hosts like any other VMM host.

Simplified Provisioning and Migration

Windows Server 2012 includes enhancements in File and Storage Services, including storage spaces. This means that you can use industry-standard storage that can be managed completely by server software. You can also use industry-standard servers, instead of specialty hardware technologies, for your more expensive infrastructure for storage and disaster recovery. Industry-standard server technology provides the same performance and capabilities of specialty hardware technologies at a much lower price. Using VMM 2012 R2, you can support a large, companywide storage technology infrastructure, such as a bare-metal provisioning of scaled-out Windows file server clusters, discovery of physical disks, and creation of virtualized storage pools.

Another new feature of VMM 2012 R2 is that it simplifies cross-datacenter disaster recovery of a virtual machine-based infrastructure services. It does this by providing the private cloud abstraction layer in the source and destination data centers.

Multitenant Cloud Infrastructure

Many organizations want to include the ability to quickly scale resources and increase efficiency to enhance their data center infrastructure. Additionally, organizations want the ability to provide multitenancy with increased IP flexibility, chargeback, and infrastructure standardization. VMM 2012 R2 provides greater support for multitenant environments through support for virtual networks and the ability to combine multiple instances of VMM infrastructures with the Service Provider Framework Application Programming Interface.

In addition, the latest VMM version enables you to add multitenant edge gateways to link an organization's physical and virtual data centers. This means that you can combine private cloud elements with certain elements in the public cloud, resulting in better hybrid cloud integration with enhanced mobility and more flexible workloads. VMM 2012 R2 offers multitenant enhanced chargeback with increased granular infrastructure metering, along with the ability to perform analytics on various business and operational metrics.

Provisioning a Windows Azure Infrastructure

VMM 2012 R2 is well integrated into the other System Center 2012 R2 products. In combination with these products, VMM offers a unified set of tools to provision and manage virtual machines into on-premise and Windows Azure environments. This includes easy workload portability without requiring format conversion. By using Microsoft System Center 2012 R2-App Controller, you can migrate on premise VMM virtual machines into Windows Azure virtual machines and manage those virtual machines from within the App Controller console.

Prerequisites for Installing VMM

Before you deploy VMM and its components, you should be certain that your system meets hardware and software requirements. Although software requirements do not change based on the number of hosts that VMM manages, hardware prerequisites may vary. In addition, not all VMM components have the same hardware and software requirements. However, Windows Server 2012 and Windows Server 2008 R2 are the only supported operating systems for Microsoft System Center 2012 Virtual Machine Manager.

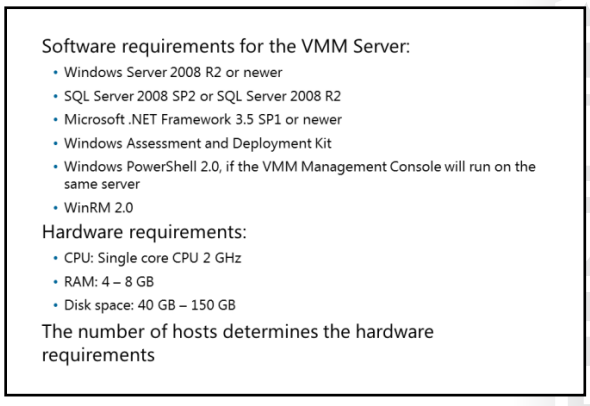

VMM Server

In addition to having Windows Server 2008 R2 or Windows Server 2012 installed, you must ensure that the following software is installed on the server that will run the VMM server:

- Microsoft .NET Framework 3.5 Service Pack 1 (SP1) or newer.
- Windows Assessment and Deployment Kit.
- Windows PowerShell® 2.0, if the VMM management console will run on the same server.
- Windows Remote Management 2.0. This is installed by default in Windows Server 2008 R2, so you should just verify that the service is running.
- SQL Server® 2008 Service Pack 2 (SP2), Standard or Enterprise, or SQL Server 2008 R2 SP1 Standard, Enterprise, or Datacenter version. This is necessary only when you install the VMM management server and SQL Server on the same computer.

Hardware requirements vary, depending on number of hosts, and include the following, at a minimum:

- CPU: Single core CPU 2 gigahertz (GHz) or Dual core CPU 2.8 GHz.
- RAM: 4 to 8 gigabytes (GB).
- Disk space: 40 to 150 GB, depending on whether an SQL Server database is installed on the same server. In addition, if the library is on the same server, then disk space will also depend on library content.

VMM Database

The VMM database stores all VMM configuration information, which you can access and modify by using the VMM management console. The VMM database requires SQL Server 2008 SP2 or newer. Because of this, the base hardware requirements for the VMM database are equal to the minimum system requirements for installing SQL Server. In addition, if you are managing more than 150 hosts, you should have at least 4 GB of RAM on the database server. Software requirements for the VMM database are the same as for SQL Server.

VMM Library

The VMM library is the server that hosts resources for building virtual machines, services, and business-unit clouds. In smaller environments, you usually install the VMM library on the VMM management server. If this is the case, the hardware and software requirements are the same as for the VMM Management Server. In larger and more complex environments, we recommend that you have VMM library on a separate server in a highly available configuration. If you want to deploy another VMM library server, the server should fulfill the following requirements:

- Supported operating system: Windows Server 2008 R2 or Windows Server 2008.
- Hardware management: Windows Remote Management 2.0.
- CPU: At least 2.8 GHz.
- RAM: At least 2 GB.
- Hard disk space: Varies based on the number and size of files that are stored.

The additional requirements for installing VMM 2012 R2 focus on the operating systems on which the various server components can run and the version of SQL Server that can store the database. Refer to the following tables to determine these requirements.

The following table lists the operating system requirements for VMM 2012 R2.

System Center 2012 R2 server-side component	Windows Server 2008 R2	Windows Server 2008 R2 SP1	Windows Server 2012 Standard, Datacenter	Windows Server 2012 R2 Standard, Datacenter
VMM management server			•	•
VMM virtual machine hosts	•	•	•	•
VMMPXE server		•	•	•
VMM update server		•	•	•
VMM library		•	•	•

The following table lists the SQL Server requirements for VMM 2012 R2.

System Center 2012 R2 component	SQL Server 2008 R2 SP1 Standard, Datacenter	SQL Server 2008 R2 SP2 Standard, Datacenter	SQL Server 2012 Enterprise, Standard (64-bit)	SQL Server 2012 SP1 Enterprise, Standard (64-bit)
VMM 2012 R2 Database Server		•	•	•

Some System Center 2012 R2 components, such as the Microsoft System Center 2012 Data Protection Manager management server, the Operations Manager management server, the Microsoft System Center 2012-Service Manager management server, and the Service Manager data warehouse management server do not work correctly if they are combined on the same server. The other components, including App Controller, Microsoft System Center 2012 Orchestrator, and VMM can run together on the same computer without issues. Keep this in mind when deploying VMM 2012 R2 and other System Center 2012 R2 components.

Private Cloud Infrastructure Components in VMM

The key architectural concept in VMM is a private cloud infrastructure. Similar to public-cloud solutions, such as in Windows Azure, a private cloud infrastructure in VMM is an abstraction layer that shields the underlying technical complexities, and enables you to manage defined resource pools of servers, networking, and storage in the enterprise infrastructure.

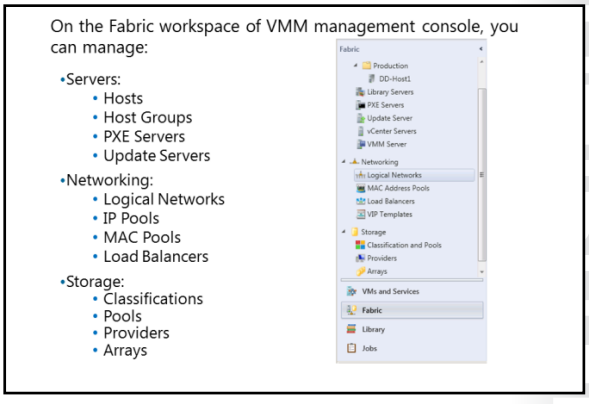

This concept is presented explicitly in the VMM management console user interface. With VMM, you can create a private cloud from Hyper-V, VMware ESX, and Citrix XenServer hosts, and benefit from cloud computing attributes, which include self-servicing, resource pooling, and elasticity.

You can configure the following resources from the VMM management console Fabric workspace:

- Servers. In the Servers node, you can configure and manage several types of servers. Host groups contain virtualization hosts, which are the destinations where you can deploy virtual machines. Library servers are the repositories of building blocks—such as images, .iso files, and templates—for creating virtual machines.
- Networking. In the VMM management console, the Networking node is where you define logical networks, assign pools of static IPs and media access control (MAC) addresses, and integrate load balancers. Logical networks are user-defined groupings of IP subnets and VLANs that organize and simplify network assignments. Logical networks provide an abstraction of the underlying physical infrastructure, and enable an administrator to provision and isolate network traffic based on selected criteria such as connectivity properties and service level agreements (SLAs).
- Storage. Using the VMM 2012 Administrator Console, an administrator can discover, classify, and provision remote storage on supported storage arrays. VMM uses the Microsoft Storage Management Service, which is enabled by default during the installation of VMM, to communicate with external arrays.

Managing Hosts and Host Groups with VMM

In addition to virtual machine management, VMM can also manage and deploy Hyper-V hosts. In VMM, you can use technologies such as Windows Deployment Services to deploy Hyper-V hosts on bare-metal machines and then manage it with VMM. When hosts are associated with VMM, you can configure several options, such as host reserves, quotas, permissions, and cloud memberships, and so VMM can also manage the Hyper-V failover clusters.

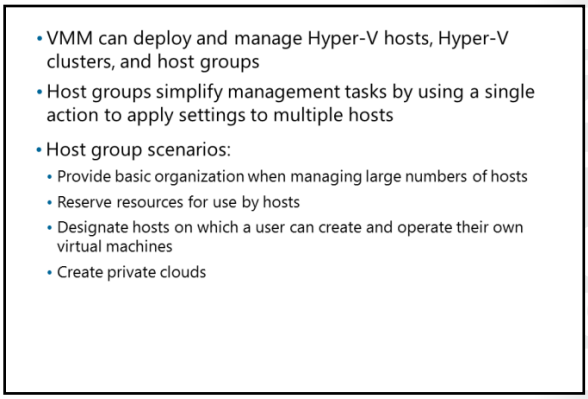

VMM provides two new features, dynamic optimization and power optimization that help optimize power and resource usage on hosts managed by VMM. Dynamic optimization balances the virtual machine load within a host cluster, while power optimization enables VMM to evacuate balanced cluster hosts, and then turn them off to save power.

We recommend that you create host groups to organize hosts in VMM. This greatly simplifies management tasks. A host group enables you to apply settings to multiple hosts with a single action. By default, there is a single host group in VMM named All Hosts. However, if necessary, you can create additional groups for your environment.

Host groups are hierarchical. When you create a new child host group, it inherits the settings from the parent host group. When a child host group moves to a new parent host group, the child host group maintains its original settings, except for Performance and Resource Optimization (PRO) settings, which are managed separately. When the settings in a parent host group change, you can apply those changes to child host groups.

Host groups are used in the following scenarios:

- Providing basic organization when you are managing multiple hosts and virtual machines. You can create custom views within the Hosts view and Virtual Machines view to provide easy monitoring and access to a host. For example, you might create a host group for each branch office in your organization.
- Reserving resources for use by hosts. Host reserves are useful when you place virtual machines on a host. Host reserves determine the CPU, memory, disk space, disk I/O capacity, and network capacity that are continuously available to the host operating system.
- Using the Host group properties action for the root host group All Hosts, to set default host reserves for all hosts that VMM manages. If you want to use more of the resources on some hosts instead of other hosts, you can set host reserves differently for each host group.
- Designating hosts on which users can create and operate their own virtual machines. When a VMM administrator adds self-service user roles, one part of role creation is to identify the hosts on which self-service users or groups in that role can create, operate, and manage their own virtual machines. We recommend that you designate a specific host group for this purpose.

Deploying Virtual Machines with VMM

One of the advantages of using a virtualized environment that is managed by VMM is the flexibility that it provides to create and deploy new virtual machines quickly.

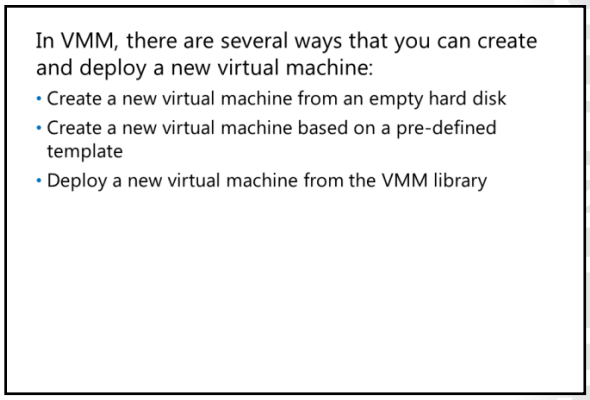

Using VMM, you can manually create a new virtual machine with new configuration settings and a new hard disk. You can then deploy the new virtual machine from one of following sources:

- An existing virtual hard disk (.vhd/.vhdx) file (blank or preconfigured)
- A virtual machine template
- A VMM library

You can create new virtual machines either by converting an existing physical computer or by cloning an existing virtual machine.

Creating a New Virtual Machine from an Existing VHD/VHDX

You can create a new virtual machine based on either a blank VHD/VHDX, or on a preconfigured VHD/VHDX that contains a guest operating system. VMM provides two blank VHD/VHDX templates that you can use to create new disks:

- Blank Disk – Small
- Blank Disk – Large

You can also use a blank VHD/VHDX when you want to use an operating system with a PXE. Alternatively, you can place an .iso image on a virtual DVD-ROM, and then install an operating system from scratch. This is an effective way to build a virtual machine's source image, which you can then use as a future template. To install the operating system on such a virtual machine, you can use an .iso image file from the library or from the local disk, then map a physical drive from the host computer, or start the guest operating system setup through a network service boot.

If you have a library of VHDs/VHDXs that you want to use in your VMM environment, you can create a virtual machine from an existing VHD/VHDX. You can also select existing VHDs/VHDX when you deploy any operating system from which VMM cannot create a template, such as an operating system that is not Windows based.

When you create a new virtual machine using an existing VHD/VHDX, you are basically creating a new virtual machine configuration that is associated with the VHD/VHDX file. VMM will create a copy of the source VHD/VHDX so that you do not have to move or modify the original.

In this scenario, the source VHD/VHDX must meet the following requirements:

- Leave the Administrator password blank on the VHD/VHDX as part of the System Preparation Tool (Sysprep) process.
- Install the Virtual Machine Additions on the virtual machine.
- Use Sysprep to prepare the operating system for duplication.

Deploying from a Template

This method creates a new virtual machine based on a template from the VMM library. The template is a library resource, which links to a virtual hard disk drive that has a generalized operating system, hardware settings, and guest operating system settings. You use the guest operating system settings to configure settings such as computer name, local administrator password, and domain membership.

The deployment process does not modify the template, which you can reuse multiple times. If you are creating virtual machines in the Self-Service Portal, you must use a template.

The following requirements apply if you want to deploy a new virtual machine from a template:

- You must install a supported operating system on the VHD/VHDX.
- You must leave the Administrator password blank on the VHD/VHDX as part of the Sysprep process. However, you do not have to leave blank the Administrator password for the guest operating system profile.
- For customized templates, you must prepare the operating system on the VHD/VHDX by removing computer identity information. For Windows operating systems, you can prepare the VHD/VHDX by using Sysprep.

Deploying from the VMM Library

If you deploy a virtual machine from the library, the virtual machine is removed from the library, and then placed on the selected host. When you use this method, you must provide the following details in the Deploy Virtual Machine Wizard:

- The host for deployment. The template that you use provides a list of potential hosts and their ratings.
- The path of the virtual machine files on the host. The virtual networks used for the virtual machine. You are presented with a list of existing virtual networks on the host.

What Are Services and Service Templates?

Services are a new concept in VMM. You must understand services fully before you deploy a private cloud infrastructure.

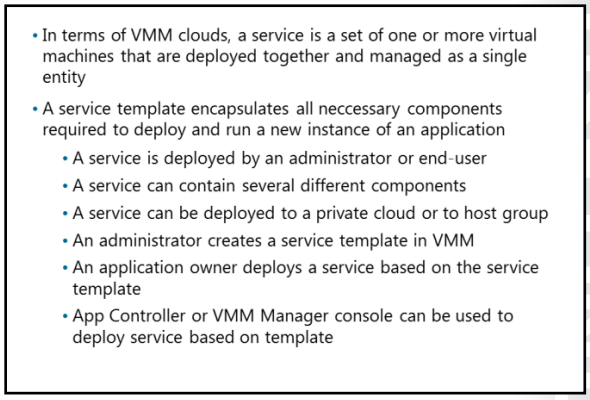

Traditional Services Scenario

When you consider services, you usually think about an application or set of applications that provide some service to end-users. For example, you can deploy various types of web-based services, but you can also implement a service such as email. In a non-cloud computing scenario, deployment of any type of service usually requires users, developers, and administrators to work together through the phases of creating a service, deploying the service, testing the service, and maintaining the service.

A service frequently includes several computers that must work together to provide a service to end-users. For example, a web-based service is usually an application that deploys on a web server, connects to a database server (which can be hosted on another computer), and performs authentication on an Active Directory domain controller. Enabling this application requires three roles, and possibly three computers: a web server, a database server, and a domain controller. Deploying a test environment for a service such as this can be time consuming and resource intensive. Ideally, developers work with IT administrators to create an environment where they can deploy and test their web application.

Concept of a Service in a Private Cloud Scenario

With the concept of a private cloud, how you deal with services can change significantly. You can prepare the environment for a service, and then let developers deploy it by using a self-service application such as App Controller.

In VMM, a service is a set of one or more virtual machines that you deploy and manage together as a single entity. You configure these machines to run together to provide a service. In VMM in Windows Server 2008, users were able to deploy new virtual machines by using the Self-Service Portal. In VMM, end-users can deploy new services. By deploying a service, users are actually deploying the entire infrastructure, including the virtual machines, network connections, and applications that are required to make the service work. However, you can use services to deploy only a single virtual machine without any specific purpose. Instead of deploying virtual machines in the historic way, you can now create a service that will deploy a virtual machine with, for example, Windows Server 2008 R2, and with several roles and features preinstalled and joined to the domain. This simplifies the process of creating and later updating new virtual machines.

Deploying a new service requires a high level of automation and predefined components, and requires management software support. This is why VMM provides service templates. A service template is a template that encapsulates everything required to deploy and run a new instance of an application. Just as a private cloud user can create new virtual machines on demand, the user can also use service templates to install and start new applications on demand.

Process for Deploying a New Service

Follow this procedure when you use service templates in VMM to deploy a new service or application:

1. The system administrator creates and configures service templates in VMM by using the Service Template Designer.

2. The end-user application owner, for example, a developer who has to deploy the application environment, opens the App Controller console, and requests a new service deployment based on available service templates that he or she can access. The developer can deploy the service to a private cloud where a user has access. As an alternative to App Controller, the user can also use the VMM Manager console.
3. A request is submitted to and evaluated by the VMM server. VMM searches for available resources in the private cloud, and then calculates the user quota and verifies that the cloud is capable of deploying the requested service.
4. Whereas the service is created automatically, the virtual machines and applications (if any) are deployed on the host selected by VMM.
5. The user application owner gains control over service virtual machines through the App Controller console, or by Remote Desktop Protocol (RDP).
6. If you require manual approval for resource creation, you can use Service Manager to create workflows for this purpose.

Information Included in the Service Template

The service template includes information about the virtual machines that are deployed as part of the service, which applications to install on the virtual machines, and the networking configuration needed for the service, including the use of a load balancer. The service template can use existing virtual machine templates. You can define the service without using any existing virtual machine templates. However, it is much easier to build a template if you have already created virtual machine templates. After you create the service template, you configure it for deployment by using the Configure Deployment option.

Considerations for Deploying a Highly Available VMM Server

VMM now supports a highly available VMM server. You can use failover clustering to achieve high availability for VMM, because VMM is now a cluster-aware application. However, you should consider several things before you deploy a VMM cluster.

The considerations for deploying a highly available VMM are:

- VMM is cluster-aware and can be highly available
- When deploying VMM in a cluster, the service account must be the domain account
- Use Distributed Key Management for encryption keys
- Make database and library servers highly available
- Do not install a Self-Service Portal on a clustered VMM server
- Use the Failover Cluster Manager to perform a planned failover

Before you begin the installation of a highly available VMM management server, ensure the following:

- You have installed and configured a failover cluster that is running Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008 R2 SP1.
- All computers on which you install the highly available VMM management server meet the minimum hardware requirements, and all prerequisite software is installed on all computers.
- You have created a domain account to be used by the VMM service. You must use a domain user account for a highly available VMM management server.
- You are prepared to use distributed key management to store encryption keys in AD DS. You must use distributed key management for a highly available VMM management server.

- You have a computer with a supported SQL Server version installed and running. Unlike VMM 2008 R2, VMM will not automatically install a SQL Server Express edition. Also, we highly recommend that you use the clustered instance of SQL Server.

Highly Available Databases and Library Servers

To achieve full redundancy, we recommend that you use a highly available server running SQL Server. You should install a highly available server running SQL Server on a separate failover cluster from the failover cluster on which you are installing the highly available VMM management server. Similarly, we also recommend that you use a highly available file server for hosting your library shares.

Self-Service Portal and Clustered VMM Server

As a best practice, do not install the VMM Self-Service Portal on the same computer as the highly available VMM management server. If your VMM Self-Service Portal currently resides on the same computer as the VMM server, we recommend that you uninstall the VMM Self-Service Portal for VMM 2008 R2 SP1 before you upgrade to VMM. We also recommend that you install the VMM Self-Service Portal on a highly available web server to achieve redundancy and load balancing.

Failover Cluster Manager

You cannot perform a planned failover, for example, to install a security update or do maintenance on a cluster node, by using the VMM console. Instead, use the Failover Cluster Manager console to perform a planned failover.

During a planned failover, ensure that there are no tasks actively running on the VMM management server. Any tasks that are executing during a failover will be stopped and will not restart automatically. Any connections to a highly available VMM management server from the VMM console or the VMM Self-Service Portal will also be lost during a failover. However, the VMM console can reconnect automatically to the highly available VMM management server after a failover, if it was opened before you performed failover to another VMM server.

Lab: Implementing Failover Clustering with Hyper-V

Scenario

The initial deployment of virtual machines on Hyper-V has been successful for A. Datum. As a next step in the deployment, A. Datum is considering ways to ensure that the services and applications deployed on the virtual machines are highly available. As part of the implementation of high availability for most network services and applications, A. Datum also is considering options for making the virtual machines that run on Hyper-V highly available.

As one of the senior network administrators at A. Datum, you are responsible for integrating Hyper-V with failover clustering to ensure that the virtual machines deployed on Hyper-V are highly available. You are responsible for planning the virtual machine and storage configuration, and for implementing the virtual machines as highly available services on the failover cluster. You also are considering other techniques, such as Hyper-V Replica, for ensuring high availability for virtual machines.

Objectives

After completing this lab, you will be able to:

- Configure a Hyper-V Replica.
- Configure a failover cluster for Hyper-V.
- Configure a highly available virtual machine.

Lab Setup

Estimated Time: 75 minutes

Virtual machines	20412C-LON-DC1-B 20412C-LON-SVR1-B
User name	**Adatum\administrator**
Password	**Pa$$w0rd**

This lab should be performed with a partner. To perform this lab, you must start the host computers to Windows Server 2012 R2. Ensure that you and your partner have started different hosts (one should start the LON-HOST1, and the other should start the LON-HOST2). Also, ensure that LON-DC1-B is imported on LON-HOST1, and LON-SVR1-B is imported on LON-HOST2, and that these virtual machines are started.

Exercise 1: Configuring Hyper-V Replicas

Scenario

Before you start with cluster deployment, you decided to evaluate new technology in Hyper-V, for replicating virtual machines between hosts. You want to be able to manually mount a copy of a virtual machine on another host if the active copy (or host) fails.

The main tasks for this exercise are as follows:

1. Import LON-CORE virtual machine on LON-HOST1

2. Configure a replica on both host machines

3. Configure replication for LON-CORE virtual machine

4. Validate a planned failover to the replica site

▶ Task 1: Import LON-CORE virtual machine on LON-HOST1

1. On LON-HOST1, open the Hyper-V Manager, and import the 20412C-LON-CORE virtual machine.
2. Use path **E:\Program Files\Microsoft Learning\20412\Drives\20412C-LON-CORE**
3. Accept default values.

Note: The drive letter may be different based upon the number of drives on the physical host machine.

▶ Task 2: Configure a replica on both host machines

1. On LON-HOST1 and LON-HOST2, configure each server to be a Hyper-V Replica server.
   - Use Kerberos (HTTP) for authentication.
   - Enable replication from any authenticated server.
   - Create and use the folder **E:\VMReplica** as a default location to store replica files.
2. Enable the firewall rule named **Hyper-V Replica HTTP Listener (TCP-In)** on both hosts.

▶ Task 3: Configure replication for LON-CORE virtual machine

1. On LON-HOST1, enable replication for the 20412C-LON-CORE virtual machine:
   - Use Kerberos (HTTP).
   - Set the replication frequency to 30 seconds.
   - Select to have only latest recovery point available.
   - Start replication immediately.
2. Wait for initial replication to finish and ensure that the 20412C-LON-CORE virtual machine has appeared in Hyper-V Manager console on LON-HOST2.

▶ Task 4: Validate a planned failover to the replica site

1. On LON-HOST2, view replication health for 20412C-LON-CORE.
2. On LON-HOST1, perform the planned failover to LON-HOST2. Verify that 20412C-LON-CORE is running on LON-HOST2.
3. On LON-HOST1, remove replication for 20412C-LON-CORE.
4. On LON-HOST2, shut down 20412C-LON-CORE.

Results: After completing this exercise, you will have Hyper-V Replica configured.

Exercise 2: Configuring a Failover Cluster for Hyper-V

Scenario

A. Datum has several virtual machines that are hosting important services that must be highly available. Because these services are not cluster-aware, A. Datum has decided to implement a failover cluster on the Hyper-V host level. You plan to use iSCSI drives as storage for these virtual machines.

The main tasks for this exercise are as follows:

1. Connect to iSCSI target from both host machines
2. Configure failover clustering on both host machines
3. Configure disks for failover cluster

▶ Task 1: Connect to iSCSI target from both host machines

1. On LON-HOST1, start the **iSCSI initiator**.
2. Use the 172.16.0.21 address to discover and connect to the iSCSI target.
3. On LON-HOST2, start the **iSCSI initiator**.
4. Use the 172.16.0.21 address to discover and connect to the iSCSI target.
5. On LON-HOST2, open Disk Management, and initialize and bring online all iSCSI drives:
   - Format the first drive, and name it **ClusterDisk**.
   - Format the second drive, and name it **ClusterVMs**.
   - Format the third drive, and name it **Quorum**.
6. On LON-HOST1, open Disk Management, and bring all three iSCSI drives online.

▶ Task 2: Configure failover clustering on both host machines

1. On LON-HOST1 and LON-HOST2, install the Failover Clustering feature.
2. On LON-HOST1, create a failover cluster:
   - Add LON-HOST1 and LON-HOST2.
   - Name it **VMCluster**.
   - Assign the address **172.16.0.126**.
   - Clear the option to **Add all eligible storage to the cluster**.

▶ Task 3: Configure disks for failover cluster

1. In Failover Cluster Manager on LON-HOST1, add all three iSCSI disks to the cluster.
2. Verify that all three iSCSI disks appear available for cluster storage.
3. Add the disk with the volume name of **ClusterVMs** to **Cluster Shared Volumes**.
4. From the **VMCluster.adatum.com** node, select **More Actions**, and then configure the Cluster Quorum Settings to use default configuration.

Results: After completing this exercise, students will have the failover clustering infrastructure configured for Hyper-V.

Exercise 3: Configuring a Highly Available Virtual Machine

Scenario

After you have configured the Hyper-V failover cluster, you want to add virtual machines as highly available resources. Also, you want to evaluate Live Migration and test Storage Migration.

The main tasks for this exercise are as follows:

1. Copy virtual machine storage to iSCSI target
2. Configure the virtual machine as highly available
3. Perform a Live Migration for the virtual machine
4. Perform a Storage Migration for the virtual machine
5. Prepare for the next module

▶ Task 1: Copy virtual machine storage to iSCSI target

1. Ensure that LON-HOST1 is the owner of the ClusterVMs disk. If it is not, move the ClusterVMs disk to LON-HOST1.
2. On LON-HOST1, open File Explorer, browse to **E:\Program Files\Microsoft Learning\20412\Drives\20412C-LON-CORE\Virtual Hard Disks**, and then copy the **20412C-LON-CORE.vhd** virtual hard disk file to the **C:\ClusterStorage\Volume1** location.

▶ Task 2: Configure the virtual machine as highly available

1. In the Failover Cluster Manager, click the **Roles** node, and then start the New Virtual Machine Wizard.
   - Select LON-HOST1 as the cluster node.
   - Name the computer as **TestClusterVM**.
   - Store the file at C:\ClusterStorage\Volume1.
   - Assign 1536 MB of RAM to the TestClusterVM.
   - Connect the machine to the existing virtual hard disk drive 20412C-LON-CORE.vhd, located at C:\ClusterStorage\Volume1.
2. Open settings for TestClusterVM.
3. Enable the option for migration to computers with a different processor version.
4. From the Roles node, start the virtual machine.

▶ Task 3: Perform a Live Migration for the virtual machine

1. On LON-HOST2, in the Failover Cluster Manager, start **Live Migration** failover of **TestClusterVM** from LON-HOST1 to LON-HOST2.
2. Connect to **TestClusterVM**, and ensure that you can operate it.

▶ Task 4: Perform a Storage Migration for the virtual machine

1. On LON-HOST1, browse to E:\Program Files\Microsoft Learning\20412\Drives\ and create new folder on this location. Name the folder LON-GUEST1.
2. Browse to E:\Program Files\Microsoft Learning\20412\Drives\20412C-LON-CORE\Virtual Hard Disks, and then copy the 20412C-LON-CORE.vhd virtual hard disk file to the E:\Program Files\Microsoft Learning\20412\Drives\LON-GUEST1 location.

3. Create new virtual machine on LON-HOST1, with following settings:
   - Name: LON-GUEST1
   - Location: E:\Program Files\Microsoft Learning\20412\Drives\LON-GUEST1
   - Memory: 1024 MB
   - Network: External Network
   - Hard disk: E:\Program Files\Microsoft Learning\20412\Drives\LON-GUEST1\20412C-LON-CORE.vhd
4. On LON-HOST1, open the Hyper-V Manager, and start LON-GUEST1.
5. Perform a Move operation on LON-GUEST1. Move the virtual machine from its current location to C:\GUEST1.
6. Check whether the machine is operational during the move process.
7. When complete, shut down all running virtual machines.

▶ Task 5: Prepare for the next module

1. Restart LON-HOST1.
2. When you are prompted with the boot menu, select **Windows Server 2012**, and then press Enter.
3. Sign in to the host machine as directed by your instructor.
4. Repeat steps one through three on LON-HOST2.

Results: After completing this exercise, the students will have configured the virtual machine as highly available.

Question: How can you extend Hyper-V Replica in Windows Server 2012 R2?

Question: What is the difference between Live Migration and Storage Migration?

Module Review and Takeaways

Review Question

Question: Do you have to implement CSV in order to provide high availability for virtual machines in VMM in Windows Server 2012?

Tools

The tools for implementing failover clustering with Hyper-V include:

Tools	Where to Find	Use
Failover Cluster Manager	Administrative Tools	Failover clustering management
Hyper-V Manager	Administrative Tools	Virtual machine management
VMM Console	Start menu	Hyper-V hosts and virtual machine management

Best Practice:

- Develop standard configurations before you implement highly available virtual machines. The host computers should be configured as close to identically as possible. To ensure that you have a consistent Hyper-V platform, you should configure standard network names, and use consistent naming standards for CSVs.
- Use new features in Hyper-V Replica to extend your replication to more than one server.
- Consider using Scale-Out File Server clusters as storage for highly available virtual machines.
- Implement VMM. VMM provides a management layer on top of Hyper-V and Failover Cluster Management that can block you from making mistakes when you manage highly available virtual machines. For example, it blocks you from creating virtual machines on storage that is inaccessible from all nodes in the cluster.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Virtual machine failover fails after you implement CSV and migrate the shared storage to CSV.	The CSV home folder is located on the host-server system drive. You cannot move it. If the host computers use different system drives, the failovers will fail because the hosts cannot access the same storage location. All failover cluster nodes should use the same hard-drive configuration.
A virtual machine fails over to another node in the host cluster, but loses all network connectivity.	All the nodes in a host cluster must have the same networks configured. If they do not, then the virtual machines cannot connect to a network when they failover to another node.
Four hours after restarting a Hyper-V host that is a member of a host cluster, there are still no virtual machines running on the host.	By default, virtual machines do not fail back to a host computer after they have migrated to another host. You can enable failback on the virtual machine properties in Failover Cluster Management, or you can implement PRO in VMM.

Module 12

Implementing Business Continuity and Disaster Recovery

Contents:

Module Overview

Organizations are always vulnerable to losing some of their data—for reasons such as unintentional deletion, file system corruption, hardware failures, malicious users, and natural disasters. Because of this, organizations must have well-defined and tested recovery strategies that will help them to bring their servers and data back to a healthy and operational state, and in the fastest time possible.

In this module, you will learn how to identify security risks for your organization. You will also learn about data protection and recovery, including how to back up specific data locally and to the cloud, how to back up servers, and how you can recover data.

Objectives

After completing this module, you will be able to:

- Describe data protection concepts.
- Implement the Windows Server Backup feature in Windows Server® 2012 R2.
- Implement server, service, and data recovery.

Lesson 1
Data Protection Overview

Data Protection is a catchall term that describes the many technologies and methods that allow you to bring data, services, and servers back to an operational state after an unplanned event has occurred. This unplanned event might involve data corruption or an application failure, or it might even include the loss of a site through flooding or fire. An effective data protection strategy addresses the organization's needs without providing an unnecessary level of coverage. While absolute protection may seem desirable, it is unlikely to be economically feasible. In developing a data protection strategy, you need to balance the cost to the organization of a particular type of data loss, with the cost to the organization of protection from that data loss.

Lesson Objectives

After completing this lesson, you will be able to:

- Identify data protection requirements.
- Describe service-level agreements (SLAs).
- Describe enterprise data protection strategies.
- Describe disaster mitigation strategies.
- Describe best practices for implementing a data protection strategy.

Identifying Recovery Requirements

Before you develop a data protection strategy, organizations must identify their recovery requirements to ensure that they will provide appropriate protection for critical resources.

There are three broad areas of data protection:

- Data recovery. Allows recovery of lost or corrupted data. This is the type of data protection and recovery that is performed most regularly in IT organizations.
- Failure recovery. Allows recovery of virtual machines, applications, or services in the event of hardware or software failure or corruption.
- Disaster protection. Allows recovery of servers, virtual machines, applications, and data in the event that these are lost to causes external to software and hardware failure such as a fire in the server room or a flood damaging a site.

Identify your data protection options by:

1. Defining organization critical resources
2. Identifying risks associated with those critical resources
3. Identifying the time needed to complete the recovery
4. Developing a protection strategy

The following is a high-level list of steps that you can use to identify data protection requirements:

1. Define organization-critical resources. These resources include data, services, and the servers which the data and services run.
2. Identify risks associated with those critical resources. For example, data can be accidentally or intentionally deleted, and a hard drive or storage controller where data is stored might fail. Additionally, services that use critical data might fail for many reasons (such as network problems), and servers might fail because of hardware failures. Major power outages could also cause entire sites to shut down.

3. Determine the amount of time needed to perform the recovery. Based on their business requirements, organizations should decide how much time is acceptable for recovering critical resources. Scenarios may vary from minutes to hours, or even one day.
4. Develop a recovery strategy. Based on the previous steps, organizations will define a service-level agreement (SLA) that will include information such as service levels and service hours. Organizations should develop a data protection strategy that will help them minimize the risks, and at the same time recover their critical resources within the minimum time acceptable for their business requirements.

Note: Organizations will have differing data protection requirements based on their business requirements and goals. Data protection requirements should not be static, but they should be evaluated and updated on a regular basis—for example, once every few months. It is also important that administrators test the data protection strategies on a regular basis. This testing should be performed in an isolated, non-production environment by using a copy of the production data.

What Are Service Level Agreements?

A *service level agreement* (SLA) is a document that describes the responsibilities of the IT department or IT service provider, with respect to a specific set of objectives. In terms of data protection SLAs, these agreements usually specify precisely which parts of the IT infrastructure and data will be protected, and how quickly they will return to service after a failure.

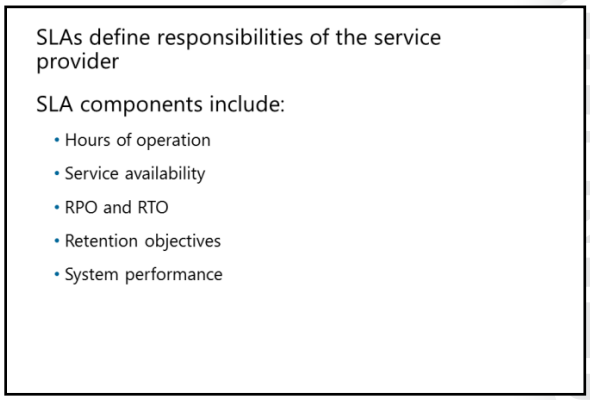

In some organizations, SLAs are formalized, and the performance of the IT department is measured against the objectives that are spelled out in the SLA. These metrics form part of the IT department's performance evaluation, and have a direct influence on items such as budgets and salaries. For managed services or cloud providers, SLAs are critical for billing purposes. In other organizations, SLAs are guidelines and are less formalized. The key to developing a SLA is to ensure that its standards are realistic and achievable, rather than unrealistic and unachievable.

Some of the elements of an SLA include:

- Hours of operation. *Hours of operation* defines how much time the data and services are available to users, and how much planned downtime there will be due to system maintenance.
- Service availability. *Service availability* is defined as a percentage of time per year that data and services will be available to users. For example, a service availability of 99.9 percent per year means that data and services will have unplanned downtime not more than 0.1 percent per year, or 8.75 hours per year on a 24 hours a day, seven days a week basis. In some cases, this will only apply to business hours, although in a globalized environment, business hours usually mean 24 hours each day.
- Recovery point objective (RPO). A *RPO* sets a limit on how much data can be lost due to failure, measured as a unit of time. For example, if an organization sets an RPO of six hours, it would be necessary to perform a backup every six hours, or to create a replication copy on different locations at six-hour intervals. In the event of a failure, it would be necessary to go back to the most recent

backup, which, in the worst-case scenario, assuming that the failure occurred just before (or during) the next backup, would be six hours earlier.

You can configure backup software to perform backups every hour, offering a theoretical RPO of 60 minutes. When calculating RPO, it is also important to take into account the time it takes to perform the backup. For example, suppose it takes 15 minutes to perform a backup and you back up every hour. If a failure occurs during the backup process, your best possible RPO will be one hour and 15 minutes. A realistic RPO must always balance the desired recovery time with the realities of the network infrastructure. You should not aim for an RPO of two hours, for example, when a backup itself takes three hours to complete.
The RPO also depends on the backup software technology. For example, when you use the snapshot feature in Windows Server Backup, or if you use another backup software that uses Volume Shadow Copy Service (VSS), you are backing up to the point in time when the backup was started.

- Recovery time objective (RTO). A *RTO* is the amount of time it takes to recover from failure. The RTO will vary depending on the type of failure. The loss of a motherboard on a critical server will have a different RTO than the loss of a disk on a critical server, because one of these components takes significantly longer to replace than the other.
- Retention objectives. *Retention* is a measure of the length of time you need to store backed-up data. For example, you may need to recover data quickly from up to a month ago, but need to store data, in some cases, for several years. The speed at which you agree to recover data in your SLA will depend on the age of the data, because some data is quickly recoverable and other data may need to be recovered from the archives.
- System performance. Although not directly related to data protection, system performance is also an important component of SLAs, because applications that are included in an SLA should be available, and they should also have acceptable response times to users' requests. If the system performance is slow, then business requirements will not be met.

 **Note:** Each organization's data protection SLA depends on the components that are important to the organization.

Overview of Enterprise Data Protection Strategies

When you plan backup for your enterprise, you need to develop strategies for recovering data, services, servers, and sites. You also need to make provisions for offsite backup.

You need strategies for recovering:

- Data
- Services
- Servers
- Sites
- Offsite backups

Data Protection Strategies

Data is the most commonly recovered category in an enterprise environment. This is because it is more likely that users will delete files accidentally, than it is that server hardware will fail or that applications will cause data corruption. Therefore, when you develop an enterprise data protection strategy, take into account small events, such as data deletion, in addition to big disasters, such as a site suffering a flood or fire.

When you consider data recovery strategies, backup is not the only technology approach that you can use. You can address many file and folder recovery scenarios by implementing previous versions of file functionality on file shares. You can also replicate data, and even entire virtual machines to different physical locations, or to a public or private cloud.

Service Protection Strategies

The functionality of the network depends on the availability of certain critical network services. Although well-designed networks build redundancy into core services such as Domain Name System (DNS) and Active Directory® Domain Services (AD DS), even those services might have issues, such as when a major fault is replicated that requires a restore from backup. In addition, an enterprise backup solution must ensure that services such as Dynamic Host Configuration Protocol (DHCP) and Active Directory Certificate Services (AD CS), and important resources such as file shares can be restored in a timely and up-to-date manner.

Full Server Protection Strategies

Developing a full-server recovery strategy involves determining which servers you need to be able to recover, and the RPO and RTO for critical servers. Suppose that you have a site with two computers functioning as domain controllers. When you develop your backup strategy, should you aim to have both servers capable of full server recovery with a 15-minute RPO? Alternatively, is it only necessary for one server to be recovered quickly if it fails, given that either server will be able to provide the same network service and ensure business continuity?

In developing the full server recovery component of your organization's enterprise backup plan, determine which servers are required to ensure business continuity, and ensure that these servers are backed up regularly.

Site Protection Strategies

Most large organizations have branch office sites. While it might be desirable to back up all the computers at those locations, it may not be economically feasible to do so. Developing a site recovery strategy involves determining which data, services, and servers at a specific site must be recoverable to ensure business continuity.

Offsite Backup Strategies

Many organizations that do not store offsite backups do not recover from a primary site disaster. If your organization's head office site has a fire, is subject to a once-in-a-100-year flood, an earthquake, or a cyclone, it will not matter which backup strategies you have in place if all those backups are stored at the location that was destroyed by the disaster.

A comprehensive enterprise data protection strategy involves moving backed-up data to a safe offsite location so that you can recover it regardless of the kind of disaster that occurs. This does not need to happen every day. The RPO for recovery at the offsite location—often called the *disaster recovery site*—is usually different from the RPO at the primary site. An alternative is to ensure that services are redundant across sites. The key is to be able to ensure that you are able to recover in the event of a primary site disaster.

Mitigation Strategies

No matter how prepared organizations are, they cannot prevent problems from occurring. Therefore, organizations must also develop mitigation strategies that will minimize the impact of an unexpected loss of data, a server, a service, or sites. To prepare mitigation strategies, organizations must create risk assessments that analyze all possible disaster scenarios, and document how to mitigate each of those scenarios.

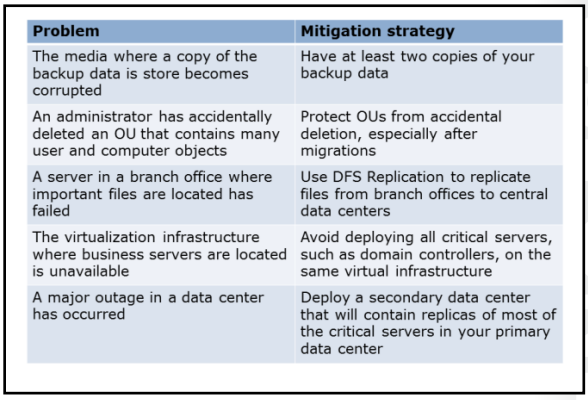

Problem	Mitigation strategy
The media where a copy of the backup data is store becomes corrupted	Have at least two copies of your backup data
An administrator has accidentally deleted an OU that contains many user and computer objects	Protect OUs from accidental deletion, especially after migrations
A server in a branch office where important files are located has failed	Use DFS Replication to replicate files from branch offices to central data centers
The virtualization infrastructure where business servers are located is unavailable	Avoid deploying all critical servers, such as domain controllers, on the same virtual infrastructure
A major outage in a data center has occurred	Deploy a secondary data center that will contain replicas of most of the critical servers in your primary data center

The following table lists some of the risks associated with data or services loss, and the appropriate mitigation strategies.

Problem	Mitigation strategy
The media where a copy of the backup data is located becomes corrupted.	Have at least two copies of your backup data, and validate your backups on a regular basis.
An administrator has accidentally deleted an organizational unit (OU) that contains many user and computer objects.	Protect OUs from accidental deletion, especially after migrations. Use recent versions of the Remote Server Administration Tools (RSATs) to ensure that newly created OUs are automatically protected from deletion.
A file server in a branch office where important files are located has failed.	Use Distributed File System (DFS) Replication to replicate files from branch offices to central data centers.
The virtualization infrastructure where business servers are located is unavailable.	Avoid deploying all critical servers—such as domain controllers—on the same virtual infrastructure.
A major outage in a data center has occurred.	Deploy a secondary data center that will contain replicas of the critical servers in your primary data center.

Best Practices when Implementing a Data Protection Strategy

When you implement a data-protection strategy, your organization should follow these best practices:

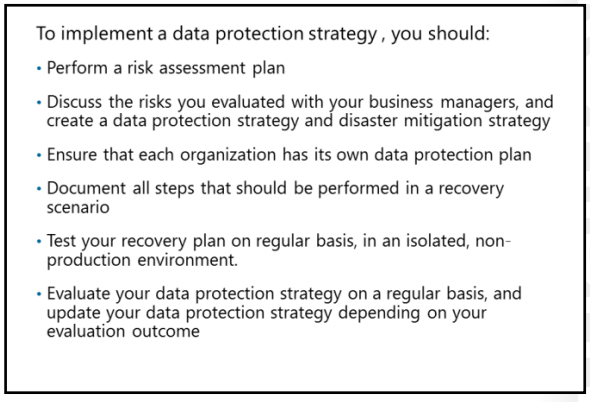

To implement a data protection strategy , you should:

- Perform a risk assessment plan
- Discuss the risks you evaluated with your business managers, and create a data protection strategy and disaster mitigation strategy
- Ensure that each organization has its own data protection plan
- Document all steps that should be performed in a recovery scenario
- Test your recovery plan on regular basis, in an isolated, non-production environment.
- Evaluate your data protection strategy on a regular basis, and update your data protection strategy depending on your evaluation outcome

- Perform a risk assessment plan. This will help you identify all of the risks associated with the availability of your organization data, servers, services, and sites.
- Discuss the risks you evaluated with your business managers. Decide together which resources should be protected with the data protection plan, and which resources should be protected with disaster mitigation, and at which level. The higher the requirements for data protection, the more expensive they are to

implement. You also want to have a low-level data protection plan for resources that are protected with disaster mitigation.

- Ensure that each organization has its own data protection plan.
- Document in detail all of the steps that should be performed in a disaster scenario.
- Test your data protection plan on a regular basis in an isolated, non-production environment.
- Use a production backup to test those recovery strategies, to ensure the backups contain valid data and to evaluate the amount of time needed to recover the amount of data.
- Evaluate your data protection plan on a regular basis, and update your data protection plan based on your evaluation.

Lesson 2
Implementing Windows Server Backup

To protect critical data, every organization must perform regular backups. Having a well-defined and tested backup strategy ensures that companies can restore data if unexpected failures or data loss occur. This lesson describes the Windows Server Backup feature in Windows Server 2012 and Windows Server 2012 R2, and the Windows Azure™ Online Backup for Windows Server 2012.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe data and service information that needs to be backed up in a Windows Server environment.
- Describe the backup types.
- Describe backup technologies.
- Explain how to plan for backup capacity.
- Explain how to plan backup security.
- Describe Windows Server Backup.
- Explain how to configure a scheduled backup using Windows Server Backup.
- Describe the Windows Azure Backup.
- Describe the considerations for an enterprise backup solution.
- Summarize the features available with Microsoft® System Center 2012 Data Protection Manager.

What Needs to Be Backed Up?

When you plan backups across your organization, ensure that you protect resources that are mission critical, such as:

- Critical resources
- Backup verification
- Backup security
- Compliance and regulatory requirements

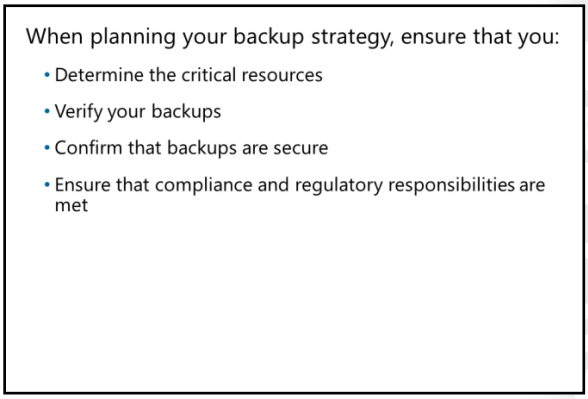

Determining Critical Resources to Back Up

In an ideal scenario, you would back up everything and restore data instantly as it existed at a particular point in time from any point in the last several years. In reality, such a backup strategy would be an expensive solution. Therefore, the first step in planning backup across the enterprise is to determine what exactly needs to be backed up.

For example, should you back up every domain controller in the domain, given that Active Directory information will be replicated back to a replacement domain controller as soon as it is promoted? Is it necessary to back up every file server in all file shares, if every file is replicated to multiple servers through a distributed file system?

You also need to distinguish between technical reasons and regulatory reasons for backing up data. Due to legal requirements, you may need to be able to provide your business with business-critical data for the previous 10 years or even longer.

To determining what to back up, consider the following:

- If the data is only stored in one place, ensure that it is backed up.
- If data is replicated, it may not be necessary to back up each replica. However, you must back up at least one location to ensure that the backup can be restored. A better strategy is to back up at multiple locations.
- Is the server or data a mission-critical component?
- If this server or disk failed, or if this data became corrupted, what steps would be required to recover it?

Many organizations ensure the availability of critical services and data through redundancy. For example, Microsoft Exchange Server® 2013 provides continuous replication of mailbox databases to other servers through a technology called Database Availability Groups (DAGs). While the use of DAGs does not mean that an organization should not back up its Exchange Server 2013 Mailbox servers, it does change how an organization should think about backing up its Mailbox servers or centralizing its backup strategies.

Verifying Your Backups

Performing a backup, and ensuring that the backup contains everything that you need, are two different tasks. You need to have a method for verifying that each backup has completed successfully. You also need to know when backups have failed. At a minimum, this will mean checking the logs on each server to determine whether a failure has occurred. If you have configured backups to occur on each server every six hours, how often should you check the logs? A better solution is to employ an alert mechanism to alert you in the event that a backup fails; such a mechanism is available in System Center 2012 R2 Operations Manager. The point is to avoid discovering that your backups for a particular server have failed just when you need those backups to perform a recovery.

One way of verifying backups is to perform regular testing of the recovery procedures, in which you simulate a particular failure. This allows you to verify the integrity of the data that you are using to perform a recovery, and that the recovery procedures that you have in place effectively resolve the failure. It is better to discover that you need to add steps to your recovery procedure during a test, rather than during an actual failure.

Confirming That Backups Are Secure

By definition, a good set of backups contains all of your organization's critical data. This data needs to be protected from unauthorized access. Although data might be protected by permissions and access controls while it is hosted on servers in a production environment, anyone who has access to the media that hosts that backup data can restore it. For example, some products, such as Windows Server Backup, do not allow administrators to encrypt backup data. This means that physical security is the only way that you can ensure that critical data does not end up in the hands of unauthorized users.

When developing an enterprise backup strategy, ensure that backup data is stored in a secure location.

You might also consider using backup software that allows you to split the backup-and-restore roles so that users who have permissions to back up data do not have permissions to restore that data, and users who have permissions to restore data do not have permissions to back it up.

Ensuring That Compliance and Regulatory Responsibilities Are Met

Systems administrators should be aware of what the organization's regulatory and compliance responsibilities are with respect to the archiving of data. For example, some jurisdictions require that business-relevant email message data be retained for a period of up to seven years. Unfortunately,

regulatory requirements vary from country to country, and even from state to state within the United States. When you develop your organization's data protection strategy, you should schedule a meeting with your legal team to determine precisely which data needs to be stored, and for how long.

Backup Types

When you use Windows Server Backup, you can perform the following types of backups:

- A full backup is a block-level replica of all blocks on all the server's volumes
- An incremental backup is a copy of only those blocks that have changed since the last full or incremental backup

- Full backup. A full backup is a block-level replica of all blocks on all of the server's volumes. Rather than copying files and folders to backup media, the underlying blocks are copied across to the backup media.
- Incremental backup. An incremental backup is a copy of only those blocks that have changed since the last full or incremental backup. During an incremental backup, these blocks are copied across to the backup media. When this process completes, the blocks are then marked as backed up. During recovery, the original set of blocks is restored. Then, each set of incremental blocks are applied, bringing the recovered data back to the appropriate state in a consistent manner.

Backup Technologies

Most backup products in use today use the VSS infrastructure that is present in Windows Server 2012 and Windows Server 2012 R2. Some older applications, however, use streaming backup. It may be necessary to support such older applications in complex, heterogeneous environments.

- The VSS backup technology solves data consistency issues by creating shadow copies
- You can also use streaming backups for older applications that are not VSS-aware
- Hyper-V replica provides you with a disaster recovery option by replicating a consistent copy of a virtual machine to another server or site.

One challenge of performing backups is to ensure consistency of the data that you are backing up. Backups do not occur instantly; they can take seconds, minutes, or hours. Unfortunately, servers are not static, and the state of a server at the beginning of a backup might not be the same state that the server is in when the backup completes. If you do not take consistency into account, this can cause problems during restoration because the configuration of the server may have changed during the backup.

VSS

VSS—a technology that Microsoft included with the server operating system since Windows Server 2003 R2, and which is present in all newer server operating systems—solves the consistency problem at the disk-block level by creating what is known as a shadow copy. A *shadow copy* is a backup of the file table, which also marks all used blocks as un-updateable. Whenever write requests occur after the snapshot is taken, the old blocks are compressed and stored before the blocks data is changed. This enables you to have an in-time view of the file system. When a backup occurs, the old blocks are backed up, which means that any changes that might have occurred since the freeze are not backed up.

Creating a shadow copy tells the operating system first to put all files, such as DHCP databases and Active Directory database files, in a consistent state for a moment. Then the current state of the file system is recorded at that specific point in time. After VSS creates the shadow copy, all write accesses that would overwrite data store the previous data blocks first. Therefore, a shadow copy is small in the beginning, and it grows over the time as data changes. By default, the operating system is configured to reserve 12 percent of the volume for VSS data, and VSS automatically deletes older snapshots when this limit is reached. You can change this default value, and you can change the default location of the VSS data. This ensures that the backup has a snapshot of the system in a consistent state, no matter how long it actually takes to write the backup data to the backup storage device.

Streaming Backup

Streaming backup is often used by older applications that do not use VSS. You back up applications that are not VSS-aware by using a method known as a *streaming backup*. In contrast to VSS, where the operating system ensures that data is kept in a consistent state and at a current point in time, when you use streaming backup, the application or the data protection application is responsible for ensuring that the data remains in a consistent state. In addition, after streaming backup completes, some files have the state they had in the beginning of the backup, while other files have the state of the end of the backup window.

Hyper-V Replica

Hyper-V® in Windows Server 2012 and Windows Server 2012 R2 supports creating replicas of virtual machines. These replicas can be stored on another server in the same site, on another server in another site, or even in a public cloud. Virtual machine replication allows you to have consistent versions of production virtual machines stored in a separate location. While Hyper-V replica does allow you to keep copies of virtual machines that are nearly up to date (there is always some lag involved when replicating across sites) a replica virtual machine only protects you against some types of failures. If an application running on the virtual machine or data hosted on the virtual machine becomes corrupt, but the virtual machine remains operational, it is likely that the corrupted files will also be replicated across to the replica virtual machine.

Planning Backup Capacity

When you develop an enterprise recovery strategy, you need to determine how much storage capacity your organization will require for backups. The following factors affect the amount of space that is required to store backup data:

- Space requirements for a full backup
- Space requirements for an incremental backup
- Amount of time required to back up
- Backup frequency
- Backup retention

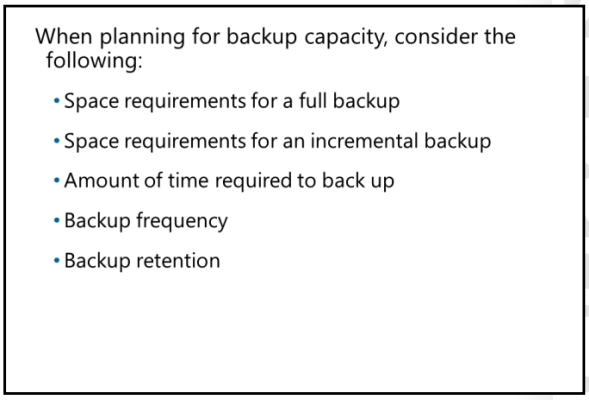

Full Backup Requirements

To calculate the space required for a full backup, determine how much space you will need to back up from all volumes. If the server has a dedicated drive for backups, you would not perform a backup on that drive.

With products that perform image-based backups, such as Windows Server Backup, this data is not compressed. On some types of servers, notably file servers, the amount of space required for a full backup tends to grow over time. You can lessen this tendency by using file expiration policies such as those found in File Server Resource Manager (FSRM).

Incremental Backup Requirements

An incremental backup on Windows Server Backup stores all of the hard disk blocks that have changed since the last full or incremental backup. Incremental backups are substantially faster than full backups and require less space. The downside of incremental backups is that they can require greater recovery time, especially if multiple incremental backups need to be restored after the last full backup.

Amount of Time Required to Back up

The amount of time required to write data from the server being backed up to the backup storage device can have an impact on projected RPO, because we do not recommend that you begin a new backup operation before you complete the current one.

Backup Frequency

Backup frequency is a measure of how often backups are taken. With incremental block-level backups, no substantial difference will exist between the amount of data written over the sum of four 30-minute sessions and one two-hour incremental session on the same server. This is because over the two hours, the same number of blocks will have changed on the server as the four 30-minute sessions. However, the four 30-minute sessions have broken it up into smaller parts. When backups occur more frequently, they reduce the time required to perform the backup by splitting it into smaller parts. The overall total will be about the same.

Backup Retention

When you attempt to determine the required backup capacity, you should determine precisely how long you need to retain backup data. For example, if you need to be able to recover to any backup point in the last 28 days, and if you have recovery points generated every hour, you will need more space than if you have recovery points generated once a day and you only need to restore data from the last 14 days.

Planning Backup Security

When planning your backup security, consider the following:

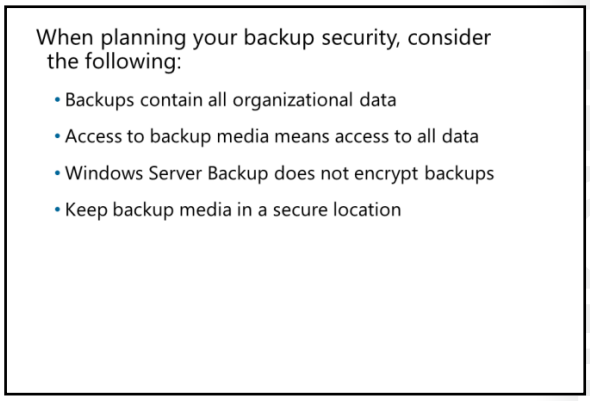

- Backups contain all organizational data. By nature, backups will contain all the data necessary to ensure your organization's continued ability to function in the event of failure. Because this data is likely to contain sensitive information, you should protect it with the same level of diligence with which it is protected when hosted on the server.
- Access to backup media means access to all data. If feasible, use administrative role separation to ensure that the users who back up the data are not the users who can restore it. In high-security environments, ensure that the backup-and-restore operations are properly audited so that you can track backups, and restore function activity.
- Windows Server Backup does not encrypt backups. Windows Server Backup writes backups in VHD format. This means that anyone who has access to Windows® 8, Windows 8.1, Windows Server 2012,

or Windows Server 2012 R2 can mount those backups as volumes, and then extract data from them. An even more sophisticated attack might include booting into the backup VHD to impersonate the backed up system on the organizational network.

- Keep backup media in a secure location. At a minimum, backups should be kept locked up in a secure location. If your organization is backing up to disk drives that are attached to servers by USB cable, ensure that those disk drives are locked in place, even if they are located in a secure server room, and even if your organization's server room has a security camera.

What Is Windows Server Backup?

The Windows Server Backup feature in Windows Server 2012 and Windows Server 2012 R2 that consists of a Microsoft Management Console (MMC) snap-in, the command **wbadmin**, and Windows PowerShell® commands. You can use the wizards in the Windows Server Backup feature to guide you through running backups and recoveries.

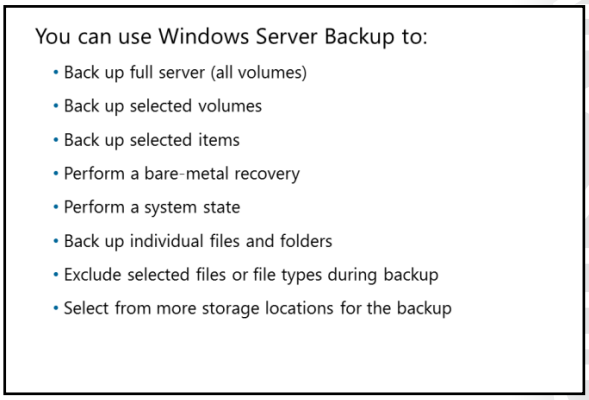

You can use Windows Server Backup to back up:

- Full server (all volumes).
- Selected volumes.
- Select specific items for backup, such as specific folders or the system state.

In addition, Windows Server Backup allows you to:

- Perform a bare-metal restore. A bare-metal backup contains at least all critical volumes, and allows you to restore without first installing an operating system. You do this by using the product media on a DVD or USB key, and the Windows Recovery Environment (Windows RE). You can use this backup type together with the Windows RE to recover from a hard disk failure, or if you have to recover the whole computer image to new hardware.
- Use system state. The backup contains all information to roll back a server to a specific point in time. However, you need an operating system installed prior to recovering the system state.
- Recover individual files and folders or volumes. The Individual files and folders option enables you to select to back up and restore specific files, folders, or volumes, or you can add specific files, folders, or volumes to the backup when you use an option such as critical volume or system state.
- Exclude selected files or file types. For example, you can exclude temporary files from the backup.
- Select from more storage locations. You can store backups on remote shares or non-dedicated volumes.

If there are events such as hard disk failures, you can perform system recovery by using a full server backup and Windows RE. This will restore your complete system onto the new hard disk.

Windows Server Backup is a single-server backup solution. You cannot use one instance of Windows Server Backup to back up multiple servers. You would need to install and configure Windows Server Backup on each server.

Demonstration: Configuring a Scheduled Backup

In this demonstration, you will see how to configure Windows Server Backup to perform a scheduled backup of specific folders that includes a filter to exclude specific file types.

Demonstration Steps

1. On LON-SVR1, start Windows Server Backup.
2. Configure the backup schedule with the following options:
   - Backup Configuration: Custom
   - Select Items for Backup: **C:\HR Data**
   - Add Exclusion: **C:\HR Data\Old HR file.txt**
   - Backup Time: **Once a day, 1:00 AM**
   - Destination Type: **Back up to a shared network folder**
   - Remote Shared Folder: **\\LON-DC1\Backup**:
     - Register Backup Schedule: Username: **Administrator**
     - Password: **Pa$$w0rd**
3. Run the Backup Once Wizard using the scheduled backup options.
4. Close Windows Server Backup.

What Is Windows Azure Backup?

Windows Azure Backup is a cloud-based backup solution for Windows Server 2012 and Windows Server 2012 R2. You can use the Windows Azure cloud subscription service to provide off-site protection against data loss caused by disasters. You back up files and folders, and then recover them from the public or private cloud as needed. You can use Windows Azure Backup to back up and protect critical data from any location.

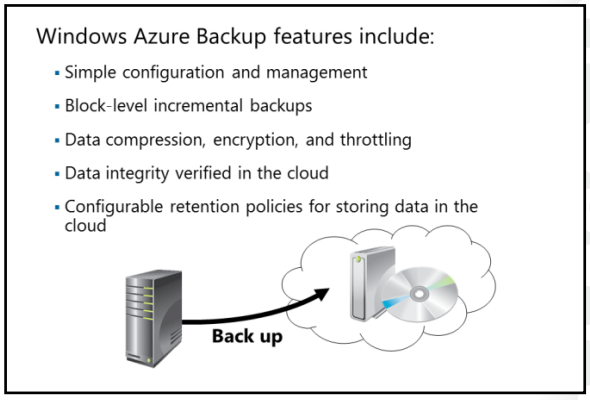

Windows Azure Backup is built on the Windows Azure platform, and uses Windows Azure Blob storage for storing customer data. Windows Server 2012 uses the downloadable Windows Azure Backup Agent to transfer file and folder data securely to the Windows Azure Backup. After you install the Windows Azure Backup Agent, the agent integrates its functionality through the Windows Server Backup interface. You can download Windows Azure Backup Agent from the Microsoft website.

Key Features

The key features of Windows Azure Backup include:

- Simple configuration and management, including:
  - Simple user interface to configure and monitor backups.
  - Integrated recovery experience to recover files and folders from a local disk or from a cloud platform.

  - Easy data recoverability for data that was backed up onto any server of your choice.
  - Scripting capability that is provided by the Windows PowerShell command-line interface.
- Block-level incremental backups. The Windows Azure Backup Agent performs incremental backups by tracking file and block-level changes, and only transferring the changed blocks, which reduces the storage and bandwidth usage. Different point-in-time versions of the backups use storage efficiently by only storing the changed blocks between these versions.
- Data compression, encryption, and throttling. The Windows Azure Backup Agent ensures that data is compressed and encrypted on the server before it is sent to the Windows Azure Backup on the network. Therefore, the Windows Azure Backup only stores encrypted data in cloud storage. The encryption passphrase is not available to the Windows Azure Backup, and therefore, the data is never decrypted in the cloud. In addition, users can set up throttling and configure how the Windows Azure Online uses the network bandwidth when it backs up or restores information.
- Data integrity verified in the cloud. In addition to the secure backups, the backed-up data is also checked automatically for integrity after the backup completes. Therefore, any corruptions that may arise because of data transfer can be easily identified. These corruptions are fixed automatically in the next backup.
- Configurable retention policies for storing data in the cloud. The Windows Azure Backup accepts and implements retention policies to recycle backups that exceed the desired retention range, thereby meeting business policies and managing backup costs.

Windows Azure Backup can only be used to back up files and folders. You cannot use Windows Azure Backup to backup system state data or perform a full-server or volume backup and recovery, although you can back up all files and folders on a volume. Windows Server Backup allows you to back up a maximum of 850 gigabytes (GB) per volume during a backup session.

While each instance of Windows Azure Backup can back up to the same recovery vault in Windows Azure, you must install and configure each instance of Windows Azure Backup separately. You cannot centrally manage multiple instances of Windows Azure Backup.

To learn more about Windows Azure SQL Database, go to:

http://go.microsoft.com/fwlink/?LinkId=270041

At this time, Windows Azure Backup is not available in all countries. For updated information, go to:

http://go.microsoft.com/fwlink/?LinkID=386645

Considerations for an Enterprise Backup Solution

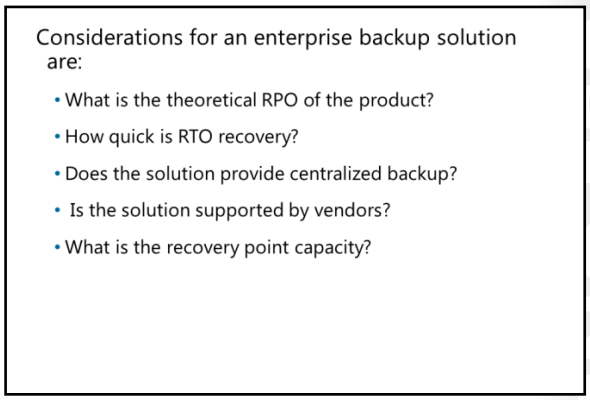

Windows Server Backup and Windows Azure Backup are single-server backup solutions. When you plan backup for an enterprise, consider the following points:

- Maximum amount of data lost. What is the theoretical RPO of the product? Products that offer restoration closer to the point of the failure are likely to cost more than products that offer 15-minute or 30-minute RPOs. You need to determine your organization's needs. Does your organization need to be able to recover to the last SQL Server transaction, or is a 15-minute recovery window an acceptable compromise?
- How quick is RTO recovery? How long does it take to go from failure to restored functionality? Being able to restore to the last SQL Server transaction is the optimal solution, but if it takes two days to recover to that point, the solution is not as helpful as it may appear.
- Does the solution provide centralized backup? Does the product allow you to centralize your backup solution on one server, or must backups be performed directly on each server in the organization?
- Is the solution supported by vendors? Some vendors use undocumented application programming interfaces (APIs) to back up and recover specific products, or to back up files without ensuring that the service is at a consistent state.
- Is the backup solution compatible with your applications? For example, a new update to a product may make your backup solution incompatible with the application. Check with the application vendor to determine whether the enterprise backup solution is supported.
- Recovery-point capacity. Determine the product's recovery-point capacity. How many restore points does the enterprise data protection solution offer, and is this adequate for your organization's needs?

What Is Data Protection Manager?

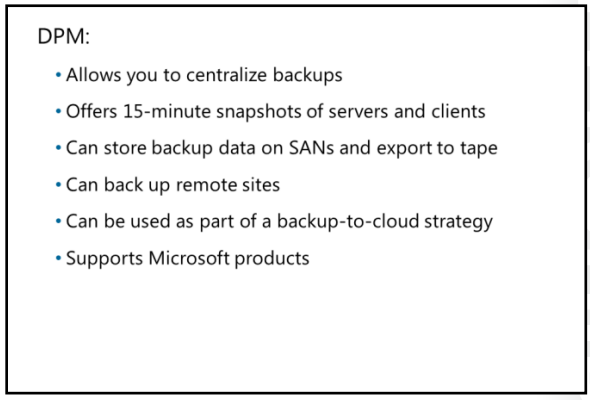

Data Protection Manager (DPM) is a Microsoft System Center enterprise data protection and recovery product. DPM has the following features:

- Backup centralization. DPM uses a client/server architecture, in which the client software is installed on all the computers that are to be backed up. Those clients stream backup data to the DPM server. This allows each DPM server to support entire small to medium-sized organizations. You can also manage multiple DPM servers from one centralized DPM console.
- 15-minute RPO. DPM allows 15-minute snapshots of supported products. This includes most of the Microsoft enterprise suite of products, including Windows Server with its roles and services, Exchange Server, Hyper-V, and Microsoft SQL Server.

- Microsoft workload support. DPM was designed specifically by Microsoft to support Microsoft applications such as Exchange Server, SQL Server, and Hyper-V. However, DPM has not been specifically designed to support non-Microsoft server applications that do not have consistent states on disk, or that do not support VSS.
- Disk-based backup. DPM can perform scheduled backups to disk arrays and storage area networks (SANs). You can also configure DPM to export specific backup data to tape for retention and compliance-related tasks.
- Remote-site backup. DPM uses an architecture that allows it to back up clients that are located in remote sites. This means that a DPM server that is located in a head office site can perform backups of servers and clients that are located across wide area network (WAN) links.
- Backup-to-cloud strategy support. DPM supports backup of DPM servers to a cloud platform. This means that a DPM server at a cloud-based hosting facility can be used to back up the contents of a head office DPM server. For disaster redundancy, you can also configure DPM servers to back up each other.

Lesson 3
Implementing Server and Data Recovery

Recovering servers and data requires well-defined and documented procedures that administrators can follow when failures occur. The recovery process also requires knowledge of the backup-and-restore hardware and software, such as DPM, and tape library devices.

This lesson describes how to restore data and servers by using the Windows Server Backup feature in Windows Server 2012, and Windows Azure Backup in Windows Server 2012.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the options for server recovery.
- Describe the options for server restore.
- Describe the options for data recovery.
- Explain how to perform a restore with Windows Server Backup.
- Explain how to perform a restore with Windows Azure Online.

Options for Server Recovery

Windows Server Backup provides the following recovery options:

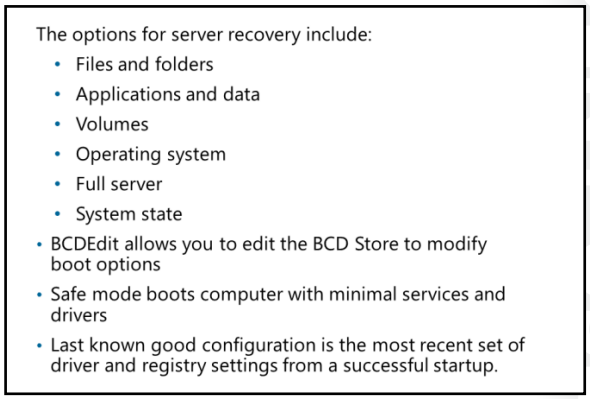

- Files and folders. You can back up individual files or folders as long as the backup is on a separate volume or in a remote shared folder.
- Applications and data. You can recover applications and data if the application has a VSS writer, and is registered with Windows Server Backup.
- Volumes. When you restore a volume, the procedure always restores all the contents of the volume. When you choose to restore a volume, you cannot restore individual files or folders.
- Operating system. You can recover the operating system through the Windows Recovery Environment (Windows RE), the product DVD, or a USB flash drive.
- Full server. You can recover the full server through Windows RE.
- System state. *System state* creates a point-in-time backup that you can use to restore a server to a previous working state.

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

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

- **Conflict Resolution**. When you restore data from a backup, it frequently conflicts with existing versions of the data. Conflict resolution allows you to determine how to handle those conflicts. When conflicts occur, you have the following options:
  - **Create copies and retain both versions.**
  - **Overwrite existing version with recovered version.**
  - **Do not recover items if they already exist in the recovery location.**
- **Security Settings**. Use this option to restore permissions to the data that is being recovered.

BCDEdit and the BCD Store

The Boot Configuration Data store (BCD) stores boot applications and boot application settings. During recovery, it may be necessary to modify the BCD store using the BCDEdit utility, either to enable debugging options or to configure the server to boot to VHD or another environment to allow data to be recovered from Windows Server 2012 or Windows Server 2012 deployment that has become corrupted in such a way that it does not boot. You can also use BCDEdit to enable Emergency Management Services (EMS) for a specific boot option. EMS allows out-of-band management of Windows Server 2012 and Windows Server 2012 R2 in the event that the server remains operational, but you cannot interact with the server using the keyboard and mouse or establish remote connections. When configured, EMS allows connections through a terminal program connected to a server's USB port. When connected, you can run a limited set of commands, allowing you to terminate process or initiate server shutdown; this is preferable to turning off the server's power, which may cause damage.

Safe Mode and Last Known Good Configuration

Safe mode is a boot option available by pressing F8 when the server starts. Safe mode allows you to start the computer with a minimal set of services and drivers, and is useful when a service or driver may be blocking a server from starting properly. Using the Safe mode with networking option includes services and drivers that allow network connectivity. The Safe mode with command prompt is a special version of safe mode in which a command prompt is available, but the Windows Desktop is not loaded.

You can also use Last Known Good Configuration as a boot option. This is the most recent set of driver and registry settings that allowed the server to start up and users to sign on. When a user signs on successfully for the first time after a server starts up, a new Last Known Good Configuration is written.

Options for Server Restore with Windows RE

The Windows RE is a special environment that allows you to perform server repair tasks including performing full server recovery. You can enter Windows RE by taking the following steps:

- Issue the shutdown /r /0 command from a command prompt or Windows PowerShell.
- Hold the Shift key when clicking the Restart option on the Settings charm.
- Boot from the installation or recovery media.

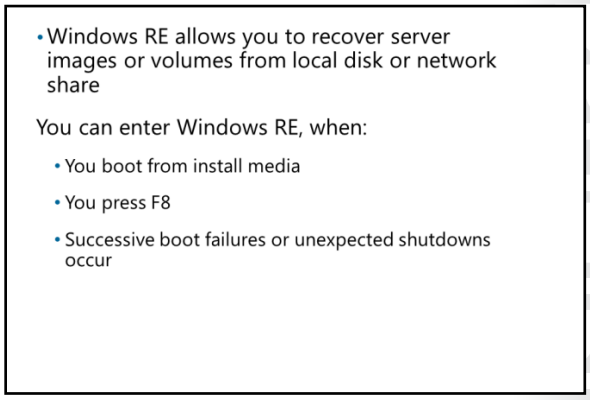

A server will also automatically enter Windows RE under the following circumstances:

- Two successive failed attempts to start Windows Server 2012 or Windows Server 2012 R2 occur.
- Two successive unexpected shutdowns occur within 120 seconds of successful boot.
- A secure boot error occurs.

You can use Windows RE to recover volumes or server images from locally attached disks or from network locations.

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

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

Options for Data Recovery

Data is the most frequently recovered component of an IT infrastructure. This is because users may accidentally delete data, and will need you to recover that data. There are several strategies that you can pursue when you are developing a data recovery procedure. You can:

- Allow users to recover their own data.
- Perform a recovery to an alternative location.
- Perform a recovery to the original location.
- Perform a full-volume recovery.

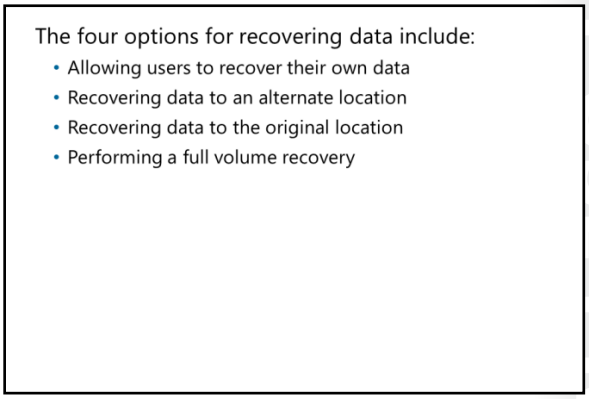

Users Recover Their Own Data

The most common data recovery performed by IT departments is the recovery of files and folders that users have deleted, lost, or in some way corrupted. The Previous Versions of Files functionality that was introduced in Windows Server 2003, (which you can also enable on all computers running Windows Server 2012 and Windows Server 2012 R2) lets users recover their own files using the file or folder properties right from their workstation. After end-users are trained how to do this, the IT department spends less time recovering user data, which allows them to focus on more valuable tasks.

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

Recover Data to an Alternative Location

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

When you perform a recovery to an alternative location, always ensure that permissions are also restored. A common problem occurs when administrators recover data that includes restricted material, to a location where permissions are not applied, thereby enabling unintended access to data for users who should not have access. Similarly, you want to ensure that users who should have access to data are able to access it.

Recover Data to the Original Location

During some types of failures, such as data corruption or deletion, you will have to restore data to the original location. This is the case when applications or users who access the data are preconfigured with information about where the data is located.

Recover a Volume

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

 **Note:** During the restore process, you should copy event logs before you start the restore process. If you overwrite the event log files—for example, with a system recovery—you will be not able to read event-log information that occurred before the restore started. That data could lead you to information about what caused the issue.

Demonstration: Using Windows Server Backup to Restore a Folder

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

Demonstration Steps

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

Restoring with Windows Azure Backup

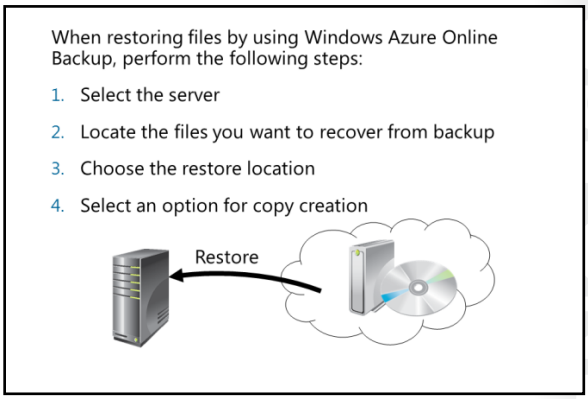

You can use the Windows Azure Backup to back up only Windows Server 2012 servers. However, you do not have to restore data on to the same server from which you backed it up.

You can recover files and folders by using both Windows Azure Backup Microsoft Management Console (MMC) in Server Manager, or by using the Windows PowerShell command-line interface. To use the Windows Azure Backup MMC, perform the following steps:

1. Select the server on which backup data was created originally. This server could be a local server or another server. If you select the option for another server, you must provide your Windows Azure Backup administrator credentials.
2. Browse for files that have to be restored, or you can search for them in the Windows Azure Backup.
3. After you locate the files, select them for recovery, and select a location where the files will be restored.
4. When you restore files, select one of the following options:
    - Create copies so that you have both the restored file and original file in the same location. The restored file's name will be in the following format: *Recovery Date*+Copy of+*Original File Name*.
    - Overwrite the existing versions with the recovered version.
    - Do not recover the items that already exist on the recovery destination.

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

Lab: Implementing Windows Server Backup and Restore

Scenario

Much of the data that is stored on the A. Datum Corporation's network is extremely valuable to the organization. Losing this data would be a significant loss to the organization. Additionally, many of the servers that are running on the network provide extremely valuable services for the organization, which means that losing these servers for a significant time would also result in losses to the organization. Because of the significance of the data and services, it is critical that they can be restored in the event of disaster.

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

As one of the senior network administrators at A. Datum, you are responsible for planning and implementing a data-protection solution that will ensure that critical data and services can be recovered in the event of any type of failure. You need to implement a backup-and-restore process that can recover lost data and services.

Lab Setup

Estimated Time: 60 minutes

Virtual machines: 20412C-LON-DC1,

20412C-LON-SVR1

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 **20412C-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 **20412C-LON-SVR1**.

Exercise 1: Backing Up Data on a Windows Server 2012 R2 Server

Scenario

The LON-SVR1 server contains financial data that must be backed up on a regular basis. This data is critical to the organization. You decided to use Windows Server Backup to back up critical data. You will to install this feature and configure scheduled backups.

The main tasks for this exercise are as follows:

- Install Windows Server Backup.
- Configure a scheduled backup.
- Complete an on-demand backup.

▶ Task 1: Install Windows Server Backup

1. Switch to LON-SVR1.
2. From the Server Manager, install the Windows Server Backup feature. Accept the default values on the Add Roles and Features Wizard.

▶ Task 2: Configure a scheduled backup

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

 **Note:** In a production environment, you will not store backup to a domain controller. You do it here for lab purposes only.

▶ Task 3: Complete an on-demand backup

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

Results: After you complete this exercise, you will have configured the Windows Server Backup feature, scheduled a backup task, and completed an on-demand backup.

Exercise 2: Restoring Files Using Windows Server Backup

Scenario

To ensure that the financial data can be restored, you must validate the procedure for restoring the data to an alternate location.

The main tasks for this exercise are as follows:

- Delete a file from the server.
- Restore a file from backup.
- Prepare for the next module.

▶ Task 1: Delete a file from the server

- On LON-SVR1, open Windows Explorer, and then delete the **C:\Financial Data** folder.

▶ Task 2: Restore a file from backup

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

▶ Task 3: Prepare for the next module

When you finish the lab, revert the virtual machines to their initial state.

1. On the host computer, start the Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20412C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps two and three for 20412C-LON-SVR1.

Results: After completing this exercise, you will have tested and validated the procedure for restoring a file from backup.

Question: You are concerned about business-critical data that is located on your company's servers. You want to perform backups every day, but not during business hours. What should you do?

Question: Users report that they can no longer access data that is located on the server. You connect to the server, and you realize that the shared folder where users were accessing data is missing. What should you do?

Module Review and Takeaways

Best Practice

- Analyze your important infrastructure resources and mission-critical and business-critical data. Based on that analysis, create a backup strategy that will protect the company's critical infrastructure resources and business data.
- Work with the organization's business managers to identify the minimum recovery time for business-critical data. Based on that information, create an optimal restore strategy.
- Always test backup-and-restore procedures regularly. Perform testing in a non-production and isolated environment.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
The server has suffered a major failure on its components.	

Review Questions

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

Question: How frequently should you perform backups on critical data?

Real-world Issues and Scenarios

Your organization needs information about which data to back up, how frequently to back up different types of data and technologies, where to store backed up data (onsite or in the cloud), and how fast it can restore backed-up data. How would you improve your organization's ability to restore data efficiently when it is necessary?

Answer: Your company should develop backup-and-restore strategies based on multiple parameters, such as business-continuity needs, risk-assessment procedures, and critical resource and data identification. You must develop strategies that should then be evaluated and tested. These strategies should take into consideration the dynamic changes that are occurring with new technologies, and the changes that may occur with your organization's growth.

Tools

Tool	Use	Where to find it
Windows Server Backup	Perform on-demand or scheduled backup and restore of data and servers.	Server Manager Tools
Windows Azure Backup	Perform on-demand or scheduled backup to the cloud, and restore data from the backup located in the cloud.	Server Manager Tools

Course Evaluation

Keep this evaluation topic page if this is the final module in this course. Insert the Product_Evaluation.ppt on this page.

If this is not the final module in the course, delete this page

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.

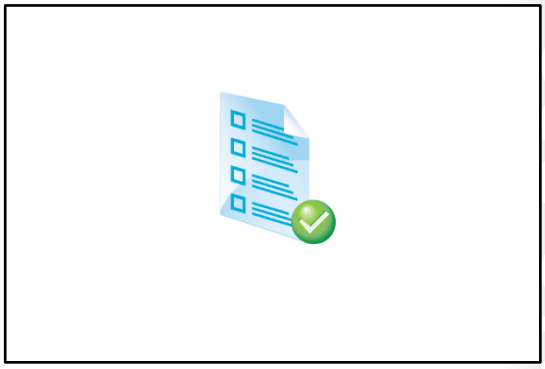

Microsoft will keep your answers to this survey private and confidential and will use your responses to improve your future learning experience. Your open and honest feedback is valuable and appreciated.

Module 1: Implementing Advanced Network Services

Lab: Implementing Advanced Network Services

Exercise 1: Configuring Advanced DHCP Settings

▶ Task 1: Configure a superscope

1. On LON-DC1, in Server Manager, click **Tools**, and then click **DHCP**.
2. In the DHCP console, click **lon-dc1.adatum.com**, select and then right-click **IPv4**, and then click **New Scope**.
3. In the New Scope Wizard, click **Next**.
4. On the **Scope Name** page, in the **Name** box, type **Scope1**, and then click **Next**.
5. On the **IP Address Range** page, in the **Start IP address** box, type **192.168.0.50**, and then in the **End IP address** box, type **192.168.0.100**.
6. In the **Subnet mask** box, ensure that **255.255.255.0** is entered, and then click **Next**.
7. On the **Add Exclusions and Delay** page, click **Next**.
8. On the **Lease Duration** page, click **Next**.
9. On the **Configure DHCP Options** page, select **Yes, I want to configure these options now**, and then click **Next**.
10. On the **Router (Default Gateway)** page, in the **IP address** box, type **192.168.0.1**, click **Add**, and then click **Next**.
11. On the **Domain Name and DNS Servers** page, ensure that the parent domain is **Adatum.com**, and then click **Next**.
12. On the **WINS Servers** page, click **Next**.
13. On the **Activate Scope** page, click **No, I will activate this scope later**, and then click **Next**.
14. On the **Completing the New Scope Wizard** page, click **Finish**.
15. Right-click **IPv4**, and then click **New Scope**.
16. In the New Scope Wizard, click **Next**.
17. On the **Scope Name** page, in the **Name** box, type **Scope2**, and then click **Next**.
18. On the **IP Address Range** page, in the **Start IP address** box, type **192.168.1.50**, and then in the **End IP address** box, type **192.168.1.100**.
19. In the **Subnet mask** box, ensure that **255.255.255.0** is entered, and then click **Next**.
20. On the **Add Exclusions and Delay** page, click **Next**.
21. On the **Lease Duration** page, click **Next**.
22. On the **Configure DHCP Options** page, select **Yes, I want to configure these options now**, and then click **Next**.
23. On the **Router (Default Gateway)** page, in the **IP address** box, type **192.168.1.1**, click **Add**, and then click **Next**.

24. On the **Domain Name and DNS servers** page, ensure the parent domain is **Adatum.com**, and then click **Next**.
25. On the **WINS Servers** page, click **Next**.
26. On the **Activate Scope** page, click **No, I will activate this scope later**, and then click **Next**.
27. On the **Completing the New Scope Wizard** page, click **Finish**.
28. Right-click the **IPv4** node, and then click **New Superscope**.
29. In the New Superscope Wizard, click **Next**.
30. On the **Superscope Name** page, in the **Name** box, type **AdatumSuper**, and then click **Next**.
31. On the **Select Scopes** page, select **Scope1**, hold down the Ctrl key, select **Scope2**, and then click **Next**.
32. On the **Completing the New Superscope Wizard** page, click **Finish**.
33. In the DHCP console, under IPv4, select and then right-click **Superscope Adatum Super**, and then click **Activate**.

▶ Task 2: Configure DHCP name protection

1. On LON-DC1, in the DHCP console, expand **Lon-DC1.adatum.com**.
2. Right-click **IPv4**, and then click **Properties**.
3. In the **IPv4 Properties** dialog box, click the **DNS** tab.
4. In the Name Protection pane, click **Configure**.
5. Select the **Enable Name Protection** check box, and then click **OK**.
6. Click **OK** again.

▶ Task 3: Configure and verify DHCP failover

1. On LON-SVR1, in Server Manager, click **Tools**, and then from the drop-down list, click **DHCP**. Note that the server is authorized, but that no scopes are configured.
2. On LON-DC1, in the DHCP console, right-click the **IPv4** node, and then click **Configure Failover**.
3. In the Configure Failover Wizard, click **Next**.
4. On the **Specify the partner server to use for failover** page, in the **Partner Server** box, type **172.16.0.21**, and then click **Next**.
5. On the **Create a new failover relationship** page, in the **Relationship Name** box, type **Adatum**.
6. In the **Maximum Client Lead Time** field, set the hours to **0**, and set the minutes to **15**.
7. Ensure that the **Mode** field is set to **Load balance**, and that the **Load Balance Percentage** is set to **50%**.
8. Select the **State Switchover Interval** check box. Keep the default value of 60 minutes.
9. In the **Enable Message Authentication Shared Secret** box, type **Pa$$w0rd**, and then click **Next**.
10. Click **Finish**, and then click **Close**.
11. On LON-SVR1, refresh the IPv4 node, and then note that the IPv4 node is active.
12. Expand the IPv4 node, expand **Scope [172.16.0.0] Adatum**, click the **Address Pool** node, and note that the address pool is configured.
13. Click the **Scope Options** node, and note that the scope options are configured.

14. Start **20412C-LON-CL1**, and then sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
15. On the Start screen, type **Control Panel**.
16. In the **Apps Results** box, click **Control Panel**.
17. In Control Panel, click **Network and Internet**, click **Network and Sharing Center**, click **Change adapter settings**, right-click **Ethernet**, and then click **Properties**.
18. In the **Ethernet Properties** dialog box, click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
19. In the **Properties** dialog box, select the **Obtain an IP address automatically** radio button, click **Obtain DNS server address automatically**, and then click **OK**.
20. In the **Ethernet Properties** dialog box, click **Close**.
21. Hover over the bottom right corner to expose the fly-out menu, and then click the **Search** charm.
22. In the **Apps** search box, type **Cmd**, and then press Enter.
23. In the command prompt window, type **ipconfig**, and then press Enter. Record your IP address.
24. On **LON-DC1**, on the taskbar, click the **Server Manager** icon.
25. In Server Manager, click **Tools**, and then click **Services**.
26. In the Services window, right-click the **DHCP Server** service, and then click **Stop** to stop the service.
27. Close the Services window, and close the DHCP console.
28. On **LON-CL1**, in the command prompt window, type **ipconfig /release**, and then press Enter.
29. Type **ipconfig /renew**, and then press Enter.
30. Type **ipconfig**, and then press Enter. What is your IP address? Answers may vary.
31. On LON-DC1, in the Services console, start the DHCP server service.

Results: After completing this exercise, you will have configured a superscope, configured DHCP Name Protection, and configured and verified DHCP failover.

Exercise 2: Configuring Advanced DNS Settings

▶ Task 1: Configure DNSSEC

1. On LON-DC1, in Server Manager, click **Tools**, and then in the drop-down list, click **DNS**.
2. Expand **LON-DC1**, expand **Forward Lookup Zones**, click **Adatum.com**, and then right-click **Adatum.com**.
3. On the menu, click **DNSSEC>Sign the Zone**.
4. In the Zone Signing Wizard, click **Next**.
5. On the **Signing options** page, click **Customize zone signing parameters**, and then click **Next**.
6. On the **Key Master** page, ensure that the Domain Name System (DNS) server **LON-DC1** is selected as the Key Master, and then click **Next**.
7. On the **Key Signing Key (KSK)** page, click **Next**.
8. On the **Key Signing Key (KSK)** page, click **Add**.
9. On the **New Key Signing Key (KSK)** page, click **OK**.
10. On the **Key Signing Key (KSK)** page, click **Next**.
11. On the **Zone Signing Key (ZSK)** page, click **Next**.
12. On the **Zone Signing Key (ZSK)** page, click **Add**.
13. On the **New Zone Signing Key (ZSK)** page, click **OK**.
14. On the **Zone Signing Key (ZSK)** page, click **Next**.
15. On the **Next Secure (NSEC)** page, click **Next**.
16. On the **Trust Anchors (TAs)** page, check the **Enable the distribution of trust anchors for this zone** check box, and then click **Next**.
17. On the **Signing and Polling Parameters** page, click **Next**.
18. On the **DNS Security Extensions (DNSSEC)** page, click **Next**, and then click **Finish**.
19. In the DNS console, expand **Trust Points**, expand **com**, and then click **Adatum**. Ensure that the DNSKEY resource records display, and that their status is valid.
20. Minimize the DNS Manager.
21. In Server Manager, click **Tools**, and then on the drop-down list, click **Group Policy Management**.
22. Expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**, right-click **Default Domain Policy**, and then click **Edit**.
23. In the Group Policy Management Editor, under Computer Configuration, expand **Policies**, expand **Windows Settings**, and then click **Name Resolution Policy**.
24. In the right pane, under Create Rules, in the **Suffix** box, type **Adatum.com** to apply the rule to the suffix of the namespace.
25. Select both the **Enable DNSSEC in this rule** check box and the **Require DNS clients to check that the name and address data has been validated by the DNS server** check box, and then click **Create**.
26. Close the Group Policy Management Editor and Group Policy Management Console.

▶ Task 2: Configure the DNS socket pool

1. On LON-DC1, on the taskbar, click the Windows PowerShell icon.
2. In the Windows PowerShell window, type the following command, and then press Enter:

```
Get-DNSServer
```

   This command displays the current size of the DNS socket pool (on the fourth line in the ServerSetting section). Note that the current size is 2,500.

3. Type the following command, and then press Enter to change the socket pool size to 3,000.

```
dnscmd /config /socketpoolsize 3000
```

4. Type the following command, and then press Enter to stop the DNS server:

```
net stop dns
```

5. Type the following command, and then press Enter to start the DNS server.

```
net start dns
```

6. Type the following command, and then press Enter to confirm the new socket pool size.

```
Get-DnsServer
```

▶ Task 3: Configure DNS cache locking

1. In the Windows PowerShell window, type the following command, and then press Enter.

```
Get-Dnsserver
```

   This displays the current percentage value of the DNS cache lock. Note that the current value is 100 percent. The value displays in the ServerCache section.

2. Type the following command, and then press Enter:

```
Set-DnsServerCache –LockingPercent 75
```

   This changes the cache lock value to 75 percent.

3. Type the following command, and then press Enter to stop the DNS server.

```
net stop dns
```

4. Type the following command, and then press Enter to start the DNS server:

```
net start dns
```

5. Type the following command, and then press Enter:

```
Get-DnsServer
```

   This command displays the current percentage value of the DNS cache lock. Note that the new value is 75 percent.

6. Leave the command prompt window open for the next task.

▶ Task 4: Configure a GlobalNames zone

1. Create an Active Directory -integrated forward lookup zone named **Contoso.com** by running the following cmdlet in Windows PowerShell:

```
Add-DnsServerPrimaryZone –Name Contoso.com –ReplicationScope Forest
```

2. In the Windows PowerShell window, type the following command, and then press Enter to enable support for GlobalName zones:

```
Set-DnsServerGlobalNameZone –AlwaysQueryServer $true
```

3. Create an Active Directory- integrated forward lookup zone named **GlobalNames** by running the following command:

```
Add-DnsServerPrimaryZone –Name GlobalNames –ReplicationScope Forest
```

4. Minimize the command prompt window.
5. From the taskbar, restore the DNS console.
6. In the DNS console, click **Action**, and then click **Refresh**.
7. In the DNS console, refresh and then expand **Forward Lookup Zones**, click the **Contoso.com** zone, right-click **Contoso.com**, and then click **New Host (A or AAAA)**.
8. In the **New Host** dialog box, in the **Name** box, type **App1**.

Note: The **Name** box uses the parent domain name if it is left blank.

9. In the **IP address** box, type **192.168.1.200**, and then click **Add Host**.
10. Click **OK**, and then click **Done**.
11. Select and then right-click the **GlobalNames** zone, and then click **New Alias (CNAME)**.
12. In the **New Resource Record** dialog box, in the **Alias name** box, type **App1**.
13. In the **Fully qualified domain name (FQDN) for target host** box, type **App1.Contoso.com**, and then click **OK**.
14. Close DNS Manager, and close the command prompt.

Results: After completing this exercise, you will have configured DNSSEC, the DNS socket pool, DNS cache locking, and the GlobalName zone.

Exercise 3: Configuring IPAM

▶ Task 1: Install the IPAM feature

1. On LON-SVR2, in the Server Manager Dashboard, click **Add roles and features**.
2. In the Add Roles and Features Wizard, click **Next**.
3. On the **Select installation type** page, click **Next**.
4. On the **Select destination server** page, click **Next**.
5. On the **Select server roles** page, click **Next**.
6. On the **Select features** page, select the **IP Address Management (IPAM) Server** check box.
7. In the **Add features that are required for IP Address Management (IPAM) Server** popup, click **Add Features**, and then click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. Close the Add Roles and Features Wizard when complete.

▶ Task 2: Configure IPAM–related GPOs

1. On LON-SVR2, in the Server Manager navigation pane, click **IPAM**.
2. In the IPAM Overview pane, click **Connect to IPAM server**, click **LON-SVR2.Adatum.com**, and then click **OK**.
3. Click **Provision the IPAM server**.
4. In the Provision IPAM Wizard, on the **Before you begin** page, click **Next**.
5. On the **Configure database** page, click **Next**.
6. On the **Select provisioning method** page, ensure that the **Group Policy Based** method is selected. In the **GPO name prefix** box, type **IPAM**, and then click **Next**.
7. On the **Confirm the Settings** page, click **Apply**. Provisioning will take a few minutes to complete.
8. When provisioning completes, click **Close**.

▶ Task 3: Configure IP management server discovery

1. On the IPAM Overview pane, click **Configure server discovery**.
2. In the **Configure Server Discovery** settings dialog box, click **Add**, and then click **OK**.
3. In the IPAM Overview pane, click **Start server discovery**. Discovery may take 5 to 10 minutes to run. The yellow bar will indicate when discovery is complete.

▶ Task 4: Configure managed servers

1. In the IPAM Overview pane, click **Select or add servers to manage and verify IPAM access**. Notice that the IPAM Access Status is blocked.
2. Scroll down to the Details view, and note the status report, which is that the IPAM server has not yet been granted permission to manage LON-DC1 via Group Policy.
3. On the taskbar, right-click **Windows PowerShell**, and then click **Run as Administrator**.
4. At the Windows PowerShell prompt, type the following command, and then press Enter:

```
Invoke-IpamGpoProvisioning –Domain Adatum.com –GpoPrefixName IPAM –IpamServerFqdn
LON-SVR2.adatum.com –DelegatedGpoUser Administrator
```

5. When you are prompted to confirm the action, type **Y**, and then press Enter. The command will take a few minutes to complete.
6. Close Windows PowerShell.
7. In Server Manager, in the SERVER INVENTORY>IPv4 pane, right-click **LON-DC1**, and then click **Edit Server**.
8. In the **Add or Edit Server** dialog box, set the **Manageability status** to **Managed**, and then click **OK**.
9. Switch to LON-DC1.
10. From the start screen, start a command prompt.
11. In the command prompt window, type **Gpupdate /force**, and then press Enter.
12. Wait for the gpupdate to finish.
13. Close the Windows command prompt.
14. Switch to LON-SVR2.
15. Repeat Steps 7 through 13 for LON-SVR1.
16. In Server Manager, in the IPAM console, right-click **LON-DC1**, and then click **Refresh Server Access Status**.
17. In Server Manager, in the IPAM console, right-click **LON-SVR1**, and then click **Refresh Server Access Status**.
18. After the refresh completes, click the Server Manager console refresh button. It may take up to 10 minutes for the status to change. If necessary, repeat both refresh tasks as needed until a green check mark displays next to LON-DC1 and the IPAM Access Status shows **Unblocked**.
19. In the Server Inventory Page, right click LON-DC1 and then click **Retrieve ALL Server Data**. This action will take a few minutes to complete.
20. In the IPAM Overview pane, right click LON-SVR1 and then click **Retrieve ALL Server Data**. This action will take a few minutes to complete.

▶ Task 5: Configure and verify a new DHCP scope with IPAM

1. On LON-SVR2, in the IPAM navigation pane, under MONITOR AND MANAGE, click **DNS and DHCP Servers**.
2. In the details pane, right-click the instance of **LON-DC1.Adatum.com** that contains the DHCP server role, and then click **Create DHCP Scope**.
3. In the **Create DHCP Scope** dialog box, in the **Scope Name** box, type **TestScope**.
4. In the **Start IP address** box, type **10.0.0.50**.
5. In the **End IP address** box, type **10.0.0.100**.
6. Ensure the subnet mask is **255.0.0.0**.
7. In the Create scope pane, click **Options**.
8. On the DHCP Scope Options page, click **New**.
9. In the Configure options pane, click the **Option** drop-down arrow, and then select **003 Router**.
10. Under **Values**, in the **IP Address** box, type **10.0.0.1**, click **Add Configuration**, and then click **OK**.
11. In the navigation pane, click **DHCP Scopes**.
12. Right-click **Test Scope**, and then click **Configure DHCP Failover**.

13. In the **Configure DHCP Failover Relationship** dialog box, for the **Partner server** field, click the **Select** drop-down arrow, and then click **lon-svr1.adatum.com**.
14. In the **Relationship Name** field, type **TestFailover**.
15. In the **Enable Message Authentication Secret** field, type **Pa$$w0rd**.
16. In the **Maximum Client Lead Time** field, set the hours to **zero**, and then set the minutes to **15**.
17. Ensure the **Mode** field is set to **Load balance**.
18. Ensure that the **Load Balance Percentage** is set to **50%**.
19. Select the **Enable state switchover** check box. Leave the default value of 60 minutes.
20. Click **OK**.
21. On LON-DC1, in the Server Manager toolbar, click **Tools**, and then click **DHCP**.
22. In the DHCP console, expand **lon-dc1.adatum.com**, expand **IPv4**, and confirm that TestScope exists.

▶ Task 6: Configure IP address blocks, record IP addresses, and create DHCP reservations and DNS records

1. On LON-SVR2, in the Server Manager, in the IPAM console tree, click **IP Address Blocks**.
2. In the right pane, click the **Tasks** drop-down arrow, and then click **Add IP Address Block**.
3. In the **Add or Edit IPv4 Address Block** dialog box, provide the following values, and then click **OK**:
   - Network ID: **172.16.0.0**
   - Prefix length: **16**
   - Description: **Head Office**
4. In the IPAM console tree, click **IP Address Inventory**.
5. In the right pane, click the **Tasks** drop-down arrow, and then click **Add IP Address**.
6. In the **Add IP Address** dialog box, under **Basic Configurations**, provide the following values, and then click **OK**:
   - IP address: **172.16.0.1**
   - MAC address: **112233445566**
   - Device type: **Routers**
   - Description: **Head Office Router**
7. Click the **Tasks** drop-down arrow, and then click **Add IP Address**.
8. In the **Add IP Address** dialog box, under **Basic Configuration**, provide the following values:
   - IP address: 172.16.0.10
   - MAC address: 223344556677
   - Device type: Host
9. In the Add IPv4 Address pane, click **DHCP Reservation**, and then enter the following values:
   - Client ID: Check the Associate MAC to Client ID checkbox
   - Reservation server name: LON-DC1.Adatum.com
   - Reservation name: Webserver
   - Reservation type: Both

10. In the Add IPv4 Address pane, click **DNS Record**, enter the following values, and then click **OK**:
    - Device name: **Webserver**
    - Forward lookup zone: **Adatum.com**
    - Forward lookup primary server: **LON-DC1.adatum.com**
    - Check the **Automatically create DNS records for this IP address** check box.
11. On LON-DC1, open the DHCP console, expand **IPv4**, expand **Scope (172.16.0.0) Adatum**, and then click **Reservations**. Ensure that the Webserver reservation for 172.16.0.10 displays.
12. Open the DNS console, expand **Forward Lookup Zones**, and then click **Adatum.com**. Ensure that a host record displays for Webserver.

▶ Task 7: To prepare for the next module

1. On the host computer, start the Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20412C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20412C-LON-SVR1**, **20412C-LON-SVR2**, and **20412C-LON-CL1**.

Results: After completing this exercise, you will have installed IPAM and configured IPAM with IPAM-related GPOs, IP management server discovery, managed servers, a new DHCP scope, IP address blocks, IP addresses, DHCP reservations, and DNS records.

Module 2: Implementing Advanced File Services

Lab A: Implementing Advanced File Services

Exercise 1: Configuring iSCSI Storage

▶ Task 1: Install the iSCSI target feature

1. Sign in on LON-DC1 with the username **Adatum\Administrator** and the password **Pa$$w0rd**.
2. In the Server Manager, click **Add roles and features**.
3. In the Add Roles and Features Wizard, on the **Before You Begin** page, click **Next**.
4. On the **Select installation type** page, click **Next**.
5. On the **Select destination server** page, ensure that **Select server from the server pool** is selected, and then click **Next**.
6. On the **Select server roles** page, expand **File And Storage Services (2 of 12 Installed)**, expand **File and iSCSI Services (1 of 11 Installed)**, select the **iSCSI Target Server** check box, and then click **Next**.
7. On the **Select features** page, click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. When the installation completes, click **Close**.

▶ Task 2: Configure the iSCSI targets

1. On LON-DC1, in the Server Manager, in the navigation pane, click **File and Storage Services**.
2. In the File and Storage Services pane, click **iSCSI**.
3. In the iSCSI VIRTUAL DISKS pane, click **TASKS**, and then in the **TASKS** drop-down list box, click **New iSCSI Virtual Disk**.
4. In the New iSCSI Virtual Disk Wizard, on the **Select iSCSI virtual disk location** page, under **Storage location**, click drive **C**, and then click **Next**.
5. On the **Specify iSCSI virtual disk name** page, in the **Name** text box, type **iSCSIDisk1**, and then click **Next**.
6. On the **Specify iSCSI virtual disk size** page, in the **Size** text box, type **5**, in the drop-down list box, ensure that **GB** is selected, and then click **Next**.
7. On the **Assign iSCSI target** page, click **New iSCSI target**, and then click **Next**.
8. On the **Specify target name** page, in the **Name** box, type **LON-DC1**, and then click **Next**.
9. On the **Specify access servers** page, click **Add**.
10. In the **Select a method to identify the initiator** dialog box, click **Enter a value for the selected type**, in the **Type** drop-down list box, click **IP Address**, in the **Value** text box, type **172.16.0.22**, and then click **OK**.
11. On the **Specify access servers** page, click **Add**.
12. In the **Select a method to identify the initiator** dialog box, click **Enter a value for the selected type**, in the **Type** drop-down list box, click **IP Address**, in the **Value** text box, type **131.107.0.2**, and then click **OK**.
13. On the **Specify access servers** page, click **Next**.

14. On the **Enable Authentication** page, click **Next**.
15. On the **Confirm selections** page, click **Create**.
16. On the **View results** page, wait until creation completes, and then click **Close**.
17. In the iSCSI VIRTUAL DISKS pane, click **TASKS**, and then in the **TASKS** drop-down list box, click **New iSCSI Virtual Disk**.
18. In the New iSCSI Virtual Disk Wizard, on the **Select iSCSI virtual disk location** page, under **Storage location**, click drive **C**, and then click **Next**.
19. On the **Specify iSCSI virtual disk name** page, in the **Name** box, type **iSCSIDisk2**, and then click **Next**.
20. On the **Specify iSCSI virtual disk size** page, in the **Size** box, type **5**, in the drop-down list box, ensure that **GB** is selected, and then click **Next**.
21. On the **Assign iSCSI target** page, click **lon-dc1**, and then click **Next**.
22. On the **Confirm selection** page, click **Create**.
23. On the **View results** page, wait until creation completes, and then click **Close**.
24. In the iSCSI VIRTUAL DISKS pane, click **TASKS**, and then in the **TASKS** drop-down list box, click **New iSCSI Virtual Disk**.
25. In the New iSCSI Virtual Disk Wizard, on the **Select iSCSI virtual disk location** page, under **Storage**, click drive **C**, and then click **Next**.
26. On the **Specify iSCSI virtual disk name** page, in the **Name** text box, type **iSCSIDisk3**, and then click **Next**.
27. On the **Specify iSCSI virtual disk size** page, in the **Size** text box, type **5**, in the drop-down list box, ensure that **GB** is selected, and then click **Next**.
28. On the **Assign iSCSI target** page, click **lon-dc1**, and then click **Next**.
29. On the **Confirm selection** page, click **Create**.
30. On the **View results** page, wait until creation completes, and then click **Close**.

▶ Task 3: Configure MPIO

1. Sign in on LON-SVR2 with the username **Adatum\Administrator** and the password **Pa$$w0rd**.
2. In the Server Manager, on the menu bar, click **Tools**, and then in the **Tools** drop-down list, click **Routing and Remote Access**.
3. Right-click **LON-SVR2 (local)**, and then click **Disable Routing and Remote Access**. Click **Yes**, after it has stopped, close the Routing and Remote Access console.

Note: Normally, you do not disable Routing and Remote Access (RRAS) before configuring Multipath input/output (MPIO). You do it here because of lab requirements.

4. In the Server Manager, click **Add roles and features**.
5. In the Add Roles and Features Wizard, on the **Before You Begin** page, click **Next**.
6. On the **Select installation type** page, click **Next**.
7. On the **Select destination server** page, ensure that **Select server from the server pool** is selected, and then click **Next**.

8. On the **Select server roles** page, click **Next**.
9. On the **Select features** page, click **Multipath I/O**, and then click **Next**.
10. On the **Confirm installation selections** page, click **Install**.
11. When installation is complete, click **Close**.
12. In Server Manager, on the menu bar, click **Tools**, and then in the **Tools** drop-down list box, select **iSCSI Initiator**.
13. In the **Microsoft iSCSI** dialog box, click **Yes**.
14. In the **iSCSI Initiator Properties** dialog box, on the **Targets** tab, in the **Target** box, type **LON-DC1**, and then click **Quick Connect**.
15. In the **Quick Connect** box, click **Done**.
16. Click **OK** to close the **iSCSI Initiator Properties** dialog box.
17. In Server Manager, on the menu bar, click **Tools**, and then in the **Tools** drop-down list box, click **MPIO**.
18. In **MPIO Properties** dialog box, click the **Discover Multi-Paths** tab.
19. On the **Discover Multi-Paths** tab, select the **Add support for iSCSI devices** check box, and then click **Add**. When you are prompted to restart the computer, click **Yes**.
20. After the computer restarts, sign in on LON-SVR2 with the username **Adatum\Administrator** and the password **Pa$$w0rd**.
21. In the Server Manager, on the menu bar, click **Tools**, and then in the **Tools** drop-down list box, click **MPIO**.
22. In the **MPIO Properties** dialog box, on the **MPIO Devices** tab, notice that additional **Device Hardware ID MSFT2005iSCSIBusType_0x9** is added to the list.
23. Click **OK** to close the **MPIO Properties** dialog box.

▶ Task 4: Connect to and configure the iSCSI targets

1. On LON-SVR2, in the Server Manager, on the menu bar, click **Tools**, and then in the **Tools** drop-down list box, click **iSCSI Initiator**.
2. In the **iSCSI Initiator Properties** dialog box, on the **Targets** tab, click **Disconnect**.
3. In the **Disconnect From All Sessions** dialog box, click **Yes**.
4. In the **iSCSI Initiator Properties** dialog box, on the **Targets** tab, click **Connect**.
5. In the **Connect to Target** window, click **Enable multi-path**, verify that the **Add this connection to the list of Favorite Targets** check box is selected, and then click **Advanced**.
6. In the **Advanced Settings** dialog box, on the **General** tab, change the **Local Adapter** from **Default** to **Microsoft iSCSI Initiator**. In the **Initiator IP** drop-down list box, click **172.16.0.22**, and in the **Target Portal IP** drop-down list box, click **172.16.0.10 / 3260**.
7. In the **Advanced Settings** dialog box, click **OK**.
8. In the **Connect to Target** window, click **OK**.
9. In the **iSCSI Initiator Properties** dialog box, on the **Targets** tab, click **Connect**.
10. In **Connect to Target** window, click **Enable multi-path**, verify that the **Add this connection to the list of Favorite Targets** check box is selected, and then click **Advanced**.

11. In the **Advanced Settings** dialog box, on the **General** tab, change the **Local Adapter** from **Default** to **Microsoft iSCSI Initiator**. In the Initiator **IP** drop-down list box, select **131.107.0.2**, and in the **Target Portal IP** drop-down list box, select **131.107.0.1 / 3260**.
12. In the **Advanced Settings** dialog box, click **OK**.
13. In the **Connect to Target** window, click **OK**.
14. In the **iSCSI Initiator Properties** dialog box, click the **Volumes and Devices** tab.
15. In the **iSCSI Initiator Properties** dialog box, on the **Volumes and Devices** tab, click **Auto Configure**.
16. In the **iSCSI Initiator Properties** dialog box, click the **Targets** tab.
17. In the **Targets** list, select **iqn.1991-05.com.microsoft:lon-dc1-lon-dc1-target**, and then click **Devices**.
18. In the **Devices** dialog box, click **MPIO**.
19. Verify that in **Load balance policy**, **Round Robin** is selected. Under **This device has the following paths**, notice that two paths are listed. Select the first path, and then click **Details**.
20. Note the IP address of the Source and Target portals, and then click **OK**.
21. Select the second path, and then click **Details**.
22. Verify that the Source IP address is of the second network adapter, and then click **OK**.
23. Click **OK** to close the **Device Details** dialog box.
24. Click **OK** to close the **Devices** dialog box.
25. Click **OK** to close the **iSCSI Initiator Properties** dialog box.

Results: After completing this exercise, you will have configured and connected to iSCSI targets.

Exercise 2: Configuring the File Classification Infrastructure

▶ Task 1: Create a classification property for corporate documentation

1. On LON-SVR1, in the Server Manager, in the upper-right corner, click **Tools**, and then click **File Server Resource Manager**.
2. In the **File Server Resource Manager** window, expand **Classification Management**, select and then right-click **Classification Properties**, and then click **Create Local Property**.
3. In the **Create Local Classification Property** window, in the **Name** text box, type **Corporate Documentation**, in the **Property Type** drop-down list box, ensure that **Yes/No** is selected, and then click **OK**.
4. Leave the File Server Resource Manager console open.

▶ Task 2: Create a classification rule for corporate documentation

1. In the File Server Resource Manager, click **Classification Rules**, and then in the Actions pane, click **Create Classification Rule**.
2. In the **Create Classification Rule** window, on the **General** tab, in the **Rule name** text box, type **Corporate Documents Rule**, and then ensure that the **Enabled** check box is selected.
3. Click the **Scope** tab, and then click **Add**.
4. In the **Browse For Folder** window, expand **Allfiles (E:)**, expand **Labfiles**, click **Corporate Documentation**, and then click **OK**.
5. In the **Create Classification Rule** window, on the **Classification** tab, in the **Classification method** drop-down list box, click **Folder Classifier**. In the **Property-Choose a property to assign to files** drop-down list box, click **Corporate Documentation**, and then in the **Property-Specify a value** drop-down list box, click **Yes**.
6. Click the **Evaluation type** tab, click **Re-evaluate existing property values**, ensure that the **Aggregate the values** radio button is selected, and then click **OK**.
7. In the File Server Resource Manager, in the Actions pane, click **Run classification with all rules now**.
8. In the **Run classification** window, select the **Wait for classification to complete** radio button, and then click **OK**.
9. Review the Automatic classification report that displays in Internet Explorer, and ensure that the report lists the same number of classified files as in the Corporate Documentation folder.
10. Close Internet Explorer, but leave the File Server Resource Manager console open.

▶ Task 3: Create a classification rule that applies to a shared folder

1. In the File Server Resource Manager console, expand **Classification Management**, right-click **Classification Properties**, and then click **Create Local Property**.
2. In the **Create Local Classification Property** window, in the **Name** text box, type **Expiration Date**. In the **Property Type** drop-down list box, ensure that **Date-time** is selected, and then click **OK**.
3. In the File Server Resource Manager, expand **Classification Management**, click **Classification Rules**, and then in the Actions pane, click **Create Classification Rule**.
4. In the Create Classification Rule window, on the **General** tab, in the **Rule name** text box, type **Expiration Rule**, and ensure that the **Enabled** check box is selected.
5. Click the **Scope** tab, and then click **Add**.

6. In the **Browse For Folder** window, expand **Allfiles (E:)**, expand **Labfiles**, click **Corporate Documentation**, and then click **OK**.
7. Click the **Classification** tab. In the **Classification method** drop-down list box, click **Folder Classifier**, and then in the **Property-Choose a property to assign to files** drop-down list box, click **Expiration Date**.
8. Click the **Evaluation type** tab. Click **Re-evaluate existing property values**, ensure that the **Aggregate the values** radio button is selected, and then click **OK**.
9. In the File Server Resource Manager console, in the Actions pane, click **Run classification with all rules now**.
10. In the Run classification window, select the **Wait for classification to complete** radio button, and then click **OK**.
11. Review the Automatic classification report that displays in Internet Explorer, and ensure that the report lists the same number of classified files as in the Corporate Documentation folder.
12. Close Internet Explorer, but leave the File Server Resource Manager console open.

▶ Task 4: Create a file management task to expire corporate documents

1. In File Server Resource Manager, select and then right-click **File Management Tasks**, and then click **Create File Management Task**.
2. In the **Create File Management Task** window, on the **General** tab, in the **Task name** text box, type **Expired Corporate Documents**, and then ensure that the **Enabled** check box is selected.
3. Click the **Scope** tab, and then click **Add**.
4. In the **Browse For Folder** window, click **E:\Labfiles\Corporate Documentation**, and then click **OK**.
5. In the **Create File Management Task** window, on the **Action** tab, in the **Type** drop-down list box, ensure that **File expiration** is selected, and then in the **Expiration directory** box, type **E:\Labfiles\Expired**.
6. Click the **Notification** tab, and then click **Add**.
7. In the **Add Notification** window, on the **Event Log** tab, select the **Send warning to event log** check box, and then click **OK**.
8. Click the **Condition** tab, select the **Days since file was last modified** check box, and then in the same row, replace the default value of 0 with **1**.

 **Note:** This value is for lab purposes only. In a real scenario, the value would be 365 days or more, depending on the organization's policy.

9. Click the **Schedule** tab, ensure that the **Weekly** radio button is selected, select the **Sunday** check box, and then click **OK**.
10. Leave the File Server Resource Manager console open.

▶ Task 5: Verify that corporate documents are expired

1. In the File Server Resource Manager, click **File Management Tasks**, right-click **Expired Corporate Documents**, and then click **Run File Management Task Now**.
2. In the **Run File Management Task** window, click **Wait for the task to complete**, and then click **OK**.
3. Review the File management task report that displays in Internet Explorer, and ensure that the report lists the same number of classified files as in the Corporate Documentation folder.

4. In Server Manager, click **Tools**, and then click **Event Viewer**.
5. In the Event Viewer console, expand **Windows Logs**, and then click **Application**.
6. Review events with numbers **908** and **909**. Notice that 908 – FSRM started a file management job, and that 909 – FSRM finished a file management job.
7. Close open Windows.

▶ Task 6: Prepare for the next lab

When you finish the lab, revert 20412C-LON-SVR1. To do this, complete the following steps.

1. On the host computer, start Hyper-V Manager.
2. On the **Virtual Machines** list, right-click **20412C-LON-SVR1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.

   Keep all other virtual machines running for the next lab.

Results: After completing this exercise, you will have configured a File Classification Infrastructure so that the latest version of the documentation is always available to users.

Lab B: Implementing BranchCache

Exercise 1: Configuring the Main Office Servers for BranchCache

▶ Task 1: Configure LON-DC1 to use Windows BranchCache

1. On LON-DC1, on the taskbar, click the Server Manager icon.
2. In the Server Manager, click **Add roles and features**.
3. In the Add Roles and Features Wizard, on the **Before You Begin** page, click **Next**.
4. On the **Select installation type** page, click **Next**.
5. On the **Select destination server** page, ensure that **Select server from the server pool** is selected, and then click **Next**.
6. On the **Select server roles** page, expand **File And Storage Services (2 of 12 installed)**, expand **File and iSCSI Services (1 of 11 installed)**, select the **BranchCache for Network Files** check box, and then click **Next**.
7. On the **Select features** page, click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. Click **Close**, and then close Server Manager.
10. Right-click the **Start** charm and click **Run**, in the **Run** text box, type **gpedit.msc**, and then press Enter.
11. In the Local Group Policy Editor console, in the navigation pane, under **Computer Configuration**, expand **Administrative Templates**, expand **Network**, and then click **Lanman Server**.
12. On the Lanman Server result pane, in the **Setting** list, right-click **Hash Publication for BranchCache**, and then click **Edit**.
13. In the **Hash Publication for BranchCache** dialog box, click **Enabled**, in the **Hash publication actions** list, select the **Allow hash publication only for shared folders on which BranchCache is enabled**, and then click **OK**.

▶ Task 2: Simulate a slow link to the branch office

1. In the Local Group Policy Editor console, in the navigation pane, under **Computer Configuration**, expand **Windows Settings**, right-click **Policy-based QoS**, and then click **Create new policy**.
2. In the Policy-based QoS Wizard, on the **Create a QoS policy** page, in the **Policy name** text box, type **Limit to 100 Kbps**, select the **Specify Outbound Throttle Rate** check box. In the **Specify Outbound Throttle Rate** text box, type **100**, and then click **Next**.
3. On the **This QoS policy applies to** page, click **Next**.
4. On the **Specify the source and destination IP addresses** page, click **Next**.
5. On the **Specify the protocol and port numbers** page, click **Finish**.
6. Close the Local Group Policy Editor console.

▶ Task 3: Enable a file share for BranchCache

1. On LON-DC1, on the taskbar, click the **File Explorer** icon.
2. In the File Explorer window, browse to Local Disk **(C:)**.
3. In the Local Disk **(C:)** window, on the menu, click the **Home** tab, and then click **New Folder**.

4. Type **Share**, and then press Enter.
5. Right-click **Share**, and then click **Properties**.
6. In the **Share Properties** dialog box, on the **Sharing** tab, click **Advanced Sharing**.
7. Select the **Share this folder** check box, and then click **Caching**.
8. In the **Offline Settings** dialog box, select the **Enable BranchCache** check box, and then click **OK**.
9. In the **Advanced Sharing** dialog box, click **OK**.
10. In the **Share Properties** dialog box, click **Close**.
11. Right-click the **Start** charm and click **Search**, and in the **Search** text box, type **cmd**, and then press Enter.
12. At the command prompt, type the following command, and then press Enter:

```
Copy C:\windows\system32\mspaint.exe c:\share
```

13. Close the command prompt.
14. Close File Explorer.

▶ Task 4: Configure client firewall rules for BranchCache

1. On LON-DC1, in the Server Manager, click **Tools**, and then click **Group Policy Management**.
2. In the Group Policy Management console, expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**, right-click **Default Domain Policy** and then click **Edit**.
3. In the Group Policy Management Editor, in the navigation pane, under **Computer Configuration**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, and then expand **Windows Firewall with Advanced Security**.
4. In Windows Firewall with Advanced Security, in the navigation pane, expand **Windows Firewall with Advanced Security**, and then click **Inbound Rules**.
5. In the Group Policy Management Editor, right-click Inbound Rules, click **New Rule**.
6. In the New Inbound Rule Wizard, on the **Rule Type** page, click **Predefined**, click **BranchCache – Content Retrieval (Uses HTTP)**, and then click **Next**.
7. On the **Predefined Rules** page, click **Next**.
8. On the **Action** page, click **Finish** to create the firewall inbound rule.
9. In the Group Policy Management Editor, in the navigation pane, click **Inbound Rules**, and then in the Group Policy Management Editor, on the **Action** menu, click **New Rule**.
10. On the **Rule Type** page, click **Predefined**, click **BranchCache – Peer Discovery (Uses WSD)**, and then click **Next**.
11. On the **Predefined Rules** page, click **Next**.
12. On the **Action** page, click **Finish**.
13. Close the Group Policy Management Editor and Group Policy Management Console.

14. Right-click the **Start** charm and click **Run**. Type **CMD** in the Run box and click Enter.
15. At the command prompt type **gpupdate /force** and press Enter.

Results: At the end of this exercise, you will have deployed BranchCache, configured a slow link, and enabled BranchCache on a file share.

Exercise 2: Configuring the Branch Office Servers for BranchCache

▶ Task 1: Install the BranchCache feature on LON-SVR2

1. On LON-SVR2, in Server Manager, click **Add roles and features**.
2. In the Add Roles and Features Wizard, on the **Before You Begin** page, click **Next**.
3. On the **Select installation type** page, click **Next**.
4. On the **Select destination server** page, ensure that **Select server from the server pool** is selected, and then click **Next**.
5. On the **Select server roles** page, expand **File And Storage Services (1 of 12 Installed)**, expand **File and iSCSI Services**, and then select the **BranchCache for Network Files** check box.
6. In the **Add Roles and Features Wizard** dialog box click **Add Features**.
7. On the **Select server roles** page, click **Next**.
8. On the **Select features** page, click **BranchCache**, and then click **Next**.
9. On the **Confirm installation selections** page, click **Install**, and then click **Close**.

▶ Task 2: Start the BranchCache host server

1. On LON-SVR2, on the taskbar, click the **Windows PowerShell** icon.
2. In the Windows PowerShell window, type the following cmdlet, and then press Enter:

```
Enable-BCHostedServer
-RegisterSCP
```

3. In the Windows PowerShell window, type the following cmdlet, and then press Enter:

```
Get-BCStatus
```

4. Ensure that BranchCache is enabled and running. Note in the DataCache section, the current active cache size is zero.
5. At the prompt type **gpupdate /force** and press Enter

Results: At the end of this exercise, you will have enabled the BranchCache server in the branch office.

Exercise 3: Configuring Client Computers for BranchCache

▶ Task 1: Configure client computers to use BranchCache in hosted cache mode

1. On LON-DC1, on the taskbar, click the **Server Manager** icon.
2. In Server Manager, on the menu bar, click **Tools**, and then in the **Tools** drop-down list box, select **Group Policy Management**.
3. In the Group Policy Management console, expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**, right-click **Adatum.com** and click **New Organizational Unit**.
4. In the New Organizational Unit dialog box type **Branch** in the Name field and click **OK**.
5. Right-click the **Branch** OU and click **Create a GPO in this domain, and link it here**.
6. In the New GPO dialog box type **BranchCache** and click **OK**.
7. Expand the Branch OU and right-click the **BranchCache** GPO and click **Edit**.
8. In the Group Policy Management Editor, in the navigation pane, under **Computer Configuration**, expand **Policies**, expand **Administrative Templates**, expand **Network**, and then click **BranchCache**.
9. In the BranchCache results pane, in the **Setting** list, right-click **Turn on BranchCache**, and then click **Edit**.
10. In the **Turn on BranchCache** dialog box, click **Enabled**, and then click **OK**.
11. In the BranchCache results pane, in the **Setting** list, right-click **Enable Automatic Hosted Cache Discovery by Service Connection Point**, and then click **Edit**.
12. In the **Enable Automatic Hosted Cache Discovery by Service Connection Point** dialog box, click **Enabled**, and then click **OK**.
13. In the BranchCache results pane, in the **Setting** list, right-click **Configure BranchCache for network files**, and then click **Edit**.
14. In the **Configure BranchCache for network files** dialog box, click **Enabled**, in the **Type the maximum round trip network latency (milliseconds) after which caching begins** text box, type **0**, and then click **OK**. This setting is required to simulate access from a branch office and is not typically required.
15. Close the Group Policy Management Editor.
16. Close the Group Policy Management Console.
17. In Server Manager, on the menu bar, click **Tools**, and then in the **Tools** drop-down list box, select **Active Directory Users and Computers**.
18. Expand **Adatum.com** and click the **Computers** container.
19. While pressing the Ctrl key, select both **LON-CL1** and **LON-CL2**. Right-click the selection and click **Move**.
20. Click the **Branch** OU and then click **OK**.
21. Close **Active Directory Users and Computers**.
22. Start **20412C-LON-CL1**, and sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
23. On the Start screen, in the lower-right corner of the screen, click **Search**, in the **Search** text box, type **cmd**, and then press Enter.

24. At the command prompt, type the following command, and then press Enter:

```
netsh branchcache show status all
```

25. Verify that the BranchCache Current Status is **Running**. If the status is **Stopped**, restart the client machines.
26. Start **20412C-LON-CL2**, and sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
27. On the Start screen, in the lower-right corner of the screen, click **Search**, in the **Search** text box, type **power**, and then press Enter.
28. In the **Windows PowerShell** window, type the following command, and then press Enter:

```
netsh branchcache show status all
```

29. Verify that BranchCache Current Status is **Running**. If the status is **Stopped**, restart the client.

Results: At the end of this exercise, you will have configured the client computers for BranchCache.

Exercise 4: Monitoring BranchCache

▶ Task 1: Configure Performance Monitor on LON-SVR1

1. Switch to LON-SVR2.
2. In Server Manager, on the menu bar, click **Tools**, and then from the **Tools** drop-down list box, click **Performance Monitor**.
3. In the Performance Monitor console, in the navigation pane, under **Monitoring Tools**, click **Performance Monitor**.
4. In the Performance Monitor results pane, click the **Delete** icon.
5. In the Performance Monitor results pane, click the **Add** (Ctrl+N) icon.
6. In the **Add Counters** dialog box, under **Select counters from computer**, click **BranchCache**, click **Add**, and then click **OK**.
7. On the **Change Graph type** button, select **Report**.

▶ Task 2: Configure performance statistics on LON-CL1

1. Switch to LON-CL1.
2. Point to the lower-right corner of the screen, click **Search**, in the **Search** text box, type **perfmon**, and then press Enter.
3. In the Performance Monitor console, in the navigation pane, under **Monitoring Tools**, click **Performance Monitor**.
4. In the Performance Monitor results pane, click the **Delete** icon.
5. In the Performance Monitor results pane, click the **Add** (Ctrl+N) icon.
6. In the **Add Counters** dialog box, under **Select counters from computer**, click **BranchCache**, click **Add**, and then click **OK**.
7. Change graph type to **Report**. Notice that the value of all performance statistics is zero.

▶ Task 3: Configure performance statistics on LON-CL2

1. Switch to LON-CL2.
2. Point to the lower-right corner of the screen, click **Search**, in the **Search** text box, type **perfmon**, and then press Enter.
3. In the Performance Monitor console, in the navigation pane, under **Monitoring Tools**, click **Performance Monitor**.
4. In the Performance Monitor results pane, click the **Delete** icon.
5. In the Performance Monitor results pane, click the **Add** (Ctrl+N) icon.
6. In the **Add Counters** dialog box, under **Select counters from computer**, click **BranchCache**, click **Add**, and then click **OK**.
7. Change graph type to **Report**. Notice that the value for all performance statistics is zero.

▶ Task 4: Test BranchCache in the hosted cache mode

1. Switch to LON-CL1.
2. On the taskbar, click the **File Explorer** icon.
3. In File Explorer address bar, type **\\LON-DC1.adatum.com\Share**, and then press Enter.

4. In the Share window, in the **Name** list, right-click **mspaint.exe**, and then click **Copy**.
5. In the Share window, click **Minimize**.
6. On the desktop, right-click anywhere, and then click **Paste**.
7. Read the performance statistics on LON-CL1. This file was retrieved from LON-DC1 (Retrieval: Bytes from Server). After the file was cached locally, it was passed up to the hosted cache. (Retrieval: Bytes Served)
8. Switch to LON-CL2.
9. On the taskbar, click the **File Explorer** icon.
10. In the File Explorer address bar, type **\\LON-DC1.adatum.com\Share**, and then press Enter.
11. In the Share window, in the **Name** list, right-click **mspaint.exe**, and then click **Copy**.
12. In the **Share** window, click **Minimize**.
13. On the **desktop**, right-click anywhere, and then click **Paste**.
14. Read the performance statistics on **LON-CL2**. This file was obtained from the hosted cache (Retrieval: Bytes from Cache).
15. Read the performance statistics on **LON-SVR2**. This server has offered cached data to clients (Hosted Cache: Client file segment offers made).
16. On LON-SVR2, on the taskbar, click the **Windows PowerShell** icon.
17. In the Windows PowerShell window, type the following cmdlet, and then press Enter:

```
Get-BCStatus
```

Note: In the DataCache section, the current active cache size is no longer zero, it is 6560896.

▶ Task 5: Prepare for the next module

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

1. On the host computer, start **Hyper-V Manager**.
2. On the **Virtual Machines** list, right-click **20412C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps two and three for **20412C-LON-SVR2**, **20412C-LON-CL1**, and **20412C-LON-CL2**.

Results: At the end of this exercise, you will have verified that BranchCache is working as expected.

Module 3: Implementing Dynamic Access Control

Lab: Implementing Secure Data Access

Exercise 1: Preparing for DAC Deployment

▶ Task 1: Preparing AD DS for DAC deployment

1. On LON-DC1, in Server Manager, click on **Tools**, and then click **Active Directory Domains and Trusts**.
2. In the **Active Directory Domains and Trusts** console, right-click **Adatum.com** and select **Raise Domain Functional Level.**
3. In the **Raise domain functional level** window, in the **Select an available domain functional level** window, select **Windows Server 2012** and click **Raise**.
4. Click **OK** twice.
5. Right-click **Active Directory Domains and Trusts [LON-DC1.Adatum.com]** and click **Raise Forest Functional Level…**
6. In the **Raise forest functional level** window, in the **Select an available forest functional level** window, select **Windows Server 2012** and click **Raise**.
7. Click **OK** twice.
8. Close **Active Directory Domains and Trusts** console.
9. On LON-DC1, in Server Manager, click **Tools**, and then click **Active Directory Users and Computers**.
10. In Active Directory Users and Computers, right-click **Adatum.com**, click **New**, and then click **Organizational Unit**.
11. In the **New Object – Organizational Unit** dialog box, in the **Name** field, type **DAC-Protected**, and then click **OK**.
12. Click the **Computers** container.
13. Right-click the **LON-SVR1** and then click **Move**.
14. In the Move window, click **DAC-Protected**, and then click **OK**.
15. Repeat steps 13 and 14 for **LON-CL1** computer
16. Close Active Directory Users and Computers.
17. On LON-DC1, in Server Manager, click **Tools**, and then click **Group Policy Management**.
18. Expand **Forest: Adatum.com**, expand **Domains**, and then expand **Adatum.com**.
19. Click the **Group Policy Objects** container.
20. In the results pane, right-click **Default Domain Controllers Policy**, and then click **Edit**.
21. In the Group Policy Management Editor, under **Computer Configuration**, expand **Policies**, expand **Administrative Templates**, expand **System**, and then click **KDC**.
22. In the details pane, double-click **KDC support for claims, compound authentication and Kerberos armoring**.
23. In the KDC support for claims, compound authentication and Kerberos armoring window, select **Enabled**, in the **Options** section, click the drop-down list box, select **Always provide claims**, and then click **OK**.

24. Close Group Policy Management Editor and the Group Policy Management Console.
25. On the taskbar, click the **Windows PowerShell** icon.
26. In the Windows PowerShell window, type **gpupdate /force**, and then press Enter. After Group Policy updates, close Windows PowerShell.
27. On LON-DC1, in Server Manager, click **Tools**, and then click **Active Directory Users and Computers**.
28. Expand **Adatum.com**, right-click **Users**, click **New**, and then click **Group**.
29. In the **Group name** field, type **ManagersWKS**, and then click **OK**.
30. Click the **DAC-Protected** container, right-click **LON-CL1**, and then click **Properties**.
31. Click the **Member Of** tab, and then click **Add**.
32. In Select Groups window, type **ManagersWKS**, click **Check Names**, and then click **OK** twice.
33. Click the **Managers** OU, right-click **Aidan Delaney**, and then click **Properties**.
34. In the **Aidan Delaney Properties** window, click the **Organization** tab. Ensure that the **Department** field is populated with the value **Managers**, and then click **Cancel**.
35. Click the **Research** OU, right-click **Allie Bellew**, and then click **Properties**.
36. In the **Allie Bellew Properties** window, click the **Organization** tab. Ensure that the **Department** field is populated with the value **Research**, and then click **Cancel**.
37. Close Active Directory Users and Computers.

▶ Task 2: Configuring user and device claims

1. On LON-DC1, click **Tools**, and then click **Active Directory Administrative Center**.
2. In the Active Directory Administrative Center, in the navigation pane, click **Dynamic Access Control**, and then double-click **Claim Types**.
3. In the Claim Types container, in the Tasks pane, click **New**, and then click **Claim Type**.
4. In the Create Claim Type window, in the Source Attribute section, select **department**.
5. In the **Display name** text box, type **Company Department**.
6. Select both **User** and **Computer** check boxes.
7. Scroll down to **Suggested Values** section and select **The following values are suggested:** option
8. Click **Add...**
9. In the **Add a suggested value** window, type **Managers** in both **Value** and **Display name** fields, and click **OK**.
10. Click **Add**...
11. In the **Add a suggested value** window type **Research** in both **Value** and **Display name** fields, and click **OK**.
12. Click **OK**.
13. In the Active Directory Administrative Center, in the Tasks pane, click **New**, and then select **Claim Type**.
14. In the Create Claim Type window, in the Source Attribute section, click **description**.
15. Clear the **User** check box, select the **Computer** check box, and then click **OK**.

▶ Task 3: Configuring resource properties and resource property lists

1. In the Active Directory Administrative Center, click **Dynamic Access Control**.
2. In the central pane, double-click **Resource Properties**.
3. In the **Resource properties** list, right-click **Department**, and then click **Enable**.
4. In the **Resource properties** list, right-click **Confidentiality**, and then click **Enable**.
5. In the Resource Property List, ensure that both the **Department** and **Confidentiality** properties are enabled.
6. Double-click **Department**, scroll down to the Suggested Values section, and then click **Add**.
7. In the Add a suggested value window, in both **Value** and **Display name** text boxes, type **Research**, and then click **OK** twice.
8. Click **Dynamic Access Control**, and then double-click **Resource Property Lists**.
9. In the central pane, double-click **Global Resource Property List**, ensure that both **Department** and **Confidentiality** display, and then click **Cancel**. If they do not display, click **Add**, add these two properties, and then click **OK**.
10. Close the Active Directory Administrative Center.

▶ Task 4: Implement file classifications

1. On LON-SVR1, in Server Manager, click **Tools**, and then click **File Server Resource Manager**.
2. In File Server Resource Manager, expand **Classification Management**.
3. Select and right-click **Classification Properties**, and then click **Refresh**.
4. Verify that **Confidentiality** and **Department** properties are listed.
5. Click **Classification Rules**, and in the Actions pane, click **Create Classification Rule**.
6. In the **Create Classification Rule** window, for the Rule name, type **Set Confidentiality**.
7. Click the **Scope** tab, and then click **Add**.
8. In the **Browse For Folder** dialog box, expand **Local Disk (C:)**, click the **Docs** folder, and then click **OK**.
9. Click the **Classification** tab. Make sure that following settings are set, and then click **Configure**:
   - Classification method: **Content Classifier**
   - Property: **Confidentiality**
   - Value: **High**
10. In the **Classification Parameters** dialog box, click the **Regular expression** drop-down list box, and then click **String**.
11. In the **Expression** field next to the word String, type **secret**, and then click **OK**.
12. Click the **Evaluation Type** tab, select **Re-evaluate existing property values**, click **Overwrite the existing value**, and then click **OK**.
13. In File Server Resource Manager, in the Actions pane, click **Run Classification with all rules now**.
14. Click **Wait for classification to complete**, and then click **OK**.
15. After the classification is complete, you will be presented with a report. Verify that two files were classified. You can confirm this in the Report Totals section.

16. Close the report.
17. On the taskbar, click the **File Explorer** icon.
18. In the File Explorer window, expand drive **C:**, and then expand the **Docs** folder.
19. In the Docs folder, right-click **Doc1.txt**, click **Properties**, and then click the **Classification** tab. Verify that Confidentiality is set to **High**.
20. Repeat step 19 on files Doc2.txt and Doc3.txt. Doc2.txt should have same Confidentiality as Doc1.txt, while Doc3.txt should have no value. This is because only Doc1.txt and Doc2.txt have the word "secret" in their content.

▶ Task 5: Assign property to the Research folder

1. On LON-SVR1, open File Explorer, and browse to Local Disk (C:).
2. Right-click the **Research** folder, and then click **Properties**.
3. Click **Classification** tab.
4. Click **Department**.
5. In the Value section, click **Research**. Click **Apply**.
6. Click **OK**.

Results: After completing this exercise, you will have prepared Active Directory Domain Services (AD DS) for Dynamic Access Control (DAC) deployment, configured claims for users and devices, and configured resource properties to classify files.

Exercise 2: Implementing DAC

▶ Task 1: Configure central access rules

1. On LON-DC1, in Server Manager, click **Tools**, and then click **Active Directory Administrative Center**.
2. In the Active Directory Administrative Center, in the navigation pane, click **Dynamic Access Control**, and then double-click **Central Access Rules**.
3. In the Tasks pane, click **New**, and then click **Central Access Rule**.
4. In the **Create Central Access Rule** dialog box, in the **Name** field, type **Department Match**.
5. In the Target Resources section, click **Edit**.
6. In the **Central Access Rule** dialog box, click **Add a condition**.
7. Set a condition as follows: **Resource-Department-Equals-Value-Research**, and then click **OK**.
8. In the Permissions section, click **Use following permissions as current permissions**.
9. In the Permissions section, click **Edit**.
10. Remove permission for Administrators.
11. In Advanced Security Settings for Permissions, click **Add**.
12. In Permission Entry for Permissions, click **Select a principal**.
13. In the Select User, Computer, Service Account, or Group window, type **Authenticated Users**, click **Check Names**, and then click **OK**.
14. In the Basic permissions section, select the **Modify**, **Read and Execute**, **Read** and **Write** check boxes.
15. Click **Add a condition** and then click the **Group** drop-down list box, and then click **Company Department**.
16. Click the **Value** drop-down list box, and then click **Resource**.
17. In the last drop-down list box, click **Department**. Click **OK** three times.

Note: You should have this expression as a result: **User-Company Department-Equals-Resource-Department**.

18. In the **Tasks** pane, click **New**, and then click **Central Access Rule**.
19. For the name of the rule, type **Access Confidential Docs**.
20. In the Target Resources section, click **Edit**.
21. In the Central Access Rule window, click **Add a condition**.
22. In the last drop-down list box, click **High**. Click **OK**.

Note: You should have this expression as a result: **Resource-Confidentiality-Equals-Value-High**.

23. In the Permissions section, click **Use following permissions as current permissions**.
24. In the Permissions section, click **Edit**.
25. Remove permission for Administrators.
26. In Advanced Security Settings for Permissions, click **Add**.

27. In Permission Entry for Permissions, click **Select a principal**.
28. In the Select User, Computer, Service Account, or Group window, type **Authenticated Users**, click **Check Names**, and then click **OK**.
29. In the Basic permissions section, select the **Modify**, **Read and Execute**, **Read**, and **Write** check boxes.
30. Click **Add a condition**. Set the first condition to: **User-Company Department-Equals-Value-Managers**, and then click **Add a condition**.
31. Set the second condition to: **Device-Group-Member of each-Value-ManagersWKS**. Click **OK** three times.

 **Note:** If you cannot find ManagersWKS in the last drop-down list box, click **Add items**. Then in the Select User, Computer, Service Account, or Group window, type **ManagersWKS**, click **Check Names**, and then click **OK**.

▶ Task 2: Configure central access policies

1. On LON-DC1, in the Active Directory Administrative Center, click **Dynamic Access Control**, and then double-click **Central Access Policies**.
2. In the Tasks pane, click **New**, and then click **Central Access Policy**.
3. In the **Name** field, type **Protect confidential docs**, and then click **Add**.
4. Click the **Access Confidential Docs** rule, click **>>**, and then click **OK** twice.
5. In the Tasks pane, click **New**, and then click **Central Access Policy**.
6. In the **Name** field, type **Department Match**, and then click **Add**.
7. Click the **Department Match** rule, click **>>**, and then click **OK** twice.
8. Close the Active Directory Administrative Center.

▶ Task 3: Apply central access policies to a file server

1. On LON-DC1, in Server Manager, click **Tools**, and then click **Group Policy Management**.
2. In the Group Policy Management Console, under **Domains**, expand **Adatum.com**, right-click **DAC-Protected**, and then click **Create a GPO in this domain, and link it here**.
3. Type **DAC Policy**, and then click **OK**.
4. Right-click **DAC Policy**, and then click **Edit**.
5. Expand **Computer Configuration**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, expand **File System**, right-click **Central Access Policy**, and then click **Manage Central Access Policies**.
6. Press and hold the Ctrl button and click both **Department Match** and **Protect confidential docs**, click **Add**, and then click **OK**.
7. Close the Group Policy Management Editor and the Group Policy Management Console.
8. On LON-SVR1, on the taskbar, click the **Windows PowerShell** icon.
9. At a Windows PowerShell command prompt, type **gpupdate /force**, and then press Enter. Wait until Group Policy is updated.
10. Close Windows PowerShell when you get the message that both Computer and User policies update completed successfully.

11. On the taskbar, click the **File Explorer** icon.
12. In File Explorer, browse to **Local Disk** (**C:)**, right-click the **Docs** folder, and then click **Properties**.
13. In the **Properties** dialog box, click the **Security** tab, and then click **Advanced**.
14. In the **Advanced Security Settings for Docs** window, click the **Central Policy** tab, and then click **Change**.
15. In the drop-down list box, select **Protect confidential docs**. Ensure that in the **Applies to** drop down box, value **This folder, subfolders and files** is selected, and then click **OK** twice.
16. Right-click the **Research** folder, and then click **Properties**.
17. In the **Research Properties** dialog box, click the **Security** tab, and then click **Advanced**.
18. In the Advanced Security Settings for Research window, click the **Central Policy** tab, and then click **Change**.
19. In the drop-down list box, click **Department Match**. Ensure that in the **Applies to** drop-down box, the value **This folder, subfolders and files** is selected, and then click **OK** twice.

Results: After completing this exercise, you will have implemented DAC.

Exercise 3: Validating and Remediating DAC

▶ Task 1: Access file resources as an approved user

1. Start LON-CL1 and LON-CL2 virtual machines.
2. Sign in on LON-CL1 as **Adatum\Allie** with the password **Pa$$w0rd**.
3. Click the **Desktop** tile, and then on the taskbar, click the **File Explorer** icon.
4. In the File Explorer address bar, type **\\LON-SVR1\Research**, and then press Enter.
5. Because Allie is a member of the Research team, verify that you can access this folder and open the documents inside.
6. Sign out of LON-CL1.
7. Sign in to LON-CL1 as **Adatum\Aidan** with the password **Pa$$w0rd**.
8. Click the **Desktop** tile, and then on the taskbar, click the **File Explorer** icon.
9. In the File Explorer address bar, type **\\LON-SVR1\Docs**.
10. Verify that you can access this folder and open all the files inside.
11. Sign out of LON-CL1.

▶ Task 2: Access file resources as an unapproved user

1. Sign in to LON-CL2 as **Adatum\Aidan** with the password **Pa$$w0rd**.
2. Click the **Desktop** tile, and then on the taskbar, click the **File Explorer** icon.
3. In the File Explorer address bar, type **\\LON-SVR1\Docs**. You should be unable to view Doc1.txt or Doc2.txt, because LON-CL2 is not permitted to view secret documents.
4. Sign out of LON-CL2.
5. Sign in to LON-CL2 as **Adatum\April** with the password **Pa$$w0rd**.
6. Click the **Desktop** tile, and then on the taskbar, click the **File Explorer** icon.
7. In the File Explorer address bar, type **\\LON-SVR1\Docs**, and then press Enter.
8. In the **Docs** folder, try to open Doc3.txt. You should be able to open that document. Close Notepad.
9. In the File Explorer address bar, type **\\LON-SVR1\Research**, and then press Enter. You should be unable to access the folder.
10. Sign out of LON-CL2.

▶ Task 3: Evaluate user access with DAC

1. On LON-SVR1, on the taskbar, click the **File Explorer** icon.
2. In the File Explorer window, navigate to **C:\Research**, right-click **Research**, and then click **Properties**.
3. In the **Properties** dialog box, click the **Security** tab, click **Advanced**, and then click **Effective Access**.
4. Click **select a user**, and in the Select User, Computer, Service Account, or Group window, type **April**, click **Check Names**, and then click **OK**.
5. Click **View effective access**, and then review the results. The user April should not have access to this folder.
6. Click **Include a user claim**, and then in the drop-down list box, click **Company department**.

7. In the **Value** drop-down box, select **Research**, and then click **View Effective access**. April should now have access.
8. Close all open windows.

▶ Task 4: Configure access-denied remediation

1. On LON-DC1, in Server Manager, click **Tools**, and then click **Group Policy Management**.
2. In the Group Policy Management Console, expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**, and then click **Group Policy objects**.
3. Right-click **DAC Policy**, and then click **Edit**.
4. Under **Computer Configuration**, expand **Policies**, expand **Administrative Templates**, expand **System**, and then click **Access-Denied Assistance**.
5. In the details pane, double-click **Customize Message for Access Denied errors**.
6. In the Customize Message for Access Denied errors window, click **Enabled**.
7. In the **Display the following message to users who are denied access** text box, type **You are denied access because of permission policy. Please request access.**
8. Select the **Enable users to request assistance** check box.
9. Review the other options, but do not make any changes, and then click **OK**.
10. In the details pane of the Group Policy Management Editor, double-click **Enable access-denied assistance on client for all file types**, click **Enabled**, and then click **OK**.
11. Close the Group Policy Management Editor and the Group Policy Management Console.
12. Switch to LON-SVR1, and on the taskbar, click the **Windows PowerShell** icon.
13. At the Windows PowerShell command prompt, type **gpupdate /force**, and then press Enter. Wait until Group Policy is updated.

▶ Task 5: Request access remediation

1. Sign in to LON-CL1 as **Adatum\April** with the password **Pa$$w0rd**.
2. Click the **Desktop** tile, and then on the taskbar, click the **File Explorer** icon.
3. In the File Explorer address bar, type **\\LON-SVR1\Research**, and then press Enter. You should be unable to access the folder.
4. Click **Request assistance**. Review the options for sending a message, and then click **Close**.
5. Sign out of LON-CL1.

Results: After completing this exercise, you will have validated DAC functionality.

Exercise 4: Implementing Work Folders

▶ Task 1: Install Work Folders functionality, configure SSL certificate and create WFSync group

1. On LON-SVR2, in **Server Manager**, click **Add roles and features**.
2. On the **Before you begin** page, click **Next**.
3. On the **Select installation type** page, ensure that **Role - based or feature - based installation** is selected, and then click **Next**.
4. On the **Select destination server** page, click **Next**.
5. On the **Select server roles** page, expand **File and Storage Services**, expand **File and iSCSI Services**, and then select **Work Folders**.
6. In the **Add features that are required for Work Folders** dialog box, note the features, and then click **Add Features**.
7. On the **Select server roles** page, click **Next**.
8. On the **Select features** page, click **Next**.
9. On the **Confirm installation selection** pages, click **Install**.
10. When the installation finishes, click **Close**.
11. In the **Server Manager** on LON-SVR2, click **Tools**, and then click **Internet Information Services (IIS) Manager**.
12. In the **Internet Information Services (IIS) Manager** console, click on **LON-SVR2,** click **No** when prompted**,** and then double click **Server Certificates** in the middle pane.
13. In the **Actions** pane, click **Create Domain Certificate.**
14. In the **Create Certificate** window, fill in the text fields as follows:
    a. Common name: **lon-svr2.adatum.com**
    b. Organization: **Adatum**
    c. Organizational unit: **IT**
    d. City/locality : **Seattle**
    e. State/province : **WA**
    f. Country/region: **US**
15. Click **Next.**
16. On the **Online Certification Authority** page, click **Select**.
17. In the **Select Certification Authority** window, select **AdatumCA** and click **OK**
18. In the **Friendly name** text box type **lon-svr2.adatum.com** and click **Finish**. (Note: if you receive an error, restart the LON-DC1 machine, and then repeat this step.)
19. In the **IIS** console, expand **Sites**, and then click on **Default Web Site**.
20. In the **Actions** pane, click **Bindings**
21. In the **Site Bindings** window, click **Add…**.
22. In the **Add Site Binding** window, under Type, select https. In the **SSL certificate** drop-down list, select **lon-svr2.adatum.com**.

23. Click **OK**, and then click **Close**.
24. Close **Internet Information Services (IIS) Manager**.
25. Switch to LON-DC1.
26. On LON-DC1, in Server Manager, click **Tools**, and then click **Active Directory Users and Computers**.
27. Expand **Adatum.com**, right-click **Users**, click **New**, and then click **Group**.
28. In the **Group name** field, type **WFSync**, and then click **OK**.
29. Click the **Managers** OU, right-click **Aidan Delaney**, and then click **Properties**.
30. Click the **Member Of** tab, and then click **Add**.
31. In Select Groups window, type **WFSync**, click **Check Names**, and then click **OK** twice.

▶ Task 2: Provision a share for Work Folders

1. On LON-SVR2, in Server Manager, in the navigation pane, click **File and Storage Services**.
2. Click **Shares**, and in the **SHARES** area, click **Tasks**, and then select **New Share...**.
3. In the **New Share Wizard**, on the **Select the profile for this share** page, ensure that **SMB Share – Quick** is selected, and then click **Next**.
4. On the **Select the server and path for this share** page, accept the defaults, and then click **Next**.
5. On the **Specify share name** page, in the **Share name** field, type **WF-Share**, and then click **Next**.
6. On the **Configure share settings** page, select **Enable access - based enumeration**, leave the other settings at their defaults, and then click **Next**.
7. On the **Specify permissions to control access** page, note the default settings, and then click **Next**.
8. On the **Confirm selections** page, click **Create**.
9. On the **View results** page, click **Close**.

▶ Task 3: Configuring and implementing Work Folders

1. On LON-SVR2, in Server Manager, expand **File and Storage Services**, and then click **Work Folders**.
2. In the **WORK FOLDERS** tile, click **Tasks**, and then click **New Sync Share...**.
3. In the **New Sync Share Wizard**, on the **Before you begin** page, click **Next**.
4. On the **Select the server and path** page, select **Select by file share**, ensure that the share you created in the previous task (WF-Share) is highlighted, and then click **Next**.
5. On the **Specify the structure for user folders** page, accept the default selection (user alias), and then click **Next**.
6. On the **Enter the sync share name** page, accept the default, and then click **Next**.
7. On the **Grant sync access to groups** page, note the default selection to disable inherited permissions and grant users exclusive access, and then click **Add**.
8. In the **Select User or Group** dialog box, in the **Enter the object names to select** field, type **WFsync**, click **Check Names**, and then click **OK**.
9. On the **Grant sync access to groups** page, click **Next**.
10. On the **Specify device policies** page, note the selections, accept the default selection, and then click **Next**.
11. On the **Confirm selections** page, click **Create**.

12. On the **View results** page, click **Close**.
13. Switch to LON-DC1.
14. Open **Server Manager**, click **Tools**, and then click **Group Policy Management**.
15. Expand **Forest: Adatum.com-Domains-Adatum.com**, and then click **Group Policy Objects**. Right-click **Group Policy Objects**, and then click **New**.
16. In the **New GPO** window, type **Work Folders GPO** in the **Name** field, and then click **OK**.
17. Right-click **Work Folders GPO**, and then click **Edit**.
18. In the Group Policy Management Editor, expand **User Configuration \Policies\Administrative Templates\Windows Components**, and then click **Work Folders**.
19. Double-click **Specify Work Folders settings** in the details pane, and in the **Specify Work Folders settings** dialog box, click **Enabled**.
20. In the **Work Folders URL** text box, type **https://lon-svr2.adatum.com**, and then select **Force automatic setup**.
21. Click **OK** to close the **Specify Work Folders settings** dialog box, and then close the Group Policy Management Editor.
22. In the Group Policy Management Console, right-click the **Adatum.com** domain object, and then select **Link an Existing GPO…**.
23. In the Select GPO window, select **Work Folders GPO**, and then click **OK**.
24. Close the Group Policy Management Console.

▶ Task 4: Validate Work Folders functionality

1. Sign in to LON-CL1 as **Adatum\Aidan** with the password **Pa$$w0rd**.
2. On Start screen, start typing **PowerShell**, and then click the **Windows PowerShell** icon in the Search pane.
3. At the Windows PowerShell command prompt, type **gpupdate /force**, and then press Enter.
4. Open **File Explorer** and navigate to **This PC**.
5. Verify that the **WorkFolders** folder is created.

Note: The presence of the Work Folders folder indicates that the Work Folders configuration is successful.

6. In **File Explorer**, create a few text files in the Work Folders folder.

Note: File Explorer displays the synchronization status of the files in the Work Folders folder.

7. Right-click the Windows button on the taskbar, and then click **Control Panel**.
8. In Control Panel, click **System and Security**, and then click **Work Folders**.
9. Click **Apply policies**. Click **Yes**.
10. Ensure that Work Folders are configured and working.
11. Close the Control Panel.

12. Sign in to LON-CL2 as **Adatum\Aidan** with the password **Pa$$word**.
13. On Start screen, start typing **PowerShell**, and then click the **Windows PowerShell** icon in the Search pane.
14. At the Windows PowerShell command prompt, type **gpupdate /force**, and then press Enter.
15. Open **File Explorer** and navigate to **This PC**.
16. Verify that the **WorkFolders** folder is created.
17. Right-click the Windows button on the taskbar, and then click **Control Panel**.
18. In Control Panel, click **System and Security**, and then click **Work Folders**.
19. Click **Apply policies**. Click **Yes**.
20. Open the Work Folders folder and verify that files that you created on LON-CL1 are present.

▶ Task 5: Prepare for the next module

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20412C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps two and three for **20412C-LON-SVR1**, **20412C-LON-SVR2, 20412C-LON-CL1**, and **20412C-LON-CL2**.

Results: After completing this exercise, you will have configured Work Folders.

Module 4: Implementing Distributed Active Directory® Domain Services Deployments

Lab: Implementing Distributed AD DS Deployments

Exercise 1: Implementing Child Domains in AD DS

▶ Task 1: Install a domain controller in a child domain

1. On TOR-DC1, in the Server Manager, click **Manage**, and from the drop-down list box, click **Add Roles and Features**.
2. On the **Before you begin** page, click **Next**.
3. On the **Select installation type** page, confirm that **Role-based or feature-based installation** is selected, and then click **Next**.
4. On the **Select destination server** page, ensure that **Select a server from the server pool** is selected, and that **TOR-DC1.adatum.com** is highlighted, and then click **Next**.
5. On the **Select server roles** page, click **Active Directory Domain Services**.
6. On the **Add features that are required for Active Directory Domain Services?** page, click **Add Features**.
7. On the **Select server roles** page, click **Next**.
8. On the **Select features** page, click **Next**.
9. On the **Active Directory Domain Services** page, click **Next**.
10. On the **Confirm installation selections** page, click **Install**. (This may take a few minutes to complete.)
11. When the Active Directory Domain Services (AD DS) binaries have installed, click the blue **Promote this server to a domain controller** link.
12. In the **Deployment Configuration** window, click **Add a new domain to an existing forest**.
13. Verify that **Select domain type** is set to **Child Domain**, and that **Parent domain name** is set to **Adatum.com**. In the **New domain name** text box, type **na**.
14. Confirm that **Supply the credentials to perform this operation** is set to **ADATUM\Administrator (Current user)**, and then click **Next**.

 **Note:** If this is not the case, then use the **Change** button to enter the credentials **Adatum\Administrator** and the password **Pa$$w0rd**.

15. In the Domain Controller Options window, ensure that **Domain functional level** is set to **Windows Server 2012 R2**.
16. Ensure that both the **Domain Name system (DNS) server** and **Global Catalog (GC)** check boxes are selected.
17. Confirm that **Site name:** is set to **Toronto**.
18. Under **Type the Directory Services Restore Mode (DSRM) password**, type **Pa$$w0rd** in both text boxes, and then click **Next**.

19. On the **DNS Options** page, click **Next**.
20. On the **Additional Options** page, click **Next**.
21. On the **Paths** page, click **Next**.
22. On the **Review Options** page, click **Next**.
23. On the **Prerequisites Check** page, confirm that there are no issues, and then click **Install**.

▶ Task 2: Verify the default trust configuration

1. Sign in to TOR-DC1 as **NA\Administrator** with the password **Pa$$w0rd**.
2. When the Server Manager opens, click **Local Server**.
3. Verify that **Windows Firewall** shows **Domain: Off**. If it does not, then next to **Ethernet**, click **172.16.0.25, IPv6 enabled**. Right-click **Ethernet**, and then click **Disable**. Right-click **Ethernet**, and then click **Enable**. The Local Area Connection should now show **Adatum.com**.
4. In the Server Manager, from the **Tools** menu, click **Active Directory Domains and Trusts**.
5. In the Active Directory Domains and Trusts console, expand **Adatum.com**, right-click **na.adatum.com**, and then click **Properties**.
6. In the **na.adatum.com Properties** dialog box, click the **Trusts** tab, and in the **Domain trusted by this domain (outgoing trusts)** box, click **Adatum.com**, and then click **Properties**,
7. In the **Adatum.com Properties** dialog box, click **Validate**, and then click **Yes, validate the incoming trust**.
8. In the **User name** text box, type **administrator**, and in the **Password** text box, type **Pa$$w0rd**, and then click **OK**.
9. When the message **The trust has been validated. It is in place and active** displays, click **OK**.

 **Note:** If you receive a message that the trust cannot be validated, or that the secure channel (SC) verification has failed, ensure that you have completed step 2 and then wait for at least 10 to 15 minutes. You can continue with the lab and come back later to verify this step.

10. Click **OK** twice to close the **Adatum.com Properties** dialog box.

Results: After completing this exercise, you will have implemented child domains in AD DS.

Exercise 2: Implementing Forest Trusts

▶ Task 1: Configure stub zones for DNS name resolution

1. On LON-DC1, in the Server Manager, click the **Tools** menu, and then from the drop-down menu, click **DNS**.
2. In the DNS tree pane, expand **LON-DC1**, right-click **Forward Lookup Zones**, and then click **New Zone**.
3. In the New Zone Wizard, click **Next**.
4. On the **Zone Type** page, click **Stub zone**, and then click **Next**.
5. On the **Active Directory Zone Replication Scope** page, click **To all DNS servers running on domain controllers in this forest: adatum.com**, and then click **Next**.
6. In the **Zone name:** text box, type **treyresearch.net**, and then click **Next**.
7. On the **Master DNS Servers** page, click **<Click here to add an IP Address or DNS Name>**, type **172.16.10.10**, click on the free space, and then click **Next**.
8. On the **Completing the New Zone Wizard** page, click **Next**, and then click **Finish**.
9. Select and then right-click the new stub zone **treyresearch.net**, and then click **Transfer from Master**.
10. Right-click **treyresearch.net**, and then click **Refresh**.
11. Confirm that the treyresearch.net stub zone contains records and then close DNS Manager.
12. Switch to TREY-DC1.
13. In the Server Manager, click the **Tools** menu, and from the drop-down menu, click **DNS**.
14. In the tree pane, expand **TREY-DC1**, select and then right-click **Forward Lookup Zones**, and then click **New Zone**.
15. In the New Zone Wizard, click **Next**.
16. On the **Zone Type** page, click **Stub zone**, and then click **Next**.
17. In the **Active Directory Zone Replication Scope** page, click **To all DNS servers running on domain controllers in this forest: Treyresearch.net**, and then click **Next**.
18. In the **Zone name:** text box, type **adatum.com**, and then click **Next**.
19. On the **Master DNS Servers** page, click **<Click here to add an IP Address or DNS Name>**, type **172.16.0.10**, click on the free space, and then click **Next**.
20. On the **Completing the New Zone Wizard** page, click **Next**, and then click **Finish**.
21. Select and then right-click the new stub zone **adatum.com**, and then click **Transfer from Master**.
22. Right-click **adatum.com**, and then click **Refresh**.
23. Confirm that the adatum.com stub zone contains records.
24. Close DNS Manager.

▶ Task 2: Configure a forest trust with selective authentication

1. On LON-DC1, from the **Tools** menu, click **Active Directory Domain and Trusts**.
2. In the **Active Directory Domains and Trusts** management console window, right-click **Adatum.com**, and then click **Properties**.

3. In the **Adatum.com Properties** dialog box, click the **Trusts** tab, and then click **New Trust**.
4. On the **New Trust Wizard** page, click **Next**.
5. In the **Name** text box, type **treyresearch.net**, and then click **Next**.
6. On the **Trust Type** page, click **Forest trust**, and then click **Next**.
7. On the **Direction of Trust** page, click **One-way: outgoing**, and then click **Next**.
8. On the **Sides of Trust** page, click **Both this domain and the specified domain**, and then click **Next**.
9. On the **User Name and Password** page, type **Administrator** as the user name and **Pa$$w0rd** as the password in the appropriate boxes, and then click **Next**.
10. On the **Outgoing Trust Authentication Level-Local Forest** page, click **Selective authentication**, and then click **Next**.
11. On the **Trust Selections Complete** page, click **Next**.
12. On the **Trust Creation Complete** page, click **Next**.
13. On the **Confirm Outgoing Trust** page, click **Next**.
14. Click **Finish**.
15. In the **Adatum.com Properties** dialog box, click the **Trusts** tab.
16. On the **Trusts** tab, under **Domains trusted by this domain (outgoing trusts)**, click **treyresearch.net** and then click **Properties**.
17. In the **treyresearch.net Properties** dialog box, click **Validate**.
18. Review the message that displays: **The trust has been validated. It is in place and active**.
19. Click **OK**, and then click **No** at the prompt.
20. Click **OK** twice.
21. Close Active Directory Domains and Trusts.

▶ Task 3: Configure a server for selective authentication

1. On LON-DC1, in the Server Manager, from the **Tools** menu, click **Active Directory Users and Computers**.
2. In the Active Directory Users and Computers console, from the **View** menu, click **Advanced Features**.
3. Expand **Adatum.com**, and then click **Computers**.
4. Right-click **LON-SVR2**, and then click **Properties**.
5. In the **LON-SVR2 Properties** dialog box, click the **Security** tab, and then click **Add**.
6. On the **Select Users, Computers, Service Accounts, or Groups** page, click **Locations**.
7. Click **treyresearch.net**, and then click **OK**.
8. In the **Enter the object name to select (examples:)** text box, type **IT**, and then click **Check Names**. When prompted for credentials, type **treyresearch\administrator** with the password **Pa$$w0rd**, and then click **OK**.
9. On the **Select Users, Computers, Service Accounts, or Groups** page, click **OK.**
10. In the **LON-SVR2 Properties** window, ensure that **IT (TreyResearch\IT)** is highlighted, select the **Allow** check box that is in line with **Allowed to authenticate**, and then click **OK**.
11. Switch to LON-SVR2.

12. On the taskbar, click the Windows® Explorer icon.
13. In the Windows Explorer window, click **Local Disk (C)**.
14. Right-click in the details pane, click **New**, and then click **Folder**.
15. In the **Name** text box, type **IT-Data**, and then press Enter.
16. Right-click **IT-Data**, point to **Share with**, and then click **Specific People**.
17. In the **File Sharing** dialog box, type **TreyResearch\IT**, and then click **Add**.
18. Click **Read**, and then click **Read/Write**. Click **Share**, and then click **Done**.
19. Sign out of TREY-DC1.
20. Sign in to **TREY-DC1** as **TreyResearch\Alice** with the password **Pa$$w0rd**.
21. Hover your pointer in the lower-right corner of the desktop, and when the sidebar displays, click **Search**.
22. In the **Search** text box, type **\\LON-SVR2\IT-Data**, and then press Enter. The folder will open.

▶ Task 4: To prepare for the next module

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

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20412C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20412C-TOR-DC1**, **20412C-TREY-DC1**, and **20412C-LON-SVR2**.

Results: After completing this exercise, you will have implemented forest trusts.

Module 5: Implementing Active Directory Domain Services Sites and Replication

Lab: Implementing AD DS Sites and Replication

Exercise 1: Modifying the Default Site

▶ Task 1: Install the Toronto domain controller

1. On TOR-DC1, in the Server Manager, click **Manage**, and from the drop-down list box, click **Add Roles and Features**.
2. On the **Before You Begin** page, click **Next**.
3. On the **Select installation type** page, confirm that **Role-based or feature-based installation** is selected, and then click **Next**.
4. On the **Select destination server** page, ensure that **Select a server from the server pool** is selected, and that **TOR-DC1.adatum.com** is highlighted, and then click **Next**.
5. On the **Select server roles** page, click **Active Directory Domain Services**.
6. On the **Add features that are required for Active Directory Domain Services?** page, 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**. (This may take a few minutes to complete.)
10. When the AD DS binaries have installed, do not click Close, but click the blue **Promote this server to a domain controller** link.
11. In the Deployment Configuration window, click **Add a domain controller to an existing domain**, and then click **Next**.
12. In the Domain Controller Options window, ensure that both the **Domain Name system (DNS) server** and **Global Catalog (GC)** check boxes are selected.
13. Confirm that **Site name:** is set to **Default-First-Site-Name**, and then under **Type the Directory Services Restore Mode (DSRM) password**, type **Pa$$w0rd** in both the **Password** and **Confirm password** boxes. Click **Next**.
14. On the **DNS Options** page, click **Next**.
15. In the **Additional Options** page, click **Next**.
16. In the Paths window, click **Next**.
17. In the **Review Options** window, click **Next**.
18. In the **Prerequisites Check** window, confirm that there are no issues, and then click **Install**. The server will restart automatically.
19. After TOR-DC1 restarts, sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.

▶ Task 2: Rename the default site

1. If necessary, on LON-DC1, open the Server Manager console.
2. In Server Manager, click **Tools**, and then click **Active Directory Sites and Services**.
3. In Active Directory Sites and Services, in the navigation pane, expand **Sites**.
4. Right-click **Default-First-Site-Name**, and then click **Rename**.
5. Type **LondonHQ**, and then press Enter.
6. Expand **LondonHQ**, expand the **Servers** folder, and then verify that both **LON-DC1** and **TOR-DC1** belong to the **LondonHQ** site.

▶ Task 3: Configure IP subnets associated with the default site

1. If necessary, on LON-DC1, open the Server Manager console, and then open **Active Directory Site and Services**.
2. In the Active Directory Sites and Services console, in the navigation pane, expand **Sites**, and then click the **Subnets** folder.
3. Right-click **Subnets**, and then click **New Subnet**.
4. In the **New Object – Subnet** dialog box, under **Prefix**, type **172.16.0.0/24**.
5. Under **Select a site object for this prefix**, click **LondonHQ**, and then click **OK**.

Results: After completing this exercise, you will have reconfigured the default site and assigned IP address subnets to the site.

Exercise 2: Creating Additional Sites and Subnets

▶ Task 1: Create the AD DS sites for Toronto

1. If necessary, on LON-DC1, open the Server Manager console, click **Tools**, and then click **Active Directory Sites and Services**.
2. In the Active Directory Sites and Services console, in the navigation pane, right-click **Sites**, and then click **New Site**.
3. In the **New Object – Site** dialog box, next to **Name**, type **Toronto**.
4. Under **Select a site link object for this site**, select **DEFAULTIPSITELINK**, and then click **OK**.
5. In the **Active Directory Domain Services** dialog box, click **OK**. The Toronto site displays in the navigation pane.
6. In the Active Directory Sites and Services console, in the navigation pane, right-click **Sites**, and then click **New Site**.
7. In the **New Object – Site** dialog box, next to **Name**, type **TestSite**.
8. Under **Select a site link object for this site**, select **DEFAULTIPSITELINK**, and then click **OK**. The TestSite site displays in the navigation pane.

▶ Task 2: Create IP subnets associated with the Toronto sites

1. If necessary, on LON-DC1, open the Server Manager console, click **Tools** and then click **Active Directory Sites and Services**.
2. In the Active Directory Sites and Services console, in the navigation pane, expand **Sites**, and then click the **Subnets** folder.
3. Right-click **Subnets**, and then click **New Subnet**.
4. In the **New Object – Subnet** dialog box, under **Prefix**, type **172.16.1.0/24**.
5. Under **Select a site object for this prefix**, click **Toronto**, and then click **OK**.
6. Right-click **Subnets**, and then click **New Subnet**.
7. In the **New Object – Subnet** dialog box, under **Prefix**, type **172.16.100.0/24**.
8. Under **Select a site object for this prefix**, click **TestSite**, and then click **OK**.
9. In the navigation pane, click the **Subnets** folder. Verify in the details pane that the three subnets are created and associated with their appropriate site.

Results: After this exercise, you will have created two additional sites representing the IP subnet addresses located in Toronto.

Exercise 3: Configuring AD DS Replication

▶ Task 1: Configure site links between AD DS sites

1. If necessary, on LON-DC1, open the Server Manager console, click **Tools**, and then click **Active Directory Sites and Services**.
2. In the Active Directory Sites and Services console, in the navigation pane, expand **Sites**, expand **Inter-Site Transports**, and then click the **IP** folder.
3. Right-click **IP**, and then click **New Site Link**.
4. In the **New Object – Site Link** dialog box, next to **Name**, type **TOR-TEST**.
5. Under **Sites not in this site link**, press CTRL on the keyboard, click **Toronto**, click **TestSite**, click **Add**, and then click **OK**.
6. Right-click **TOR-TEST**, and then click **Properties**.
7. In the **TOR-TEST Properties** dialog box, click **Change Schedule**.
8. In the **Schedule for TOR-TEST** dialog box, highlight the range from **Monday 9 AM to Friday 3 PM**, as follows:
   - Using the mouse, click at the Monday at 9:00 AM tile.
   - With the mouse button still pressed down, drag the cursor to the Friday at 3:00 PM tile.
9. Click **Replication Not Available** and then click **OK**.
10. Click **OK** to close **TOR-TEST Properties**.
11. Right-click **DEFAULTIPSITELINK**, and then click **Rename**.
12. Type **LON-TOR**, and then press Enter.
13. Right-click **LON-TOR**, and then click **Properties**.
14. Under **Sites in this link**, click **TestSite**, and then click **Remove**.
15. Next to **Replicate Every**, change the value to **60** minutes, and then click **OK**.

▶ Task 2: Move TOR-DC1 to the Toronto site

1. If necessary, on LON-DC1, click **Tools**, and then click **Active Directory Sites and Services**.
2. In Active Directory Sites and Services, in the navigation pane, expand **Sites**, expand **LondonHQ**, and then expand the **Servers** folder.
3. Right-click **TOR-DC1**, and then click **Move**.
4. In the **Move Server** dialog box, click **Toronto**, and then click **OK**.
5. In the navigation pane, expand the **Toronto** site, expand **Servers**, and then click **TOR-DC1**.

▶ Task 3: Monitor AD DS site replication

1. On LON-DC1, on the taskbar, click the Windows PowerShell icon.
2. At the Windows PowerShell prompt, type the following, and then press Enter:

```
Repadmin /kcc
```

   This command recalculates the inbound replication topology for the server.

3. At the Windows PowerShell prompt, type the following command, and then press Enter:

```
Repadmin /showrepl
```

4. Verify that the last replication with TOR-DC1 was successful.
5. At the Windows PowerShell prompt, type the following command, and then press Enter:

```
Repadmin /bridgeheads
```

This command displays the bridgehead servers for the site topology.

6. At the Windows PowerShell command prompt, type the following, and then press Enter:

```
Repadmin /replsummary
```

This command displays a summary of replication tasks. Verify that no errors appear.

7. At the Windows PowerShell command prompt, type the following, and then press Enter:

```
DCDiag /test:replications
```

8. Verify that all connectivity and replication tests pass successfully.
9. Switch to TOR-DC1, and then repeat steps 1 through 8 to view information from TOR-DC1. For step 4, verify that the last replication with LON-DC1 was successful.

Results: After this exercise, you will have configured site links and monitored replication.

Exercise 4: Monitoring and Troubleshooting AD DS Replication

▶ Task 1: Produce an error

1. If necessary, on LON-DC1, click **Tools**, and then click **Active Directory Sites and Services**.
2. In Active Directory Sites and Services, in the navigation pane, expand **Sites**, expand **LondonHQ**, expand the **Servers** folder, expand **LON-DC1**, and then select **NTDS Settings**.
3. In the **Details** pane, right click the**<automatically generated>** connection object and click **Replicate Now**.
4. In the **Replicate Now** pop-up, click **OK**.
5. In Active Directory Sites and Services, examine all the objects you created earlier, and on the taskbar, click the **Windows PowerShell** icon.
6. At the Windows PowerShell prompt, type the following, and then press Enter:

```
Get-ADReplicationUpToDatenessVectorTable –Target “adatum.com”
```

   This cmdlet will show you the last several replication events. Make a note of the date/time of the last (top) event.

7. Go to **TOR-DC1**.
8. At the Windows PowerShell prompt, type the following, and then press Enter:

```
\\LON-DC1\E$\Mod05\Mod05Ex4.ps1
```

▶ Task 2: Monitor AD DS site replication

1. If necessary, on TOR-DC1, open the **Server Manager** console, click **Tools** and then click **Active Directory Sites and Services**.
2. In the **Active Directory Sites and Services** console, in the navigation pane, expand **Sites**, then **Toronto**, then **Servers**, then **TOR-DC1**, and then select **NTDS Settings.**
3. In the details pane, right click the **<automatically generated>** and select **Replicate Now**.
4. The **Replicate Now** pop-up will appear, indicating an error informing you that “T**he RPC service is unavailable**.” Click **OK** on the **Replicate Now** pop-up.
5. On **TOR-DC1**, on the taskbar, click the **Windows PowerShell** icon.
6. At the Windows PowerShell prompt, type the following, and then press Enter:

```
Get-ADReplicationUpToDatenessVectorTable –Target “adatum.com”
```

   This cmdlet will show you the last several replication events. Note that the last date/time shown (Replication from LON-DC1) is not updating. This indicates that one-way replication is not occurring.

7. At the Windows PowerShell prompt, type the following, and then press Enter:

```
Get-AdReplicationSubnet –filter *
```

   This cmdlet will show detailed information about any subnets assigned to any sites.

8. Note that nothing is returned.

9. At the Windows PowerShell prompt, type the following, and then press Enter:

```
Get-AdReplicationSiteLink –filter *
```

This cmdlet will show detailed information about any site links assigned to particular sites.

10. Note that nothing is returned.

▶ Task 3: Troubleshoot AD DS replication

1. If necessary, on TOR-DC1, open the Windows PowerShell.
2. At the Windows PowerShell prompt, type the following, and then press Enter:

```
Ipconfig /all
```

3. Examine the results. Are they correct?
4. At the Windows PowerShell prompt, type the following, and then press Enter:

```
Get-DnsClient | Set-DnsClientServerAddress –ServerAddresses
("172.16.0.10","172.16.0.25")
```

5. Run the Ipconfig /all command again. You should get proper results.
6. If necessary, on **TOR-DC1**, open the **Server Manager** console, click **Tools**, and then click **Active Directory Sites and Services**.
7. In the **Active Directory Sites and Services** console, in the navigation pane, expand **Sites**, then **Toronto**, then **Servers**, then **TOR-DC1**, and then select **NTDS Settings.**
8. In the details pane, right click the **<automatically generated>**, and select **Replicate Now**.
9. The **Replicate Now** pop-up will appear, without an error. Click **OK**.
10. At the Windows PowerShell prompt, type the following, and then press Enter:

```
Get-DnsServer
```

11. You will get the following error: "Failed to retrieve DNS Server configuration information on TOR-DC1."
12. If necessary, on TOR-DC1, open the **Server Manager** console, click **Tools**, and then click **DNS**.
13. In the **Connect to DNS Server** pop-up window click **OK**, and then you should get another pop-up window stating "**The DNS Server is unavailable. Would you like to add it anyway?**" Click **No**, then Cancel and close any DNS window that appears.
14. At the Windows PowerShell prompt, type the following, and then press Enter:

```
Get-Service –DisplayName "DNS Server"
```

15. What is the status of the DNS Server service?
16. At the Windows PowerShell prompt, type the following, and then press Enter:

```
Start-Service –DisplayName "DNS Server"
```

After this completes, run the following again to ensure the service is running:

```
Get-Service –DisplayName "DNS Server"
```

17. If necessary, on TOR-DC1, open the **Server Manager** console, click **Tools**, and then click **Active Directory Sites and Services**.
18. In the **Active Directory Sites and Services** console, in the navigation pane, expand **Sites**, then **Toronto**, then **Servers**, then **TOR-DC1**, and then select **NTDS Settings.**
19. In the details pane, right click the **<automatically generated>**, and select **Replicate Now**.
20. The **Replicate Now** pop-up will appear, this time without an error. Click **OK**.
21. In **Active Directory Sites and Services**, examine all objects that you created earlier. Are any missing?
22. On **TOR-DC1**, open **File Explorer**. On the "**This PC**" URL bar, type the following, and then press enter: **\\LON-DC1\E$\Mod05**
23. Right-click the file named **Mod05EX4Fix.ps1** and select **Edit**.
24. The Windows PowerShell ISE will open. Examine the cmdlets in the script.
25. Find the section titled **#recreate site links for LON-TOR and TOR-Test**. Observe the line for the –Site Toronto. Note the value for the –Name entry. If it shows "172.16.1.0/22", change this to "172.16.1.0/24". The last number must be a 4, not a 2. Using the mouse, highlight all the text on the two lines beginning with **$schedule**, right-click them, and select **Run Selection**.
26. Find the section titled **#recreate Subnets**. Using the mouse, highlight all the text on the two lines beginning with "**New-ADReplicationSubnet**", right-click them, and select **Run Selection**.
27. In **Active Directory Sites and Services**, examine all the objects you created earlier. Ensure that the site link and subnet objects have been recreated.
28. On **LON-DC1** and **TOR-DC1**, close all open windows and sign off both virtual machines.

▶ Task 4: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. On the **Virtual Machines** list, right-click **20412C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps two and three for 20412C-TOR-DC1.

Module 6: Implementing AD CS

Lab A: Deploying and Configuring CA Hierarchy

Exercise 1: Deploying a Stand-Alone Root CA

▶ Task 1: Install and configure Active Directory Certificate Services on LON-CA1

1. Sign in to LON-CA1 as **Administrator** with the password **Pa$$w0rd**.
2. In the Server Manager, click **Add roles and features**.
3. On the **Before You Begin** page, click **Next**.
4. On the **Select installation type** page, click **Next**.
5. On the **Select destination server** page, click **Next**.
6. On the **Select server roles** page, select **Active Directory Certificate Services**. When the Add Roles and Features Wizard displays, click **Add Features**, and then click **Next**.
7. On the **Select features** page, click **Next**.
8. On the **Active Directory Certificate Services** page, click **Next**.
9. On the **Select role services** page, ensure that **Certification Authority** is selected, and then click **Next**.
10. On the **Confirm installation selections** page, click **Install**.
11. On the **Installation progress** page, after installation completes successfully, click the text **Configure Active Directory Certificate Services on the destination server**.
12. In the AD CS Configuration Wizard, on the **Credentials** page, click **Next**.
13. On the **Role Services** page, select **Certification Authority**, and then click **Next**.
14. On the **Setup Type** page, select **Stand-alone CA**, and then click **Next**.
15. On the **CA Type** page, ensure that **Root CA** is selected, and then click **Next**.
16. On the **Private Key** page, ensure that **Create a new private key** is selected, and then click **Next**.
17. On the **Cryptography for CA** page, keep the default selections for Cryptographic Service Provider (CSP) and Hash Algorithm, but set the **Key length** to **4096**, and then click **Next**.
18. On the **CA Name** page, in the **Common name for this CA** box, type **AdatumRootCA**, and then click **Next**.
19. On the **Validity Period** page, click **Next**.
20. On the **CA Database** page, click **Next**.
21. On the **Confirmation** page, click **Configure**.
22. On the **Results** page, click **Close**.
23. On the **Installation progress** page, click **Close**.
24. On LON-CA1, in Server Manager, click **Tools**, and then click **Certification Authority**.
25. In the certsrv – [Certification Authority (Local)] console, right-click **AdatumRootCA**, and then click **Properties**.

26. In the **AdatumRootCA Properties** dialog box, click the **Extensions** tab.
27. On the **Extensions** tab, in the **Select extension** drop-down list box, click **CRL Distribution Point (CDP)**, and then click **Add**.
28. In the **Location** box, type **http://lon-svr1.adatum.com/CertData/**, in the **Variable** drop-down list box, click **<CaName>**, and then click **Insert**.
29. In the **Variable** drop-down list box, click **<CRLNameSuffix>**, and then click **Insert**.
30. In the **Variable** drop-down list box, click **<DeltaCRLAllowed>**, and then click **Insert**.
31. In the **Location** box, position the cursor at the end of the URL, type **.crl**, and then click **OK**.
32. Select the following options, and then click **Apply**:
    - **Include in the CDP extension of issued certificates**
    - **Include in CRLs. Clients use this to find Delta CRL locations**
33. In the Certification Authority pop-up window, click **No**.
34. In the **Select extension** drop-down list box, click **Authority Information Access (AIA)**, and then click **Add**.
35. In the **Location** box, type **http://lon-svr1.adatum.com/CertData/**, then in the **Variable** drop-down list box, click **<ServerDNSName>**, and then click **Insert**.
36. In the **Location** box, type an underscore (_), then in the **Variable** drop-down list box, click **<CaName>**, and then click **Insert**. Position the cursor at the end of the URL.
37. In the **Variable** drop-down list box, click **<CertificateName>**, and then click **Insert**.
38. In the **Location** box, position the cursor at the end of the URL, type **.crt**, and then click **OK**.
39. Select the **Include in the AIA extension of issued certificates** check box, and then click **OK**.
40. Click **Yes** to restart the Certification Authority service.
41. In the Certification Authority console, expand **AdatumRootCA**, right-click **Revoked Certificates**, point to **All Tasks**, and then click **Publish**.
42. In the Publish CRL window, click **OK**.
43. Right-click **AdatumRootCA**, and then click **Properties**.
44. In the **AdatumRootCA Properties** dialog box, click **View Certificate**.
45. In the **Certificate** dialog box, click the **Details** tab.
46. On the **Details** tab, click **Copy to File**.
47. In the Certificate Export Wizard, on the **Welcome** page, click **Next**.
48. On the **Export File Format** page, select **DER encoded binary X.509 (.CER)**, and then click **Next**.
49. On the **File to Export** page, click **Browse**. In the **File name** box, type **\\lon-svr1\C$**, and then press Enter.
50. In the **File name** box, type **RootCA**, click **Save**, and then click **Next**.
51. Click **Finish**, and then click **OK** three times.
52. Open a File Explorer window, and browse to **C:\Windows\System32\CertSrv\CertEnroll**.
53. In the Cert Enroll folder, click both files, right-click the highlighted files, and then click **Copy**.
54. In the File Explorer address bar, type **\\lon-svr1\C$**, and then press Enter.

55. Right-click the empty space, and then click **Paste**.
56. Close File Explorer.

▶ Task 2: Creating a DNS host record for LON-CA1 and configure sharing

1. ON LON-DC1, in the Server Manager, click **Tools**, and then click **DNS**.
2. In the DNS Manager console, expand **LON-DC1**, expand **Forward Lookup Zones**, click **Adatum.com**, right-click **Adatum.com**, and then click **New Host (A or AAAA)**.
3. In the New Host window, in the **Name** box, type **LON-CA1**.
4. In the IP address window, type **172.16.0.40**, click **Add Host**, click **OK**, and then click **Done**.
5. Close the DNS Manager.
6. Switch to LON-CA1.
7. On **Start** screen, click **Control Panel**.
8. In the Control Panel window, click **View network status and tasks**.
9. In the Network and Sharing Center window, click **Change advanced sharing settings**.
10. Under **Guest or Public (current profile)**, select the **Turn on file and printer sharing** option, and then click **Save changes**.

Results: After completing this exercise, you will have deployed a root stand-alone CA.

Exercise 2: Deploying an Enterprise Subordinate CA

▶ Task 1: Install and configure AD CS on LON-SVR1

1. Sign in to LON-SVR1 as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. In the Server Manager, click **Add roles and features**.
3. On the **Before You Begin** page, click **Next**.
4. On the **Select installation type** page, click **Next**.
5. On the **Select destination server page**, click **Next**.
6. On the **Select server roles** page, select **Active Directory Certificate Services**.
7. When the Add Roles and Features Wizard displays, click **Add Features**, and then click **Next**.
8. On the **Select features** page, click **Next**.
9. On the **Active Directory Certificate Services** page, click **Next**.
10. On the **Select role services** page, ensure that **Certification Authority** is selected already, and then select **Certification Authority Web Enrollment**.
11. When the Add Roles and Features Wizard displays, click **Add Features**, and then click **Next**.
12. On the **Confirm installation selections** page, click **Install**.
13. On the **Installation progress** page, after installation is successful, click the text **Configure Active Directory Certificate Services on the destination server**.
14. In the AD CS Configuration Wizard, on the **Credentials** page, click **Next**.
15. On the **Role Services** page, select both **Certification Authority** and **Certification Authority Web Enrollment**, and then click **Next**.
16. On the **Setup Type** page, select **Enterprise CA**, and then click **Next**.
17. On the **CA Type** page, click **Subordinate CA**, and then click **Next**.
18. On the **Private Key** page, ensure that **Create a new private key** is selected, and then click **Next**.
19. On the **Cryptography for CA** page, keep the default selections, and then click **Next**.
20. On the **CA Name** page, in the **Common name for this CA** box, type **Adatum-IssuingCA**, and then click **Next**.
21. On the **Certificate Request** page, ensure that **Save a certificate request to file on the target machine** is selected, and then click **Next**.
22. On the **CA Database** page, click **Next**.
23. On the **Confirmation** page, click **Configure**.
24. On the **Results** page, click **Close**.
25. On the **Installation progress** page, click **Close**.

▶ Task 2: Install a subordinate CA certificate

1. On LON-SVR1, open a File Explorer window, and then navigate to **Local Disk (C:)**.
2. Right-click **RootCA.cer**, and then click **Install Certificate**.
3. In the Certificate Import Wizard, click **Local Machine**, and then click **Next**.

4. On the **Certificate Store** page, click **Place all certificates in the following store**, and then click **Browse**.
5. Select **Trusted Root Certification Authorities**, click **OK**, click **Next**, and then click **Finish**.
6. When the Certificate Import Wizard window appears, click **OK**.
7. In the File Explorer window, select the **AdatumRootCA.crl** and **LON-CA1_AdatumRootCA.crt** files, right-click the files, and then click **Copy**.
8. Double-click **inetpub**.
9. Double-click **wwwroot**.
10. Create a new folder, and then name it **CertData**.
11. Paste the two copied files into that folder.
12. Switch to Local Disk (C:).
13. Right-click the file **LON-SVR1.Adatum.com_Adatum-LON-SVR1-CA.req**, and then click **Copy**.
14. In the File Explorer address bar, type **\\LON-CA1\C$**, and then press Enter.
15. In the File Explorer window, right-click an empty space, and then click **Paste**. Make sure that the request file is copied to LON-CA1.
16. Switch to the LON-CA1server.
17. In the Certificate Authority console, right-click **AdatumRootCA**, point to **All Tasks**, and then click **Submit new request**.
18. In the Open Request File window, navigate to **Local Disk (C:)**, click file **LON-SVR1.Adatum.com_Adatum-LON-SVR1-CA.req**, and then click **Open**.
19. In the Certification Authority console, click the **Pending Requests** container. Right-click **Pending Requests**, and then click **Refresh**.
20. In the details pane, right-click the request (with ID 2), point to **All Tasks**, and then click **Issue**.
21. In the Certification Authority console, click the **Issued Certificates** container.
22. In the details pane, double-click the certificate, click the **Details** tab, and then click **Copy to File**.
23. In the Certificate Export Wizard, on the **Welcome** page, click **Next**.
24. On the **Export File Format** page, click **Cryptographic Message Syntax Standard – PKCS #7 Certificates (.P7B)**, click **Include all certificates in the certification path if possible**, and then click **Next**.
25. On the **File to Export** page, click **Browse**.
26. In the **File name** box, type **\\lon-svr1\C$**, and then press Enter.
27. In the **File name** box, type **SubCA**, click **Save**, click **Next**, click **Finish**, and then click **OK** twice.
28. Switch to LON-SVR1.
29. In the Server Manager, click **Tools**, and then click **Certification Authority**.
30. In the Certification Authority console, right-click **Adatum-IssuingCA**, point to **All Tasks**, and then click **Install CA Certificate**.
31. Navigate to **Local Disk (C:)**, click the **SubCA.p7b** file, and then click **Open**.
32. Wait for 15 to 20 seconds, and then on the toolbar, click the green icon to start the CA service.
33. Ensure that the CA starts successfully.

▶ Task 3: Publish the root CA certificate through Group Policy

1. On LON-DC1, on the taskbar, click the **Server Manager** icon.
2. In the Server Manager, click **Tools**, and then click **Group Policy Management**.
3. In the Group Policy Management Console, expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**, right-click **Default Domain Policy**, and then click **Edit**.
4. In the Computer Configuration node, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, expand **Public Key Policies**, right-click **Trusted Root Certification Authorities**, click **Import**, and then click **Next**.
5. On the **File to Import** page, click **Browse**.
6. In the **file name** box, type **\\lon-svr1\C$**, and then press Enter.
7. Click file **RootCA.cer**, and then click **Open**.
8. Click **Next** two times, and then click **Finish**.
9. When the Certificate Import Wizard window appears, click **OK**.
10. Close the Group Policy Management Editor and the Group Policy Management Console.

▶ Task 4: Prepare for the next lab

Keep all virtual machines running for the next lab. Do not revert any virtual machines.

Results: After completing this exercise, you will have deployed and configured an enterprise subordinate CA.

Lab B: Deploying and Managing Certificates

Exercise 1: Configuring Certificate Templates

▶ Task 1: Create a new template based on the web server template

1. On LON-SVR1, in the Certification Authority console, expand **Adatum-IssuingCA**, right-click **Certificate Templates**, and then select **Manage**.
2. In the Certificate Templates console, locate the **Web Server** template in the list, right-click it, and then click **Duplicate Template**.
3. Click the **General** tab.
4. In the **Template display name** field, type **Adatum WebSrv**, and set the **Validity period** to **3 years**.
5. Click the **Request Handling** tab, select **Allow private key to be exported**, and then click **OK**.

▶ Task 2: Create a new template for users that includes smart card logon

1. In the Certificate Templates console, right-click the **User** certificate template, and then click **Duplicate Template**.
2. In the **Properties of New Template** dialog box, click the **General** tab, and in the **Template display name** text box, type **Adatum User**.
3. On the **Subject Name** tab, clear both the **Include e-mail name in subject name** and the **E-mail name** check boxes.
4. On the **Extensions** tab, click **Application Policies**, and then click **Edit**.
5. In the **Edit Application Policies Extension** dialog box, click **Add**.
6. In the **Add Application Policy** dialog box, select **Smart Card Logon**, and then click **OK** twice.
7. Click the **Superseded Templates** tab, and then click **Add**.
8. Click the **User** template, and then click **OK**.
9. On the **Security** tab, click **Authenticated Users**. Under **Permissions for Authenticated Users**, select the **Allow** check box for **Read**, **Enroll**, and **Autoenroll**, and then click **OK**.
10. Close the Certificate Templates console.

▶ Task 3: Configure the templates so that they can be issued

1. On LON-SVR1, in the Certification Authority console, right-click **Certificate Templates**, point to **New**, and then click **Certificate Template to Issue**.
2. In the Enable Certificate Templates window, select **Adatum User** and **Adatum WebSrv**, and then click **OK**.

▶ Task 4: Update the web server certificate on the LON-SVR2 web server

1. Sign in to LON-SVR2 as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. From the taskbar, click the **Windows PowerShell** icon.
3. At the Windows PowerShell prompt, type **gpupdate /force**, and then press Enter.
4. If prompted, restart the server, and sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
5. On the taskbar, click the **Server Manager** icon.
6. From **Server Manager**, click **Tools**, and then click **Internet Information Services (IIS) Manager**.

7. In the IIS console, click **LON-SVR2**, at the **Internet Information Services (IIS) Manager** prompt, click **No**, and then in the central pane, double-click **Server Certificates**.
8. In the Actions pane, click **Create Domain Certificate**.
9. On the **Distinguished Name Properties** page, complete the following fields, and then click **Next**:
    - Common name: **lon-svr2.adatum.com**
    - Organization: **Adatum**
    - Organizational Unit: **IT**
    - City/locality: **Seattle**
    - State/province: **WA**
    - Country/region: **US**
10. On the **Online Certification Authority** page, click **Select**.
11. Click **Adatum-IssuingCA**, and then click **OK**.
12. In the **Friendly name** text box, type **lon-svr2**, and then click **Finish**.
13. Ensure that the certificate displays in the Server Certificates console.
14. In the IIS console, expand **LON-SVR2**, expand **Sites**, and then click **Default Web Site**.
15. In the Actions pane, click **Bindings**.
16. In the Site Bindings window, click **Add**.
17. In the **Type** drop-down list box, click **https**.
18. In the **SSL certificate** drop-down list box, click **lon-svr2**, click **OK**, and then click **Close**.
19. Close the IIS console.

Results: After completing this exercise, you will have created and published new certificate templates.

Exercise 2: Configuring Certificate Enrollment

▶ Task 1: Configure autoenrollment for users

1. On LON-DC1, in the Server Manager, click **Tools**, and then click **Group Policy Management**.
2. Expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**, right-click **Default Domain Policy**, and then click **Edit**.
3. Expand **User Configuration**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, and then click to highlight **Public Key Policies**.
4. In the right pane, double-click **Certificate Services Client – Auto-Enrollment**.
5. In the **Configuration Model** drop-down list box, click **Enabled**.
6. Select the **Renew expired certificates**, **update pending certificates, and remove revoked certificates** option.
7. Select the **Update certificates that use certificate templates** option.
8. Click **OK** to close the properties window.
9. In the right pane, double-click the **Certificate Services Client – Certificate Enrollment Policy** object.
10. On the **Enrollment Policy** tab, set the **Configuration Model** to **Enabled**, and ensure that the certificate enrollment policy list displays the **Active Directory Enrollment Policy** (it should have a checkmark next to it, and display a status of **Enabled**).
11. Click **OK** to close the window.
12. Close both the Group Policy Management Editor and the GPMC.

▶ Task 2: Verify autoenrollment

1. On LON-SVR1, from the taskbar, click the **Windows PowerShell** icon.
2. At the Windows PowerShell prompt, type **gpupdate /force**, and then press Enter.
3. After the policy refreshes, type **mmc.exe**, and then press Enter.
4. In Console1, click **File**, and then in the **File** menu, click **Add/Remove Snap-in**.
5. Click **Certificates**, and then click **Add**.
6. Click **Finish**, and then click **OK**.
7. Expand **Certificates – Current User**, expand **Personal**, and then click **Certificates**.
8. Verify that a certificate based on **Adatum User** template is issued for Administrator.
9. Close Console1. Click **No** without saving the changes.

▶ Task 3: Configure the Enrollment Agent for smart card certificates

1. On LON-SVR1, in the Server Manager console, click **Tools**, and then open **Certification Authority**.
2. In the certsrv console, expand **Adatum-IssuingCA**, right-click **Certificate Templates**, and then click **Manage**.
3. In the Certificate Templates console, double-click **Enrollment Agent**.
4. Click the **Security** tab, and then click **Add**.
5. In the Select Users, Computers, Service Accounts, or Groups window, type **Allie**, click **Check Names**, and then click **OK**.

6. On the **Security** tab, click **Allie Bellew**, select **Allow** for **Read** and **Enroll** permissions, and then click **OK**.
7. Close the Certificate Templates console.
8. In the certsrv console, right-click **Certificate Templates**, point to **New**, and then click **Certificate Template to Issue**.
9. In the list of templates, click **Enrollment Agent**, and then click **OK**.
10. Switch to LON-CL1, and sign in as **Adatum\Allie** with the password **Pa$$w0rd**.
11. Open a command-prompt window, and at the command prompt, type **mmc.exe**, and then press Enter.
12. In Console1, click **File**, and then click **Add/Remove Snap-in**.
13. Click **Certificates**, click **Add**, and then click **OK**.
14. Expand **Certificates – Current User**, expand **Personal**, click **Certificates**, right-click **Certificates**, point to **All Tasks**, and then click **Request New Certificate**.
15. In the Certificate Enrollment Wizard, on the **Before You Begin** page, click **Next**.
16. On the **Select Certificate Enrollment Policy** page, click **Next**.
17. On the **Request Certificates** page, select **Enrollment Agent**, and then click **Enroll**.
18. Click **Finish**.
19. Switch to **LON-SVR1**.
20. In the **Certification Authority** console, right-click **Adatum-IssuingCA**, and then click **Properties**.
21. Click the **Enrollment Agents** tab.
22. Click **Restrict Enrollment agents**.
23. On the pop-up window that displays, click **OK**.
24. In the **Enrollment agents** section, click **Add**.
25. In the **Select User, Computer or Group** field, type **Allie**, click **Check Names**, and then click **OK**.
26. Click **Everyone**, and then click **Remove**.
27. In the **certificate templates** section, click **Add**.
28. In the list of templates, select **Adatum User**, and then click **OK**.
29. In the **Certificate Templates** section, click **<All>**, and then click **Remove**.
30. In the **Permission** section, click **Add**.
31. In the **Select User, Computer or Group** field, type **Marketing**, click **Check Names**, and then click **OK**.
32. In the Permission section, click **Everyone**, and then click **Remove**.
33. Click **OK**.

Results: After completing this exercise, you will have configured and verified autoenrollment for users, and configured an Enrollment Agent for smart cards.

Exercise 3: Configuring Certificate Revocation

▶ Task 1: Configure Certified Revocation List (CRL) distribution

1. On LON-SVR1, in the Certification Authority console, right-click **Revoked Certificates**, and then click **Properties**.
2. In the **Revoked Certificates Properties** dialog box, set the **CRL publication interval** to **1 Days**, and the **Delta CRL** publication interval to **1 Hours**, and then click **OK**.
3. Right-click **Adatum-IssuingCA**, and then click **Properties**.
4. In the **Adatum-IssuingCA Properties** dialog box, click the **Extensions** tab, and review the values for **CDP**.
5. Click **Cancel**.

▶ Task 2: Install and configure an Online Responder

1. On LON-SVR1, on the taskbar, click the **Server Manager** icon.
2. In the Server Manager, click **Add roles and features**.
3. Click **Next** three times.
4. On the **Select server roles** page, expand **Active Directory Certificate Services (Installed)**, and then click **Online Responder**.
5. Click **Add Features**.
6. Click **Next** two times, and then click **Install**.
7. When the message displays that installation succeeded, click **Configure Active Directory Certificate Services on the destination server**.
8. In AD CS Configuration Wizard, click **Next**.
9. Click **Online Responder**, and then click **Next**.
10. Click **Configure**, and then click **Close** two times.
11. On LON-SVR1, open the Certification Authority console.
12. In the Certification Authority console, right-click **Adatum-IssuingCA**, and then click **Properties**.
13. In the **Adatum-IssuingCA Properties** dialog box, on the **Extensions** tab, in the **Select extension** list, click **Authority Information Access (AIA**), and then click **Add**.
14. In the **Add Location** dialog box, type **http://LON-SVR1/ocsp**, and then click **OK**.
15. Select the **Include in the AIA extension of issued certificates** check box.
16. Select the **Include in the online certificate status protocol (OCSP) extension** check box, and then click **OK**.
17. In the **Certificate Authority** dialog box, restart AD CS by clicking **Yes**.
18. In the certsrv console, expand **Adatum-IssuingCA**, right-click the **Certificate Templates** folder, and then click **Manage**.
19. In the **Certificate Templates** console, double-click the **OCSP Response Signing** template.
20. In the **OCSP Response Signing Properties** dialog box, click the **Security** tab, and under **Permissions for Authenticated Users**, select **Allow** for the **Enroll** check box, and then click **OK**.
21. Close the Certificate Templates console.

22. In the Certification Authority console, right-click the **Certificate Templates** folder, point to **New**, and then click **Certificate Template to Issue**.
23. In the **Enable Certificate Templates** dialog box, select the **OCSP Response Signing** template, and then click **OK**.
24. On LON-SVR1, in the Server Manager, click **Tools**, and then click **Online Responder Management**.
25. In the OCSP Management console, right-click **Revocation Configuration**, and then click **Add Revocation Configuration**.
26. In the Add Revocation Configuration Wizard, click **Next**.
27. On the **Name the Revocation Configuration** page, in the **Name** box, type **AdatumCA Online Responder**, and then click **Next**.
28. On the **Select CA Certificate Location** page, click **Next**.
29. On the **Choose CA Certificate** page, click **Browse**, click the **Adatum-IssuingCA** certificate, click **OK**, and then click **Next**.
30. On the **Select Signing Certificate** page, verify that **Automatically select a signing certificate is selected**, and **Auto-Enroll for an OCSP signing certificate** are both selected, and then click **Next**.
31. On the **Revocation Provider** page, click **Finish**. The revocation configuration status will appear as **Working**.
32. Close the Online Responder console.

Results: After completing this exercise, you will have configured certificate revocation settings.

Exercise 4: Configuring Key Recovery

▶ Task 1: Configure the CA to issue KRA certificates

1. On LON-SVR1, open the Certification Authority console.
2. In the Certification Authority console, expand the **Adatum-IssuingCA** node, right-click the **Certificates Templates** folder, and then click **Manage**.
3. In the Details pane, right-click the **Key Recovery Agent** certificate, and then click **Properties**.
4. In the **Key Recovery Agent Properties** dialog box, click the **Issuance Requirements** tab.
5. Clear the **CA certificate manager approval** check box.
6. Click the **Security** tab. Notice that Domain Admins and Enterprise Admins are the only groups that have the Enroll permission, and then click **OK**.
7. Close the Certificate Templates console.
8. In the Certification Authority console, right-click **Certificate Templates**, point to **New**, and then click **Certificate Template to Issue**.
9. In the **Enable Certificate Templates** dialog box, click the **Key Recovery Agent** template, and then click **OK**.
10. Close the Certification Authority console.

▶ Task 2: Acquire the KRA certificate

1. On LON-SVR1, on the taskbar, click the **Windows PowerShell** icon.
2. At the Windows PowerShell prompt, type **MMC.exe**, and then press Enter.
3. In the Console1-[Console Root] console, click **File**, and then click **Add/Remove Snap-in**.
4. In the **Add or Remove Snap-ins** dialog box, click **Certificates**, and then click **Add**.
5. In the **Certificates snap-in** dialog box, select **My user account**, click **Finish**, and then click **OK**.
6. Expand the **Certificates - Current User** node, right-click **Personal**, point to **All Tasks**, and then click **Request New Certificate**.
7. In the Certificate Enrollment Wizard, on the **Before You Begin** page, click **Next**.
8. On the **Select Certificate Enrollment Policy** page, click **Next**.
9. On the **Request Certificates** page, select the **Key Recovery Agent** check box, click **Enroll**, and then click **Finish**.
10. Refresh the console, and view the KRA in the personal store; that is, scroll across the certificate properties and verify that the Certificate Template Key Recovery Agent is present.
11. Close Console1 without saving changes.

▶ Task 3: Configure the CA to allow key recovery

1. On LON-SVR1, open the Certification Authority console.
2. On LON-SVR1, in the Certification Authority console, right-click **Adatum-IssuingCA**, and then click **Properties**.
3. In the **Adatum-IssuingCA Properties** dialog box, click the **Recovery Agents** tab, and then select **Archive the key**.
4. Under **Key recovery agent certificates**, click **Add**.

5. In the **Key Recovery Agent Selection** dialog box, click the certificate that is for Key Recovery Agent purpose (it will most likely be last on the list or you can click the link **Click here to view the certificate properties** for each certificate on the list to ensure that you select the right certificate), and then click **OK** twice.
6. When prompted to restart the CA, click **Yes**.

▶ Task 4: Configure a custom template for key archival

1. On LON-SVR1, in the Certification Authority console, right-click the **Certificates Templates** folder, and then click **Manage**.
2. In the Certificate Templates console, right-click the **User** certificate, and then click **Duplicate Template**.
3. In the **Properties of New Template** dialog box, on the **General** tab, in the **Template display name** box, type **Archive User**.
4. On the **Request Handling** tab, select the **Archive subject's encryption private key** check box.
5. If a pop-up window displays, click **OK**.
6. Click the **Subject Name** tab, clear both the **E-mail name** and **Include e-mail name in subject name** check boxes, and then click **OK**.
7. Close the Certificate Templates console.
8. In the Certification Authority console, right-click the **Certificates Templates** folder, point to **New**, and then click **Certificate Template to Issue**.
9. In the **Enable Certificate Templates** dialog box, click the **Archive User** template, and then click **OK**.
10. Close the Certification Authority console.

▶ Task 5: Verify key archival functionality

1. Sign in to LON-CL1 as **Adatum\Aidan**, using the password **Pa$$w0rd**.
2. On the Start screen, type **mmc.exe**, and then press Enter. Click **Yes** in the User Account Control dialog box.
3. In the Console1-[Console Root] console, click **File**, and then click **Add/Remove Snap-in**.
4. In the **Add or Remove Snap-ins** dialog box, click **Certificates**, click **Add**, click **Finish** and then click **OK**.
5. Expand the **Certificates - Current User** node, right-click Personal, click **All Tasks**, and then click **Request New Certificate**.
6. In the Certificate Enrollment Wizard, on the **Before You Begin** page, click **Next** twice.
7. On the **Request Certificate** page, select the **Archive User** check box, click **Enroll**, and then click **Finish**.
8. Refresh the console, and view that a certificate is issued to Aidan, based on the **Archive User** certificate template.
9. Simulate the loss of a private key by deleting the certificate. In the central pane, right-click the certificate that you just enrolled, select **Delete**, and then click **Yes** to confirm.
10. Switch to LON-SVR1.
11. Open the Certification Authority console, expand **Adatum-IssuingCA**, and then click the **Issued Certificates** store.

12. In the details pane, double-click a certificate with Requestor Name **Adatum\Aidan**, and Certificate Template name of **Archive User**.
13. Click the **Details** tab, copy the **Serial Number**, and then click **OK**. (You may either copy the number to Notepad—select it and press CTRL+C— or write it down on paper.)
14. On the taskbar, click the **Windows PowerShell** icon.
15. At the Windows PowerShell prompt, type the following command, (where <*serial number*> is the serial number that you copied), and then press Enter:

```
Certutil –getkey <serial number> outputblob
```

Note: If you paste the serial number from Notepad, remove spaces between numbers.

16. Verify that **outputblob** file now displays in the **C:\Users\Administrator.Adatum** folder.
17. To convert the **outputblob** file into a .pfx file, at the Windows PowerShell prompt, type the following command, and then press Enter:

```
Certutil –recoverkey outputblob aidan.pfx
```

18. When prompted for the new password, type **Pa$$w0rd**, and then confirm the password.
19. After the command executes, close Windows PowerShell.
20. Browse to **C:\Users\Administrator.ADATUM**, and then verify that **aidan.pfx**—the recovered key—is created.
21. Switch to LON-CL1 machine.
22. On the **Start** screen, type **Control Panel** and then click on **Control Panel**.
23. In the Control Panel window, click View network status and tasks.
24. In the Network and Sharing Center window, click Change advanced sharing settings.
25. Under Guest or Public (current profile), select the option Turn on file and printer sharing.
26. Click Save changes.
27. If asked for credentials, use **Adatum\administrator** as the user name, and **Pa$$w0rd** as the password.
28. Switch back to the LON-SVR1 machine.
29. Copy the **aidan.pfx** file to **\\lon-cl1\C$**.
30. Switch to **LON-CL1**, and ensure that you are still logged on as Aidan.
31. Browse to drive **C**, and double-click the **aidan.pfx** file.
32. On the **Welcome to the Certificate Import Wizard** page, click **Next**.
33. On the **File to Import** page, click **Next**.
34. On the **Password** page, enter **Pa$$w0rd** as the password, and then click **Next**.
35. On the **certificate store** page, click **Next**, click **Finish**, and then click **Ok**.
36. In the Console1-[Console Root\Certificates - Current User\Personal\Certificates] expand the **Certificates** - **Current User** node, expand **Personal**, and then click **Certificates**.
37. Refresh the console, and verify that the certificate for **Aidan** is restored.

▶ Task 6: Prepare for the next module

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

1. On the host computer, start **Hyper-V Manager**.
2. On the **Virtual Machines** list, right-click **20412C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps two and three for 20412C-LON-CL1, 20412C-LON-SVR1, 20412C-LON-CA1, and 20412C-LON-SVR2.

Results: After completing this exercise, you will have implemented key archival and tested private key recovery.

Module 7: Implementing Active Directory Rights Management Services

Lab: Implementing AD RMS

Exercise 1: Installing and Configuring AD RMS

▶ Task 1: Configure Domain Name System (DNS) and the Active Directory® Rights Management Services (AD RMS) service account

1. Sign in to LON-DC1 with the **Adatum\Administrator** account and the password **Pa$$w0rd**.
2. In the Server Manager, click **Tools**, and then click **Active Directory Administrative Center**.
3. Select and then right-click **Adatum (local)**, click **New**, and then click **Organizational Unit**.
4. In the **Create Organizational Unit** dialog box, in the **Name** field, type **Service Accounts**, and then click **OK**.
5. Right-click the **Service Accounts** organizational unit (OU), click **New**, and then click **User**.
6. On the **Create User** dialog box, enter the following details, and then click **OK**:
   - First name: **ADRMSSVC**
   - User UPN logon: **ADRMSSVC**
   - Password: **Pa$$w0rd**
   - Confirm Password: **Pa$$w0rd**
   - Password never expires: **Enabled**
   - User cannot change password: **Enabled**
7. Right-click the **Users** container, click **New**, and then click **Group**.
8. In the **Create Group** dialog box, enter the following details, and then click **OK**:
   - Group name: **ADRMS_SuperUsers**
   - E-mail: **ADRMS_SuperUsers@adatum.com**
9. Right-click the **Users** container, click **New**, and then click **Group**.
10. In the **Create Group** dialog box, enter the following details, and then click **OK**.
    - Group name: **Executives**
    - E-mail: **executives@adatum.com**
11. Double-click the **Managers** OU.
12. Hold down the Ctrl key, and click the following users:
    - **Aidan Delaney**
    - **Bill Malone**
13. In the Tasks pane, click **Add to group**.
14. In the **Select Groups** dialog box, type **Executives**, and then click **OK**.
15. Close the Active Directory Administrative Center.
16. In the Server Manager, click **Tools**, and then click **DNS**.

17. In the DNS Manager console, expand **LON-DC1**, and then expand **Forward Lookup Zones**.
18. Select and then right-click **Adatum.com**, and then click **New Host (A or AAAA)**.
19. In the **New Host** dialog box, enter the following information, and then click **Add Host**:
    - Name: **adrms**
    - IP address: **172.16.0.21**
20. Click **OK**, and then click **Done**.
21. Close the DNS Manager console.

▶ Task 2: Install and configure the AD RMS server role

1. Sign in to LON-SVR1 with the **Adatum\Administrator** account and the password **Pa$$w0rd**.
2. In the **Server Manager**, click **Manage**, and then click **Add roles and features**.
3. In the **Add Roles and Features Wizard**, click **Next** three times.
4. On the **Server Roles** page, click **Active Directory Rights Management Services**.
5. In the **Add Roles and Features** dialog box, click **Add Features**, and then click **Next** four times.
6. Click **Install**, and then click **Close**.
7. In the **Server Manager**, click the **AD RMS** node.
8. Next to **Configuration required for Active Directory Rights Management Services** at LON-SVR1, click **More**.
9. On the **All Servers Task Details and Notifications** page, click **Perform Additional Configuration**.
10. In the **AD RMS Configuration: LON-SVR1.Adatum.com** dialog box, click **Next**.
11. On the **AD RMS Cluster** page, click **Create a new AD RMS root cluster**, and then click **Next**.
12. On the **Configuration Database** page, click **Use Windows Internal Database on this server**, and then click **Next**.
13. On the **Service Account** page, click **Specify**.
14. In the **Windows Security** dialog box, enter the following details, click **OK**, and then click **Next**:
    - Username: **ADRMSSVC**
    - Password: **Pa$$w0rd**
15. On the **Cryptographic Mode** page, click **Cryptographic Mode 2**, and then click **Next**.
16. On the **Cluster Key Storage** page, click **Use AD RMS centrally managed key storage**, and then click **Next**.
17. On the **Cluster Key Password** page, enter the password **Pa$$w0rd** twice, and then click **Next**.
18. On the **Cluster Web Site** page, verify that **Default Web Site** is selected, and then click **Next**.
19. On the **Cluster Address** page, provide the following information, and then click **Next**:
    - Connection Type: **Use an unencrypted connection (http://)**
    - Fully Qualified Domain Name: **adrms.adatum.com**
    - Port: **80**
20. On the **Licensor Certificate** page, type **Adatum AD RMS**, and then click **Next**.
21. On the **SCP Registration** page, click **Register the SCP now**, and then click **Next**.

22. Click **Install**, and then click **Close**.

Note: The installation may take several minutes.

23. In the Server Manager, click **Tools** and click **Internet Information Services (IIS) Manager**.
24. In the Internet Information Services (IIS) Manager, expand **LON-SVR1(ADATUM\Adminisrtator)\Sites\Default Web Site and click _wmcs**.
25. Under /_wmcs Home, In the Details pane, in the IIS section, double-click **Authentication**, click **Anonymous Authentication** and in the Actions pane click **Enable**.
26. In the Connections pane, expand **_wmcs** and click **licensing**.
27. Under **/_wmcs/licensing Home**, In the Details pane, in the IIS section, double-click **Authentication**, click **Anonymous Authentication**, and in the Actions pane, click **Enable**.
28. Click to the **Start** screen, click **Administrator**, and then click **Sign Out**.

Note: You must sign out before you can manage AD RMS.

▶ Task 3: Configure the AD RMS Super Users group

1. Sign in to LON-SVR1 with the **Adatum\Administrator** account and the password **Pa$$w0rd**.
2. In the Server Manager, click **Tools**, and then click **Active Directory Rights Management Services**.
3. In the Active Directory Rights Management Services console, expand the lon-svr1(Local) node, and then click **Security Policies**.
4. In the Security Policies area, under **Super Users**, click **Change super user settings**.
5. In the Actions pane, click **Enable Super Users**.
6. In the Super Users area, click **Change super user group**.
7. In the **Super Users** dialog box, in the **Super user group** text box, type **ADRMS_Superusers@adatum.com**, and then click **OK**.

Results: After completing this exercise, you will have installed and configured AD RMS.

Exercise 2: Configuring AD RMS Templates

▶ Task 1: Configure a new rights-policy template

1. Ensure that you are logged on to LON-SVR1.
2. In the **Active Directory Rights Management Services** console, click the **lon-svr1 (local)\Rights Policy Templates** node.
3. In the **Actions** pane, click **Create Distributed Rights Policy Template**.
4. In the Create Distributed Rights Policy Template Wizard, on the **Add Template Identification information** page, click **Add**.
5. On the **Add New Template Identification Information** page, enter the following information, and then click **Add**:
   - Language: **English (United States)**
   - Name: **ReadOnly**
   - Description: **Read only access. No copy or print**
6. Click **Next**.
7. On the **Add User Rights** page, click **Add**.
8. On the **Add User or Group** page, type **executives@adatum.com**, and then click **OK**.
9. When executives@adatum.com is selected, under **Rights**, click **View**. Verify that **Grant owner (author) full control right with no expiration** is selected, and then click **Next**.
10. On the **Specify Expiration Policy** page, choose the following settings, and then click **Next**:
    - Content Expiration: Expires after the following duration (days): **7**
    - Use license expiration: Expires after the following duration (days): **7**
11. On the **Specify Extended Policy** page, click **Require a new use license every time content is consumed (disable client-side caching)**, click **Next**, and then click **Finish**.

▶ Task 2: Configure the rights-policy template distribution

1. On LON-SVR1, on the taskbar, click the Windows PowerShell icon.
2. At the Windows PowerShell prompt, type the following command, and then press Enter:

```
New-Item c:\rmstemplates -ItemType Directory
```

3. At the Windows PowerShell prompt, type the following command, and then press Enter:

```
New-SmbShare -Name RMSTEMPLATES -Path c:\rmstemplates -FullAccess ADATUM\ADRMSSVC
```

4. At the Windows PowerShell prompt, type the following command, and then press Enter:

```
New-Item c:\docshare -ItemType Directory
```

5. At the Windows PowerShell prompt, type the following command, and then press Enter:

```
New-SmbShare -Name docshare -Path c:\docshare -FullAccess Everyone
```

6. To exit Windows PowerShell, type **exit**.
7. Switch to the Active Directory Rights Management Services console.

8. Click the **Rights Policy Templates** node, and in the Distributed Rights Policy Templates area, click **Change distributed rights policy templates file location**.
9. In the **Rights Policy Templates** dialog box, click **Enable export**.
10. In the **Specify Templates File Location (UNC)**, type **\\LON-SVR1\RMSTEMPLATES**, and then click **OK**.
11. On the taskbar, click the File Explorer icon.
12. Navigate to the C:\rmstemplates folder, and verify that **ReadOnly.xml** is present.
13. Close the File Explorer window.

▶ Task 3: Configure an exclusion policy

1. Switch to the Active Directory Rights Management Services console.
2. Click the **Exclusion Policies** node, and then click **Manage application exclusion list**.
3. In the **Actions** pane, click **Enable Application Exclusion**.
4. In the **Actions** pane, click **Exclude Application**.
5. In the **Exclude Application** dialog box, enter the following information, and then click **Finish**:
   - Application File name: **Powerpnt.exe**
   - Minimum version: **14.0.0.0**
   - Maximum version: **16.0.0.0**

Results: After completing this exercise, you will have configured AD RMS templates.

Exercise 3: Implementing the AD RMS Trust Policies

▶ Task 1: Export the Trusted User Domains policy

1. On LON-SVR1, on the taskbar, click the **Windows PowerShell** icon.
2. At the Windows PowerShell prompt, type the following command, and then press Enter:

```
New-Item c:\export -ItemType Directory
```

3. At the Windows PowerShell prompt, type the following command, and then press Enter:

```
New-SmbShare -Name Export -Path c:\export -FullAccess Everyone
```

4. Close the Windows PowerShell window.
5. In the Active Directory Rights Management Services console, expand the **Trust Policies** node, and then click the **Trusted User Domains** node.
6. In the Actions pane, click **Export Trusted User Domains**.
7. In the **Export Trusted User Domains As** dialog box, navigate to **\\LON-SVR1\export**, set the file name to **ADATUM-TUD.bin**, and then click **Save**.
8. Sign in to TREY-DC1 with the **TREYRESEARCH\Administrator** account and the password **Pa$$w0rd**.
9. In the Server Manager, click **Tools**, and then click **Active Directory Rights Management Services**.
10. In the Active Directory Rights Management Services console, expand **trey-dc1(local)**, expand the **Trust Policies** node, and then click the **Trusted User Domains** node.
11. In the Actions pane, click **Export Trusted User Domains**.
12. In the **Export Trusted User Domains As** dialog box, navigate to **\\LON-SVR1\export**, set the file name to **TREYRESEARCH-TUD.bin**, and then click **Save**.
13. On TREY-DC1, on the taskbar, click the **Windows PowerShell** icon.
14. At the Windows PowerShell prompt, type the following command, and then press Enter:

```
Add-DnsServerConditionalForwarderZone -MasterServers 172.16.0.10 -Name adatum.com
```

15. Close the Windows PowerShell window.

▶ Task 2: Export the Trusted Publishing Domains policy

1. Switch to LON-SVR1.
2. In the Active Directory Rights Management Services console, under the **Trust Policies** node, click the **Trusted Publishing Domains** node.
3. In the Actions pane, click **Export Trusted Publishing Domains**.
4. In the **Export Trusted Publishing Domain** dialog box, click **Save As**.
5. In the **Export Trusted Publishing Domain File As** dialog box, navigate to **\\LON-SVR1\export**, set the file name to **ADATUM-TPD.xml**, and then click **Save**.
6. In the **Export Trusted Publishing Domain** dialog box, enter the password **Pa$$w0rd** twice, and then click **Finish**.
7. Switch to TREY-DC1.
8. In the Active Directory Rights Management Services console, under the **Trust Policies** node, click the **Trusted Publishing Domains** node.

9. In the **Actions** pane, click **Export Trusted Publishing Domains**.
10. In the **Export Trusted Publishing Domain** dialog box, click **Save As**.
11. In the **Export Trusted Publishing Domain File As** dialog box, navigate to **\\LON-SVR1\export**, set the file name to **TREYRESEARCH-TPD.xml**, and then click **Save**.
12. In the **Export Trusted Publishing Domain** dialog box, enter the password **Pa$$w0rd** twice, and then click **Finish**.

▶ Task 3: Import the Trusted User Domain policy from the partner domain

1. Switch to LON-SVR1.
2. In the Active Directory Rights Management Services console, under the **Trust Policies** node, click the **Trusted User Domains** node.
3. In the Actions pane, click **Import Trusted User Domain**.
4. In the **Import Trusted User Domain** dialog box, enter the following details, and then click **Finish**:
   - Trusted user domain file: **\\LON-SVR1\Export\TREYRESEARCH-TUD.bin**
   - Display Name: **Trey Research**
5. Switch to TREY-DC1.
6. In the Active Directory Rights Management Services console, under the **Trust Policies** node, click the **Trusted User Domains** node.
7. In the Actions pane, click **Import Trusted User Domain**.
8. In the **Import Trusted User Domain** dialog box, enter the following details, and then click **Finish**:
   - Trusted user domain file: **\\LON-SVR1\Export\ADATUM-TUD.bin**
   - Display Name: **Adatum**

▶ Task 4: Import the Trusted Publishing Domains policy from the partner domain

1. Switch to LON-SVR1.
2. In the Active Directory Rights Management Services console, under the **Trust policies** node, click the **Trusted Publishing Domains** node.
3. In the **Actions** pane, click **Import Trusted Publishing Domain**.
4. In the **Import Trusted Publishing Domain** dialog box, enter the following information, and then click **Finish**:
   - Trusted publishing domain file: **\\LON-SVR1\export\ TREYRESEARCH-TPD.xml**
   - Password: **Pa$$w0rd**
   - Display Name: **Trey Research**
5. Switch to TREY-DC1.
6. In the Active Directory Rights Management Services console, under the **Trust policies** node, click the **Trusted Publishing Domains** node.
7. In the **Actions** pane, click **Import Trusted Publishing Domain**.

8. In the **Import Trusted Publishing Domain** dialog box, provide the following information, and then click **Finish**:
   - Trusted publishing domain file: **\\LON-SVR1\export\adatum-tpd.xml**
   - Password: **Pa$$w0rd**
   - Display Name: **Adatum**

Results: After completing this exercise, you will have implemented the AD RMS trust policies.

Exercise 4: Verifying the AD RMS Deployment

▶ Task 1: Create a rights-protected document

1. Sign on to **LON-CL1** as **Adatum\administrator** with a password of **Pa$$w0rd**.
2. On the **Start** screen, select the **Desktop**
3. Click the **File Explorer** icon
4. In **file Explorer**, right-click **This PC** and select **Properties**
5. In the **System** Window, select **Remote settings** in the console tree.
6. Select the **Select Users** button.
7. Click the **Add** button.
8. In the **Select Users and Groups**, popup, **Enter the object names to select** text box, type **Aidan;Bill:Carol** and then click **OK** three times.
9. On the taskbar, click the **Windows** start icon.
10. On the **Start** screen, click **Administrator** and then **Sign out**.
11. Sign in to LON-CL1 as **Adatum\Aidan** using the password **Pa$$w0rd**.
12. In the **Start** screen, select the **Desktop** tile.
13. Open **Internet Explorer**. Close any warnings about add-ons.
14. In the URL text box, type **http://adrms.adatum.com**, click the arrow immediately to the right of the url text box.
15. Click the **Gear** icon in the far upper right of **Internet Explorer**.
16. Select Internet Options.
17. Select the **Security** tab.
18. In the **Select a zone to view or change** security settings, click the **Local intranet** icon, and then click the **Sites** button.
19. Click the **Advanced** button.
20. Click the **Add** button, and then **Close** and then **OK** twice. Close **Internet Explorer**.
21. Return to the Start screen.
22. On the **Start** screen, type **Word**. In the **Results** area, click **Word 2013**.
23. In the Word Recent window, click the Blank document icon. In the Microsoft Word document, type the following text:

This document is for executives only, it should not be modified.

24. Click **File**, click **Protect Document**, click **Restrict Access**, and then click **Connect to Digital Rights Management Servers and get templates**.
25. A Microsoft Word dialog box informing you it is connecting to the server will appear.
26. After the dialog box disappears, Click Protect Document and **Restrict Access** and then click **Restricted Access**.
27. In the **Permission** dialog box, enable **Restrict Permission to this document**.
28. In the **Read** text box, type **bill@adatum.com**, and then click **OK**.
29. Click **Save**.

30. In the **Save As** dialog box, click the Browse icon, and in the file name: Text box, type the **\\lon-svr1\docshare \Executives Only.docx**. and then click **Save**.
31. Close Microsoft Word.
32. Click to the **Start** screen, click the **Aidan Delaney** icon, and then click **Sign out**.

▶ Task 2: Verify internal access to protected content

1. Sign in to LON-CL1 as **Adatum\Bill** using the password **Pa$$w0rd**.
2. In the **Start** screen, select the **Desktop** tile.
3. Open **Internet Explorer**. Close any warnings about add-ons.
4. In the URL text box, type **http://adrms.adatum.com**, click the arrow immediately to the right of the url text box.
5. Click the **Gear** icon in the far upper right of **Internet Explorer**.
6. Select Internet Options.
7. Select the **Security** tab.
8. In the **Select a zone to view or change** security settings, click the **Local intranet** icon, and then click the **Sites** button.
9. Click the **Advanced** button.
10. Click the **Add** button, and then **Close** and then **OK** twice. Close **Internet Explorer**.
11. Return to the Start screen.
12. On the **Start** screen, click **Desktop**.
13. On the taskbar, click the **File Explorer** icon.
14. In the File Explorer window, navigate to **\\lon-svr1\docshare**.
15. In the docshare folder, double-click the **Executives Only** document.
16. In the First things first dialog box, select the Ask me later radio button, and click Accept.
    - In the Office dialog box, click the letter X in the far upper right.
17. When the document opens, verify that you are unable to modify or save the document.
18. Select a line of text in the document.
19. Right-click the text, and verify that you cannot make changes.
20. Click **View Permission on the yellow bar**, review the permissions, and then click **OK**.
21. Close Microsoft Word.
22. Click to the Start screen, click the **Bill Malone** icon, and then click **Sign out**.

▶ Task 3: Open the rights-protected document as an unauthorized user

1. Sign in to LON-CL1 as **Adatum\Carol** using the password **Pa$$w0rd**.
2. In the **Start** screen, select the **Desktop** tile.
3. Open **Internet Explorer**. Close any warnings about add-ons.
4. In the URL text box, type **http://adrms.adatum.com**, click the arrow immediately to the right of the url text box.
5. Click the **Gear** icon in the far upper right of **Internet Explorer**.

6. Select Internet Options.
7. Select the **Security** tab.
8. In the **Select a zone to view or change** security settings, click the **Local intranet** icon, and then click the **Sites** button.
9. Click the **Advanced** button.
10. Click the **Add** button, and then **Close** and then **OK** twice. Close **Internet Explorer**.
11. Return to the Start screen.
12. On the Start menu, click **Desktop**.
13. On the taskbar, click the **File Explorer** icon.
14. In the Windows Explorer window, navigate to **lon-svr1\docshare**.
15. In the docshare folder, double-click the **Executives Only** document.
16. Verify that Carol is unable to open the document. You will receive a message with option to Change User or request access.
17. Click **No**.
18. Select **Ask me later** and click **Accept**. and then select the **X** in the far upper right of the Microsoft Office window.
19. Close Microsoft Word.
20. Click to the Start screen, click the **Carol Troup** icon, and then click **Sign out**.

▶ Task 4: Open and edit the rights-protected document as an authorized user at Trey Research

1. Sign in to LON-CL1 as **Adatum\Aidan** using the password **Pa$$w0rd**.
2. On the Start screen, type **Word**. In the **Results** area, click **Word 2013**. Click Blank document.
3. In the Microsoft Word document, type the following text:

This document is for Trey Research only, it should not be modified.

4. Click **File**, click **Protect Document**, click **Restrict Access**, and then click **Restricted Access**.
5. In the **Permission** dialog box, enable **Restrict Permission to this document**.
6. In the **Read** text box, enter **april@treyresearch.net**, click **OK**, and then click **Save, and then Browse**.
7. In the **Save As** dialog box, save the document to the **\\lon-svr1\docshare** location as **TreyResearch-Confidential.docx**. Close Word 2013.
8. Click to the **Start** screen, click the **Aidan Delaney** icon, and then click **Sign Out**.
9. Sign on to **Trey-CL1** as **TREYRESEARCH\administrator** with a password of **Pa$$w0rd**.
10. On the **Start** screen, select the **Desktop**
11. Click the **File Explorer** icon
12. In **file Explorer**, right-click **This PC** and select **Properties**
13. In the **System** Window, select **Remote settings** in the console tree.
14. Select the **Select Users** button.

15. Click the **Add** button.
16. In the **Select Users and Groups**, popup, **Enter the object names to select** text box, type **April,** and then click **OK** three times.
17. On the taskbar, click the **Windows** start icon.
18. On the **Start** screen, click **Administrator** and then **Sign out**.
19. Sign in to TREY-CL1 as **TREYRESEARCH\APRIL** with the password **Pa$$w0rd**.
20. In the **Start** screen, select the **Desktop** tile.
21. Open **Internet Explorer**. Close any warnings about add-ons.
22. In the URL text box, type **http://adrms.treyresearch.net**, click the arrow immediately to the right of the url text box.
23. Click the **Gear** icon in the far upper right of **Internet Explorer**.
24. Select Internet Options.
25. Select the **Security** tab.
26. In the **Select a zone to view or change** security settings, click the **Local intranet** icon, and then click the **Sites** button.
27. Click the **Advanced** button.
28. Click the **Add** button, and then **Close** and then **OK** twice. Close **Internet Explorer**.
29. Return to the Start screen.
30. On the Start screen, click **Desktop**.
31. On the taskbar, click the **File Explorer** icon.
32. In the File Explorer window, navigate to **lon-svr1****docshare**.
33. In the **Windows Security** dialog box, enter the following credentials, and then click **OK**:
    - Username: **Adatum\Administrator**
    - Password: **Pa$$w0rd**
34. Copy the file **TreyResearch-Confidential.docx** to the desktop.
35. In the Active Directory Rights Management Services popup, click OK.
36. When the document opens, verify that you are unable to modify or save the document.
37. Select a line of text in the document and verify.
38. Right-click the text, and verify that you cannot make changes.
39. Click **View Permission**, review the permissions, and then click **OK**.

▶ Task 5: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20412C-LON-DC1**, and then click **Revert**.

3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20412C-LON-SVR1**, **20412C-TREY-DC1**, **20412C-LON-CL1**, and **20412C-TREY-CL1**.

Results: After completing this exercise, you will have verified that the AD RMS deployment is successful.

Module 8: Implementing and Administering AD FS

Lab A: Implementing AD FS

Exercise 1: Installing and Configuring AD FS

▶ Task 1: Create a DNS record for AD FS

1. On LON-DC1, in the Server Manager, click **Tools**, and then click **DNS**.
2. In the 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 **adfs**.
5. In the **IP address** box, type **172.16.0.10**, and then click **Add Host**.
6. In the DNS window, click **OK**.
7. Click **Done**, and then close the DNS Manager.

▶ Task 2: Create a service account

1. On LON-DC1, open a Windows PowerShell prompt.
2. At the Windows PowerShell prompt, type **New-ADUser –Name adfsService** and press Enter.
3. Type **Set-ADAccountPassword adfsService** and press Enter.
4. At the Password prompt, press Enter.
5. At the second Password prompt, type **Pa$$w0rd** and press Enter.
6. At the Repeat Password prompt, type **Pa$$w0rd** and press Enter.
7. Type **Enable-ADAccount adfsService** and press Enter.
8. Close the Windows PowerShell prompt.

▶ Task 3: Install AD FS

1. On LON-DC1, in the Server Manager, click **Manager**, and then click **Add Roles and Features**.
2. In the Add Roles and Features Wizard, on the **Before you begin** page, click **Next**.
3. On the **Select installation type** page, click **Role-based or feature-based installation**, and then click **Next**.
4. On the **Select destination server** page, click **Select a server from the server pool**, click **LON-DC1.Adatum.com**, and then click **Next**.
5. On the **Select server roles** page, select the **Active Directory Federation Services** check box, and then click **Next**.
6. On the **Select features** page, click **Next**.
7. On the **Active Directory Federation Services (AD FS)** page, click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. When the installation is complete, click **Close**.

▶ Task 4: Configure AD FS

1. On LON-DC1, in the Server Manager, click the **Notifications** icon, and then click **Configure the federation service on this server**.
2. In the Active Directory Federation Services Configuration Wizard, on the **Welcome** page, click **Create the first federation server in a federation server farm**, and then click **Next**.
3. On the **Connect to Active Directory Domain Services** page, click **Next** to use **Adatum\Administrator** to perform the configuration.
4. On the **Specify Service Properties** page, in the **SSL Certificate** box, select **adfs.adatum.com**.
5. In the **Federation Service Display Name** box, type **A. Datum Corporation**, and then click **Next**.
6. On the **Specify Service Account** page, click **Use an existing domain user account or group Managed Service Account**.
7. Click **Select**, type **adfsService**, and click **OK**.
8. In the **Account Password** box, type **Pa$$w0rd**, and then click **Next**.
9. On the **Specify Configuration Database** page, click **Create a database on this server using Windows Internal Database**, and then click **Next**.
10. On the **Review Options** page, click **Next**.
11. On the **Pre-requisite Checks** page, click **Configure**.
12. On the **Results** page, click **Close**.

 **Note:** The adfs.adatum.com certificate was preconfigured for this task. In your own environment, you need to obtain this certificate.

▶ Task 5: Verify AD FS functionality

1. On LON-CL1, sign in as **Adatum\Brad** with the password **Pa$$w0rd**.
2. On the taskbar, click **Internet Explorer**.
3. In Internet Explorer, in the address bar, type **https://adfs.adatum.com/federationmetadata/2007-06/federationmetadata.xml**, and then press Enter.
4. Verify that the file loads, and then close Internet Explorer.

Results: In this exercise, you installed and configured AD FS. You also verified that it is functioning by viewing the FederationMetaData.xml file contents.

Exercise 2: Configuring an Internal Application for AD FS

▶ Task 1: Configure a certificate for the application

1. On LON-SVR1, in Server Manager, click **Tools** and click **Internet Information Services (IIS) Manager**.
2. If necessary, in the prompt for connecting to Microsoft Web Platform components, select the **Do not show this message** check box and then click **No**.
3. In IIS Manager, click **LON-SVR1 (ADATUM\Administrator**) and then double-click **Server Certificates**.
4. In the Actions pane, click **Create Domain Certificate**.
5. In the Create Certificate window on the Distinguished Name Properties page, enter the following and then click **Next**.
   - Common name: **lon-svr1.adatum.com**
   - Organization: **A. Datum**
   - Organizational unit: **IT**
   - City/locality: **London**
   - State/Province: **England**
   - Country/region: **GB**
6. On the Online Certification Authority page, click **Select**.
7. In the Select Certification Authority window, click **AdatumCA** and click **OK**.
8. On the Online Certification Authority page, in the **Friendly name** box, type **AdatumTestApp Certificate** and click **Finish**.
9. In IIS Manager, expand **LON-SVR1 (ADATUM\Administrator)**, expand **Sites**, click **Default Web Site**, and in the Actions Pane, click **Bindings**.
10. In the **Site Bindings** window, click **Add**.
11. In the **Add Site Binding** window, in the **Type** box, select **https**.
12. In the **SSL certificate** box, select **AdatumTestApp Certificate** and click **OK**.
13. In the Site Bindings window, click **Close**.
14. Close IIS Manager.

▶ Task 2: Configure the Active Directory claims-provider trust

1. On LON-DC1, in the Server Manager, click **Tools**, and then click **AD FS Management**.
2. In the AD FS management console, expand **Trust Relationships**, and then click **Claims Provider Trusts**.
3. In the middle pane, right-click **Active Directory**, and then click **Edit Claim Rules**.
4. In the Edit Claims Rules for Active Directory window, on the **Acceptance Transform Rules** tab, click **Add Rule**.
5. In the Add Transform Claim Rule Wizard, on the **Select Rule Template** page, in the **Claim rule template** box, select **Send LDAP Attributes as Claims**, and then click **Next**.
6. On the **Configure Rule** page, in the **Claim rule name** box, type **Outbound LDAP Attributes Rule**.

7. In the **Attribute Store** drop-down list, select **Active Directory**.
8. In the **Mapping of LDAP attributes to outgoing claim types** section, select the following values for the LDAP Attribute and the Outgoing Claim Type, and then click **Finish**:
   - E-Mail-Addresses: **E-Mail Address**
   - User-Principal-Name: **UPN**
   - Display-Name: **Name**
9. In the Edit Claim Rules for Active Directory window, click **OK**.

▶ Task 3: Configure the application to trust incoming claims

1. On LON-SVR1, in the Server Manager, click **Tools**, and then click **Windows Identity Foundation Federation Utility**.
2. On the **Welcome to the Federation Utility Wizard** page, in the **Application configuration location** box, type **C:\inetpub\wwwroot\AdatumTestApp\web.config** for the location of the sample web.config file.
3. In the **Application URI** box, type **https://lon-svr1.adatum.com/AdatumTestApp/** to indicate the path to the sample application that will trust the incoming claims from the federation server, and then click **Next** to continue.
4. On the **Security Token Service** page, click **Use an existing STS**, in the **STS WS-Federation metadata document location** box, type **https://adfs.adatum.com/federationmetadata/2007-06/federationmetadata.xml**, and then click **Next** to continue.
5. On the **STS signing certificate chain validation error** page, click **Disable certificate chain validation**, and then click **Next**.
6. On the **Security token encryption** page, click **No encryption**, and then click **Next**.
7. On the **Offered claims** page, review the claims that will be offered by the federation server, and then click **Next**.
8. On the **Summary** page, review the changes that will be made to the sample application by the Federation Utility Wizard, scroll through the items to understand what each item is doing, and then click **Finish**.
9. In the Success window, click **OK**.

▶ Task 4: Configure a relying-party trust for the claims-aware application

1. On LON-DC1, in the AD FS console, click **Relying Party Trusts**.
2. In the Actions pane, click **Add Relying Party Trust**.
3. In the Relying Party Trust Wizard, on the **Welcome** page, click **Start**.
4. On the **Select Data Source** page, click **Import data about the relying party published online or on a local network**.
5. In the **Federation Metadata address (host name or URL)** box, type **https://lon-svr1.adatum.com/adatumtestapp/**, and then click **Next**. This downloads the metadata configured in the previous task.
6. On the **Specify Display Name** page, in the **Display name** box, type **A. Datum Test App**, and then click **Next**.
7. On the **Configure Multi-factor Authentication Now** page, click **I do not want to configure multi-factor authentication settings for this relying party trust at this time**, and then click **Next**.

8. On the **Choose Issuance Authorization Rules** page, click **Permit all users to access this relying party**, and then click **Next**.
9. On the **Ready to Add Trust** page, review the relying-party trust settings, and then click **Next**.
10. On the **Finish** page, click **Close**.
11. Leave the Edit Claims Rules for A. Datum Test App window open for the next task.

▶ Task 5: Configure claim rules for the relying-party trust

1. On LON-DC1, in the AD FS management console, in the Edit Claim Rules for A. Datum Test App window, on the **Issuance Transform Rules** tab, click **Add Rule**.
2. In the **Claim rule template** box, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
3. In the **Claim rule name** box, type **Pass through Windows account name**.
4. In the **Incoming claim type** drop-down list, click **Windows account name**, and then click **Finish**.
5. On the **Issuance Transform Rules** tab, click **Add Rule**.
6. In the **Claim rule template** box, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
7. In the **Claim rule name** box, type **Pass through E-Mail Address**.
8. In the **Incoming claim type** drop-down list, click **E-Mail Address**, and then click **Finish**.
9. On the **Issuance Transform Rules** tab, click **Add Rule**.
10. In the **Claim rule template** box, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
11. In the **Claim rule name** box, type **Pass through UPN**.
12. In the **Incoming claim type** drop-down list, click **UPN**, and then click **Finish**.
13. On the **Issuance Transform Rules** tab, click **Add Rule**.
14. In the **Claim rule template** box, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
15. In the **Claim rule name** box, type **Pass through Name**.
16. In the **Incoming claim type** drop-down list, click **Name**, and then click **Finish**.
17. On the **Issuance Transform Rules** tab, click **OK**.

▶ Task 6: Test access to the claims-aware application

1. On LON-CL1, open Internet Explorer.
2. In Internet Explorer, in the address bar, type **https://lon- svr1.adatum.com/AdatumTestApp/**, and then press Enter.

Note: It is critical to use the trailing slash in the URL for step 2.

3. In the Windows Security window, sign in as **Adatum\Brad** with the password **Pa$$w0rd**.
4. Review the claim information that is displayed by the application.
5. Close Internet Explorer.

▶ Task 7: Configure Internet Explorer to pass local credentials to the application automatically

1. On LON-CL1, on the Start screen, type **Internet Options**, and then click **Internet Options**.
2. In the Internet Properties window, on the **Security** tab, click **Local intranet**, and then click **Sites**.
3. In the **Local intranet** window, click **Advanced**.
4. In the Local intranet window, in the **Add this website to the zone** box, type **https://adfs.adatum.com**, and then click **Add**.
5. In the **Add this website to the zone** box, type **https://lon-svr1.adatum.com**, click **Add**, and then click **Close**.
6. In the **Local intranet** window, click **OK**.
7. In the **Internet Properties** window, click **OK**.
8. On LON-CL1, open Internet Explorer.
9. In Internet Explorer, in the address bar, type **https://lon-svr1.adatum.com/AdatumTestApp/**, and then press Enter.

Note: It is critical to use the trailing slash in the URL for step 9.

10. Notice that you were not prompted for credentials.
11. Review the claim information that is displayed by the application.
12. Close Internet Explorer.

Results: After completing this exercise, you will have configured AD FS to support authentication for an application.

Lab B: Implementing AD FS for External Partners and Users

Exercise 1: Configuring AD FS for a Federated Business Partner

▶ Task 1: Configure DNS forwarding between TreyResearch.net and Adatum.com

1. On LON-DC1, in the Server Manager, click **Tools**, and then click **DNS**.
2. In the DNS Manager, expand **LON-DC1**, and then click **Conditional Forwarders**.
3. Right-click **Conditional Forwarders**, and then click **New Conditional Forwarder**.
4. In the New Conditional Forwarder window, in the **DNS Domain** box, type **TreyResearch.net**.
5. In the **IP addresses of the master servers** box, type **172.16.10.10**, and then press Enter.
6. Select the **Store this conditional forwarder in Active Directory, and replicate it as follows** check box, select **All DNS servers in this forest**, and then click **OK**.
7. Close the DNS Manager.
8. On TREY-DC1, in the Server Manager, click **Tools**, and then click **DNS**.
9. In the DNS Manager, expand **TREY-DC1**, and then click **Conditional Forwarders**.
10. Right-click **Conditional Forwarders**, and then click **New Conditional Forwarder**.
11. In the New Conditional Forwarder window, in the **DNS Domain** box, type **Adatum.com**.
12. In the **IP addresses of the master servers** box, type **172.16.0.10**, and then press Enter.
13. Select the **Store this conditional forwarder in Active Directory, and replicate it as follows** check box, select **All DNS servers in this forest**, and then click **OK**.
14. Close the DNS Manager.

Note: In a production environment, it is likely that you would use Internet DNS instead of conditional forwarders.

▶ Task 2: Configure certificate trusts between TreyResearch.net and Adatum.com

1. On LON-DC1, open File Explorer, browse to **\\TREY-DC1\CertEnroll**, and copy **TREY-DC1.TreyResearch.net_TreyResearchCA.crt** to **C:**.
2. Close File Explorer.
3. In the Server Manager, click **Tools**, and then click **Group Policy Management**.
4. In Group Policy Management, expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**, right-click **Default Domain Policy**, and then click **Edit**.
5. In Group Policy Management Editor, under **Computer Configuration**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, expand **Public Key Policies**, and then click **Trusted Root Certification Authorities**.
6. Right-click **Trusted Root Certification Authorities**, and then click **Import**.
7. In the Certificate Import Wizard, on the **Welcome to the Certificate Import Wizard** page, click **Next**.

8. On the **File to Import** page, type **C:\TREY-DC1.TreyResearch.net_TreyResearchCA.crt**, and then click **Next**.
9. On the **Certificate Store** page, click **Place all certificates in the following store**, select **Trusted Root Certification Authorities**, and then click **Next**.
10. On the **Completing the Certificate Import Wizard** page, click **Finish**, and then click **OK** to close the success message.
11. Close the Group Policy Management Editor.
12. Close Group Policy Management.
13. On TREY-DC1, open File Explorer, and then browse to **\\LON-DC1\CertEnroll**.
14. Right-click **LON-DC1.Adatum.com_AdatumCA.crt**, and then click **Install Certificate**.
15. In the Certificate Import Wizard, on the **Welcome to the Certificate Import Wizard** page, click **Local Machine**, and then click **Next**.
16. On the **Certificate Store** page, click **Place all certificates in the following store**, and then click **Browse**.
17. In the **Select Certificate Store** window, click **Trusted Root Certification Authorities**, and then click **OK**.
18. On the **Certificate Store** page, click **Next**.
19. On the **Completing the Certificate Import Wizard** page, click **Finish**, and then click **OK** to close the success message.
20. Close File Explorer.
21. On LON-SVR1, on the taskbar, click **Windows PowerShell**.
22. At the Windows PowerShell command prompt, type **gpupdate**, and then press Enter.
23. Close Windows PowerShell.
24. On LON-SVR2, on the taskbar, click **Windows PowerShell**.
25. At the Windows PowerShell command prompt, type **gpupdate**, and then press Enter.
26. Close Windows PowerShell.

Note: If you obtain certificates from a trusted certification authority, you do not need to configure a certificate trust between the organizations.

▶ Task 3: Create a DNS record for AD FS in TreyResearch.net

1. On TREY-DC1, in Server Manager, click **Tools**, and then click **DNS**.
2. In DNS Manager, expand **TREY-DC1**, expand **Forward Lookup Zones**, and then click **TreyResearch.net**.
3. Right-click **TreyResearch.net**, and then click **New Host (A or AAAA)**.
4. In the New Host window, in the **Name** box, type **adfs**.
5. In the **IP address** box, type **172.16.10.10**, and then click **Add Host**.
6. In the DNS window, click **OK**, and then click **Done**.
7. Close the DNS Manager.

▶ Task 4: Create a certificate for AD FS

1. On TREY-DC1, in Server Manager, click **Tools** and click **Internet Information Services (IIS) Manager**.
2. If necessary, in the prompt for connecting to Microsoft Web Platform components, select the **Do not show this message** check box and then click **No**.
3. In IIS Manager, click **TREY-DC1 (TREYRESEARCH\Administrator**) and then double-click **Server Certificates**.
4. In the Actions pane, click **Create Domain Certificate**.
5. In the Create Certificate window on the Distinguished Name Properties page, enter the following and then click **Next**.
   - Common name: **adfs.TreyResearch.net**
   - Organization: **Trey Research**
   - Organizational unit: **IT**
   - City/locality: **London**
   - State/Province: **England**
   - Country/region: **GB**
6. On the Online Certification Authority page, click **Select**.
7. In the Select Certification Authority window, click **TreyResearchCA** and click **OK**.
8. On the Online Certification Authority page, in the **Friendly name** box, type **adfs.TreyResearch.net** and click **Finish**.
9. Close IIS Manager.

▶ Task 5: Create a service account

1. On TREY-DC1, open a Windows PowerShell prompt.
2. At the Windows PowerShell prompt, type **New-ADUser –Name adfsService** and press Enter.
3. Type **Set-ADAccountPassword adfsService** and press Enter.
4. At the Password prompt, press Enter.
5. At the second Password prompt, type **Pa$$w0rd** and press Enter.
6. At the Repeat Password prompt, type **Pa$$w0rd** and press Enter.
7. Type **Enable-ADAccount adfsService** and press Enter.
8. Close the Windows PowerShell prompt.

▶ Task 6: Install AD FS for TreyResearch.net

1. On TREY-DC1, in the Server Manager, click **Manage**, and then click **Add Roles and Features**.
2. In the Add Roles and Features Wizard, on the **Before you begin** page, click **Next**.
3. On the **Select Installation type** page, click **Role-based or feature-based installation**, and then click **Next**.
4. On the **Select destination server** page, click **Select a server from the server pool**, click **TREY-DC1.TreyResearch.net**, and then click **Next**.

5. On the **Select server roles** page, select the **Active Directory Federation Services** check box, and then click **Next**.
6. On the **Select features** page, click **Next**.
7. On the **Active Directory Federation Services (AD FS)** page, click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. When the installation is complete, click **Close**.

▶ Task 7: Configure AD FS for TreyResearch.net

1. On TREY-DC1, in the Server Manager, click the **Notifications** icon, and then click **Configure the federation service on this server**.
2. In the Active Directory Federation Services Configuration Wizard, on the **Welcome** page, click **Create the first federation server in a federation server farm**, and then click **Next**.
3. On the **Connect to Active Directory Domain Services** page, click **Next** to use TREYRESEARCH\Administrator to perform the configuration.
4. On the **Specify Service Properties** page, in the **SSL Certificate** box, select **adfs.TreyResearch.net**.
5. In the **Federation Service Display Name** box, type **Trey Research**, and then click **Next**.
6. On the **Specify Service Account** page, click **Use an existing domain user account or group Managed Service Account**.
7. Click **Select**, type **adfsService**, and click **OK**.
8. In the **Account Password** box, type **Pa$$w0rd**, and then click **Next**.
9. On the **Specify Configuration Database** page, click **Create a database on this server using Windows Internal Database**, and then click **Next**.
10. On the **Review Options** page, click **Next**.
11. On the **Pre-requisite Checks** page, click **Configure**.
12. On the **Results** page, click **Close**.

▶ Task 8: Add a claims-provider trust for the TreyResearch.net AD FS server

1. On LON-DC1, in Server Manager, click **Tools**, and then click **AD FS Management**.
2. In the AD FS management console, expand **Trust Relationships**, and then click **Claims Provider Trusts**.
3. In the Actions pane, click **Add Claims Provider Trust**.
4. In the Add Claims Provider Trust Wizard, on the **Welcome** page, click **Start**.
5. On the **Select Data Source** page, click **Import data about the claims provider published online or on a local network**.
6. In the **Federation metadata address (host name or URL)** box, type **https://adfs.treyresearch.net**, and then click **Next**.
7. On the **Specify Display Name** page, in the **Display name** box, type **Trey Research**, and then click **Next**.
8. On the **Ready to Add Trust** page, review the claims-provider trust settings, and then click **Next** to save the configuration.

9. On the **Finish** page, select the **Open the Edit Claim Rules dialog for this claims provider trust when the wizard closes** check box, and then click **Close**.
10. In the Edit Claim Rules for Trey Research window, on the **Acceptance Transform Rules** tab, click **Add Rule**.
11. In the Add Transform Claim Rule Wizard, on the **Select Rule Template** page, in the **Claim rule template** box, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
12. On the **Configure Rule** page, in the **Claim rule name** box, type **Pass through Windows account name**.
13. In the **Incoming claim type** drop-down list, select **Windows account name**.
14. Select **Pass through all claim values**, and then click **Finish**.
15. In the pop-up window, click **Yes** to acknowledge the warning.
16. In the Edit Claim Rules for Trey Research window, click **OK**, and then close the AD FS management console.

▶ Task 9: Configure a relying party trust in TreyResearch.net for the Adatum.com application

1. On TREY-DC1, in the Server Manager, click **Tools**, and then click **AD FS Management**.
2. In the AD FS management console, expand **Trust Relationships**, and then click **Relying Party Trusts**.
3. In the Actions pane, click **Add Relying Party Trust**.
4. In the Add Relying Party Trust Wizard, on the **Welcome** page, click **Start**.
5. On the **Select Data Source** page, click **Import data about the relying party published online or on a local network**.
6. In the **Federation metadata address (host or URL)** box, type **adfs.adatum.com**, and then click **Next**.
7. On the **Specify Display Name** page, in the **Display name** text box, type **A. Datum Corporation**, and then click **Next**.
8. On the **Configure Multi-Factor Authentication Now** page, click **I do not want to configure multi-factor authentication settings for this relying party trust at this time**, and then click **Next**.
9. On the **Choose Issuance Authorization Rules** page, select **Permit all users to access this relying party**, and then click **Next**.
10. On the **Ready to Add Trust** page, review the relying-party trust settings, and then click **Next** to save the configuration.
11. On the **Finish** page, select the **Open the Edit Claim Rules dialog box for the relying party trust when the wizard closes** check box, and then click **Close**.
12. In the Edit Claim Rules for A. Datum Corporation window, on the **Issuance Transform Rules** tab, click **Add Rule**.
13. In the Add Transform Claim Rule Wizard, on the **Select Rule Template** page, in the **Claim rule template** box, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
14. On the **Configure Rule** page, in the **Claim rule name** box, type **Pass through Windows account name**.
15. In the **Incoming claim type** drop-down list, select **Windows account name**.

16. Click **Pass through all claim values**, click **Finish**, and then click **OK**.
17. Close the AD FS management console.

▶ Task 10: Test access to the application

1. On TREY-DC1, open Internet Explorer.
2. In Internet Explorer, in the address bar, type **https://lon-svr1.adatum.com/adatumtestapp/**, and then press Enter.
3. On the **A. Datum Corporation** page, click **Trey Research**.
4. In the **Windows Security** dialog box, sign in as **TreyResearch\April** with the password **Pa$$w0rd**.
5. After the application loads, close Internet Explorer.
6. Open Internet Explorer.
7. In Internet Explorer, in the address bar, type **https://lon-svr1.adatum.com/adatumtestapp/**, and then press Enter.
8. In the **Windows Security** dialog box, sign in as **TreyResearch\April** with the password **Pa$$w0rd**.
9. Close Internet Explorer.

 **Note:** You are not prompted for a home realm on the second access. Once users have selected a home realm and have been authenticated by a realm authority, they are issued an _LSRealm cookie by the relying-party's federation server. The default lifetime for the cookie is 30 days. Therefore, to sign in multiple times, you should delete that cookie after each logon attempt to return to a clean state.

▶ Task 11: Configure issuance authorization rules

1. On TREY-DC1, in the Server Manager, click **Tools**, and then click **AD FS Management**.
2. In the AD FS management console, expand **Trust Relationships**, and click **Relying Party Trusts**.
3. Right-click **A. Datum Corporation**, and click **Edit claim rules**.
4. In the Edit Claim Rules for A. Datum Corporation window, on the **Issuance Authorization Rules** tab, click **Permit Access to All Users** and click **Remove Rule**.
5. Click **Yes** to confirm deleting the claim rule.
6. Click **Add Rule**.
7. In the Add Issuance Authorization Claim Rules Wizard, on the Select Rule Template page, in the **Claim rule template** box, select **Permit or Deny Users Based on an Incoming Claim**, and then click **Next**.
8. On the Configure Rule page, in the **Claim rule name** box, type **Allow Production Members**.
9. In the **Incoming claim type** box, select **Group**.
10. In the **Incoming claim value** box, type **TreyResearch-Production**.
11. Click **Permit access to users with the incoming claim**, and click **Finish**.
12. In the Edit Claim Rules for A. Datum Corporation window, click **OK**.
13. In the AD FS management console, click **Claims Provider Trusts**, right-click **Active Directory**, and click **Edit Claim Rules**.
14. In the **Edit Claim Rules for Active Directory** window, click **Add Rule**.

15. In the Add Transform Claim Rule Wizard, on the Select Rule Template page, in the **Claim rule template** box, select **Send Group Membership as a Claim**, and click **Next**.
16. On the Configure Rule page, in the **Claim rule name** box, type **Production Group Claim**.
17. To set the **User's group**, click **Browse**, type **Production**, and click **OK**.
18. In the **Outgoing claim type** box, select **Group**.
19. In the **Outgoing claim value** box, type **TreyResearch-Production** and click **Finish**.
20. In the **Edit Claim Rules for Active Directory** window, click **OK**.
21. Close the AD FS management console.

▶ Task 12: Test the application of issuance authorization rules

1. On TREY-DC1, open Internet Explorer.
2. In Internet Explorer, in the address bar, type **https://lon-svr1.adatum.com/adatumtestapp/**, and then press Enter.
3. In the **Windows Security** dialog box, sign in as **TreyResearch\April** with the password **Pa$$w0rd**.
4. Verify that you cannot access the application because April is not a member of the production group.
5. Close Internet Explorer.
6. Open Internet Explorer.
7. In Internet Explorer, in the address bar, type **https://lon-svr1.adatum.com/adatumtestapp/**, and then press Enter.
8. In the **Windows Security** dialog box, sign in as **TreyResearch\Ben** with the password **Pa$$w0rd**.
9. Verify that you can access the application because Ben is a member of the production group.
10. Close Internet Explorer.

Results: After completing this exercise, you will have configured access for a claims-aware application in a partner organization.

Exercise 2: Configuring Web Application Proxy

▶ Task 1: Install Web Application Proxy

1. On LON-SVR2, in the Server Manager, click **Manage**, and then click **Add Roles and Features**.
2. In the Add Roles and Features Wizard, on the **Before you begin** page, click **Next**.
3. On the **Select installation type** page, click **Role-based or feature-based installation**, and then click **Next**.
4. On the **Select destination server** page, click **LON-SVR2.Adatum.com**, and then click **Next**.
5. On the **Select server roles** page, expand **Remote Access**, select the **Web Application Proxy** check box, and then click **Next**.
6. On the **Select features** page, click **Next**.
7. On the **Confirm installation selections** page, click **Install**.
8. On the **Installation progress** page, click **Close**.

▶ Task 2: Add the adfs.adatum.com certificate to LON-SVR2

1. On LON-DC1, on the Start screen, type **mmc**, and then press Enter.
2. In the **Microsoft Management Console**, click **File**, and then click **Add/Remove Snap-in**.
3. In the Add or Remove Snap-ins window, in the **Available snap-ins** column, double-click **Certificates**.
4. In the **Certificates snap-in** window, click **Computer account**, and then click **Next**.
5. In the **Select Computer** window, click **Local Computer (the computer this console is running on)**, and then click **Finish**.
6. In the **Add or remove Snap-ins** window, click **OK**.
7. In the **Microsoft Management Console**, expand **Certificates (Local Computer)**, expand **Personal**, and then click **Certificates**.
8. Right-click **adfs.adatum.com**, point to **All Tasks**, and then click **Export**.
9. In the Certificate Export Wizard, click **Next**.
10. On the **Export Private Key** page, click **Yes, export the private key**, and then click **Next**.
11. On the **Export File Format** page, click **Next**.
12. On the **Security** page, select the **Password** check box.
13. In the **Password** and **Confirm password** boxes, type **Pa$$w0rd**, and then click **Next**.
14. On the **File to Export** page, in the **File name** box, type **C:\adfs.pfx**, and then click **Next**.
15. On the **Completing the Certificate Export Wizard** page, click **Finish**, and then click **OK** to close the success message.
16. Close the Microsoft Management Console, and then do not save the changes.
17. On LON-SVR2, on the Start screen, type **mmc**, and then press Enter.
18. In the Microsoft Management Console, click **File**, and then click **Add/Remove Snap-in**.
19. In the Add or Remove Snap-ins window, in the **Available snap-ins** column, double-click **Certificates**.
20. In the **Certificates snap-in** window, click **Computer account**, and then click **Next**.

21. In the **Select Computer** window, click **Local Computer (the computer this console is running on)**, and then click **Finish**.
22. In the **Add or remove Snap-ins** window, click **OK**.
23. In the **Microsoft Management Console**, expand **Certificates (Local Computer)**, and then click **Personal**.
24. Right-click **Personal**, point to **All Tasks**, and then click **Import**.
25. In the **Certificate Import Wizard**, click **Next**.
26. On the File to Import page, in the **File name** box, type **\\LON-DC1\c$\adfs.pfx**, and then click **Next**.
27. On the Private key protection page, in the **Password** box, type **Pa$$w0rd**.
28. Select the **Mark this key as exportable** check box, and then click **Next**.
29. On the **Certificate Store** page, click **Place all certificates in the following store**.
30. In the **Certificate store** box, select **Personal**, and then click **Next**.
31. On the **Completing the Certificate Import Wizard** page, click **Finish**, and then click **OK** to clear the success message.
32. Close the Microsoft Management Console, and then do not save the changes.

▶ Task 3: Add the LON-SVR1.adatum.com certificate to LON-SVR2

1. On LON-SVR1, on the Start screen, type **mmc**, and then press Enter.
2. In the **Microsoft Management Console**, click **File**, and then click **Add/Remove Snap-in**.
3. In the Add or Remove Snap-ins window, in the **Available snap-ins** column, double-click **Certificates**.
4. In the **Certificates snap-in** window, click **Computer account**, and then click **Next**.
5. In the **Select Computer** window, click **Local Computer (the computer this console is running on)**, and then click **Finish**.
6. In the **Add or remove Snap-ins** window, click **OK**.
7. In the **Microsoft Management Console,** expand **Certificates (Local Computer)**, expand **Personal**, and then click **Certificates**.
8. Right-click **LON-SVR1.adatum.com**, point to **All Tasks**, and then click **Export**.
9. In the **Certificate Export Wizard**, click **Next**.
10. On the **Export Private Key** page, click **Yes, export the private key**, and then click **Next**.
11. On the **Export File Format** page, click **Next**.
12. On the **Security** page, select the **Password** check box.
13. In the **Password** and **Confirm password** boxes, type **Pa$$w0rd**, and then click **Next**.
14. On the **File to Export** page, in the **File name** box, type **C:\lon-svr1.pfx**, and then click **Next**.
15. On the **Completing the Certificate Export Wizard** page, click **Finish**, and then click **OK** to close the success message.
16. Close the Microsoft Management Console, and then do not save the changes.
17. On LON-SVR2, on the Start screen, type **mmc**, and then press Enter.
18. In the Microsoft Management Console, click **File**, and then click **Add/Remove Snap-in**.
19. In the Add or Remove Snap-ins window, in the **Available snap-ins** column, double-click **Certificates**.

20. In the **Certificates snap-in** window, click **Computer account**, and then click **Next**.
21. In the Select Computer window, click **Local Computer (the computer this console is running on)**, and then click **Finish**.
22. In the Add or remove Snap-ins window, click **OK**.
23. In the **Microsoft Management Console**, expand **Certificates (Local Computer)**, and then click **Personal**.
24. Right-click **Personal**, point to **All Tasks**, and then click **Import**.
25. In the **Certificate Import Wizard**, click **Next**.
26. On the **File to Import** page, in the **File name** box, type **\\LON-SVR1\c$\lon-svr1.pfx**, and then click **Next**.
27. On the **Private key protection** page, in the **Password** box, type **Pa$$w0rd**.
28. Select the **Mark this key as exportable** check box, and then click **Next**.
29. On the **Certificate Store** page, click **Place all certificates in the following store**.
30. In the **Certificate store** box, select **Personal**, and then click **Next**.
31. On the **Completing the Certificate Import Wizard** page, click **Finish**, and then click **OK** to clear the success message.
32. Close the Microsoft Management Console, and then do not save the changes.

▶ Task 4: Configure Web Application Proxy

1. In the Server Manager, click the **Notifications** icon, and then click **Open the Web Application Proxy Wizard**.
2. In the Web Application Proxy Wizard, on the **Welcome** page, click **Next**.
3. On the **Federation Server** page, enter the following, and then click **Next**:
   - Federation service name: **adfs.adatum.com**
   - User name: **Adatum\Administrator**
   - Password: **Pa$$w0rd**
4. On the **AD FS Proxy Certificate** page, in the **Select a certificate to be used by the AD FS proxy** box, select **adfs.adatum.com**, and then click **Next**.
5. On the **Confirmation** page, click **Configure**.
6. On the **Results** page, click **Close**.
7. The Remote Access Management Console opens automatically. Leave it open for the next task.

▶ Task 5: Configure the test application in Web Application Proxy

1. On LON-SVR2, in the **Remote Access Management Console**, click **Web Application Proxy**.
2. In the **Tasks** pane, click **Publish**.
3. In the Publish New Application Wizard, on the **Welcome** page, click **Next**.
4. On the **Preauthentication** page, click **Pass-through**, and then click **Next**.
5. On the **Publishing Settings** page, in the **Name** box, type **A. Datum Test App**.
6. In the **External URL** box, type **https://lon-svr1.adatum.com/adatumtestapp/**.
7. In the **External certificate** box, select **lon-svr1.adatum.com**.

8. In the **Backend server URL** box, type **https://lon-svr1.adatum.com/adatumtestapp/**, and then click **Next**.
9. On the **Confirmation** page, click **Publish**.
10. On the **Results** page, click **Close**.

▶ Task 6: Test Web Application Proxy

1. On TREY-DC1, on Start screen, type **Notepad**.
2. Right-click **Notepad**, and then click **Run as administrator**.
3. In Notepad, click **File**, and then click **Open**.
4. In the **File name** box, type **C:\Windows\System32\Drivers\etc\hosts**, and then click **Open**.
5. At the bottom of the file, add the following two lines, click **File**, and then click **Save**:
   - **172.16.0.22 adfs.adatum.com**
   - **172.16.0.22 lon-svr1.adatum.com**
6. Close Notepad.
7. Open Internet Explorer.
8. In Internet Explorer, in the address bar, type **https://lon-svr1.adatum.com/adatumtestapp/**, and then press Enter.
9. In the **Windows Security** dialog box, sign in as **TreyResearch\Ben** with password **Pa$$w0rd**.
10. After the application loads, close Internet Explorer.

Note: You edit the hosts to force TREY-DC1 to access the application through Web Application Proxy. In a production environment, you would do this by using split DNS.

▶ Task 7: To prepare for the next module

1. When you finish the lab, revert the virtual machines to their initial state. To do this, perform the following steps:
2. On the host computer, start the Hyper-V Manager.
3. In the **Virtual Machines** list, right-click **20412C-LON-DC1**, and then click **Revert**.
4. In the **Revert Virtual Machine** dialog box, click **Revert**.
5. Repeat steps 2 and 3 to revert **20412C-LON-SVR1, 20412C-LON-SVR2**, **20412C-LON-CL1** and **20412C-TREY-DC1**.

Results: After completing this exercise, you will have configured Web Application Proxy to secure access to AdatumTestApp from the Internet.

Module 9: Implementing Network Load Balancing

Lab: Implementing NLB

Exercise 1: Implementing an NLB Cluster

▶ Task 1: Verify website functionality for stand-alone servers

1. On LON-SVR1, on the taskbar, click the **File Explorer** icon.
2. Navigate to the folder **c:\inetpub\wwwroot**.
3. Double-click the file **iis-8.png**. This will open the file in Microsoft Paint.
4. Ensure that the **Paintbrush** tool is selected, and then in the palette, click the color **Red**.
5. Use the mouse to mark the IIS logo distinctively, using the color red.
6. Save the changes that you made to **iis-8.png**, and then close Microsoft Paint.
7. Close File Explorer.
8. Switch to LON-DC1.
9. Click Start.
10. On the Start screen click the **Windows Internet Explorer** icon.
11. In the Internet Explorer address bar, type the address **http://LON-SVR1**, and then press Enter. Verify that the webpage displays the IIS logo with the distinctive color red mark that you added.
12. In the Internet Explorer address bar, enter the address **http://LON-SVR2**, and then press Enter. Verify that the webpage does not display the marked IIS logo.
13. Close Internet Explorer.

▶ Task 2: Install NLB

1. Switch to LON-SVR1.
2. On the desktop, click the **Server Manager** icon.
3. In the Server Manager console, click the **Tools** menu, and then click **Windows PowerShell ISE**.
4. In the Windows PowerShell ISE window, enter the following command, and then press Enter:

```
Invoke-Command -Computername LON-SVR1,LON-SVR2 -command {Install-WindowsFeature
NLB,RSAT-NLB}
```

▶ Task 3: Create a new Windows Server 2012 NLB cluster

1. On LON-SVR1, in the Windows PowerShell ISE window, type the following command, and then press Enter:

```
New-NlbCluster -InterfaceName "Ethernet" -OperationMode Multicast -ClusterPrimaryIP
172.16.0.42 -ClusterName LON-NLB
```

2. In the Windows PowerShell ISE window, type the following command, and then press Enter:

```
Invoke-Command -Computername LON-DC1 -command {Add-DNSServerResourceRecordA      –
zonename adatum.com –name LON-NLB –Ipv4Address 172.16.0.42}
```

▶ Task 4: Add a second host to the cluster

1. On LON-SVR1, in the Windows PowerShell ISE window, type the following command and then press Enter:

```
Add-NlbClusterNode -InterfaceName "Ethernet" -NewNodeName "LON-SVR2" -
NewNodeInterface "Ethernet"
```

▶ Task 5: Validate the NLB cluster

1. On LON-SVR1, in the Server Manager console, click the **Tools** menu, and then click **Network Load Balancing Manager**.
2. In the Network Load Balancing Manager console, verify that nodes LON-SVR1 and LON-SVR2 display with the status of **Converged** for the LON-NLB cluster.
3. Right-click the **LON-NLB** cluster, and then click **Cluster properties**.
4. In the **LON-NLB(172.16.0.42)**, on the **Cluster Parameters** tab, verify that the cluster is set to use the **Multicast** operations mode.
5. On the **Port Rules** tab, verify that there is a single port rule named **All** that starts at port **0** and ends at port **65535** for both **TCP** and **UDP** protocols, and that it uses **Single** affinity.
6. Click **OK** to close the **LON-NLB(172.16.0.42)**.

Results: After completing this exercise, you will have successfully implemented an NLB cluster.

Exercise 2: Configuring and Managing the NLB Cluster

▶ Task 1: Configure port rules and affinity

1. On LON-SVR2, on the taskbar, click the **Windows PowerShell** icon.
2. At the Windows PowerShell prompt, type each of the following commands, and then press Enter after each command:

```
Cmd.exe
Mkdir c:\porttest
Xcopy /s c:\inetpub\wwwroot c:\porttest
Exit
New-Website –Name PortTest –PhysicalPath “C:\porttest” –Port 5678
New-NetFirewallRule –DisplayName PortTest –Protocol TCP –LocalPort 5678
```

3. On the taskbar, click the **File Explorer** icon.
4. Click drive **C**, double-click the **porttest** folder, and then double-click **iis-8.png**. This will open the file in Microsoft Paint.
5. Select the color blue from the palette, and use the **Blue** paintbrush to mark the IIS logo in a distinctive manner.
6. Save the changes to **iis-8.png**, and then close Microsoft Paint.
7. Switch to LON-DC1.
8. Click Start.
9. On the Start screen, click the **Internet Explorer** icon.
10. In the Internet Explorer address bar, type **http://LON-SVR2:5678**, and then press Enter.
11. Verify that the IIS Start page with the IIS logo distinctively marked with blue displays.
12. Switch to LON-SVR1.
13. On LON-SVR1, switch to **Network Load Balancing Manager**.
14. In the Network Load Balancing Manager console, right-click **LON-NLB**, and then click **Cluster Properties**.
15. In the **LON-NLB(172.16.0.42)**, on the **Port Rules** tab, select the **All** port rule, and then click **Remove**.
16. On the **Port Rules** tab, click **Add**.
17. In the **Add/Edit Port Rule** dialog box, enter the following information, and then click **OK**:
    - Port range: **80 to 80**
    - Protocols: **Both**
    - Filtering mode: **Multiple Host**
    - Affinity: **None**
18. On the **Port Rules** tab, click **Add**.
19. In the **Add/Edit Port Rule** dialog box, enter the following information, and then click **OK**:
    - Port range: **5678 to 5678**
    - Protocols: **Both**
    - Filtering mode: **Single Host**

20. Click OK to close the **LON-NLB(172.16.0.42)**.
21. In the Network Load Balancing Manager console, right-click **LON-SVR1**, and then click **Host Properties**.
22. On the **Port Rules** tab, click the port rule that has **5678** as the Start and End value, and then click **Edit**.
23. Click the **Handling priority** value, and change it to **10**.
24. Click **OK** twice to close both the **Add/Edit Port Rule** dialog box and the **Host Properties** dialog box.

▶ Task 2: Validate port rules

1. Switch to LON-DC1.
2. Click Start.
3. On the Start screen, click the **Internet Explorer** icon.
4. In the Internet Explorer address bar, type **http://lon-nlb**, and then press Enter.
5. Click the **Refresh** icon 20 times. Verify that you see web pages with and without the distinctive red marking.
6. On LON-DC1, verify that you have Internet Explorer open.
7. In the address bar, enter the address **http://LON-NLB:5678**, and then press Enter.
8. In the address bar, click the **Refresh** icon 20 times. Verify that you are able to view only the web page with the distinctive blue marking.

▶ Task 3: Manage host availability in the NLB cluster

1. Switch to LON-SVR1.
2. Select the Network Load Balancing Manager console.
3. Right-click **LON-SVR1**, click **Control Host**, and then click **Suspend**.
4. Click the **LON-NLB** node. Verify that node LON-SVR1 displays as **Suspended**, and that node LON-SVR2 displays as **Converged**.
5. Right-click LON-SVR1, click **Control Host**, and then click **Resume**.
6. Right-click LON-SVR1, click **Control Host**, and then click **Start**.
7. Click the LON-NLB node. Verify that both nodes LON-SVR1 and LON-SVR2 now display as **Converged**. You may have to refresh the view.

Results: After completing this exercise, you will have successfully configured and managed an NLB cluster.

Exercise 3: Validating High Availability for the NLB Cluster

▶ Task 1: Validate website availability when the host is unavailable

1. On LON-SVR1, on the taskbar, click the **Windows PowerShell** icon.
2. Type the following command, and then press Enter:

```
restart-computer
```

3. Switch to LON-DC1.
4. On LON-DC1, on the desktop, click the **Internet Explorer** icon.
5. In the Internet Explorer address bar, type the address **http://LON-NLB**, and then press Enter.
6. Refresh the website 20 times. Verify that the website is available while LON-SVR1 reboots, but that it does not display the distinctive red mark on the IIS logo until LON-SVR1 has restarted.

▶ Task 2: Configure and validate Drainstop

1. Sign in to LON-SVR1 with the username **Adatum\Administrator** and the password **Pa$$word**.
2. On the desktop, click the **Server Manager** icon.
3. In Server Manager, click the **Tools** menu, and then click **Network Load Balancing Manager**.
4. In the Network Load Balancing Manager console, right-click **LON-SVR2**, click **Control Host**, and then click **Drainstop**.
5. Switch to LON-DC1.
6. In Internet Explorer, in the address bar, type **http://lon-nlb**, and then press Enter.
7. Refresh the site 20 times, and verify that only the welcome page with the red IIS logo displays.

▶ Task 3: Prepare for the next module

When you finish the lab, revert the virtual machines to their initial state.

1. On the host computer, start the Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20412C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps two and three for 20412C-LON-SVR1 and 20412C-LON-SVR2.

Results: After completing this exercise, you will have successfully validated high availability for the NLB cluster.

Module 10: Implementing Failover Clustering

Lab: Implementing Failover Clustering

Exercise 1: Configuring a Failover Cluster

▶ Task 1: Connect cluster nodes to the iSCSI targets

1. On LON-SVR3, in the Server Manager, click **Tools**, and then click the **iSCSI Initiator**.
2. In the **Microsoft iSCSI** dialog box, click **Yes**.
3. Click the **Discovery** tab.
4. Click **Discover Portal**.
5. In the **IP address or DNS name** box, type **172.16.0.21**, and then click **OK**.
6. Click the **Targets** tab.
7. Click **Refresh**.
8. In the **Targets** list, select **iqn.1991-05.com.microsoft:lon-svr1-target1-target**, and then click **Connect**.
9. Select **Add this connection to the list of Favorite Targets**, and then click **OK** two times.
10. On LON-SVR4, in the Server Manager, click **Tools**, and then click **iSCSI Initiator**.
11. In the **Microsoft iSCSI** dialog box, click **Yes**.
12. Click the **Discovery** tab.
13. Click **Discover Portal**.
14. In the **IP address or DNS name** box, type **172.16.0.21**, and then click **OK**.
15. Click the **Targets** tab.
16. Click **Refresh**.
17. In the **Targets** list, select **iqn.1991-05.com.microsoft:lon-svr1-target1-target**, and then click **Connect**.
18. Select **Add this connection to the list of Favorite Targets**, and then click **OK** two times.
19. On LON-SVR3, in the **Server Manager**, click **Tools**, and then click **Computer Management**.
20. Expand **Storage**, and then click **Disk Management**.
21. Right-click **Disk 1**, and then click **Online**.
22. Right-click **Disk 1**, and then click **Initialize disk**. In the **Initialize Disk** dialog box, click **OK**.
23. Right-click the unallocated space next to **Disk 1**, and then click **New Simple Volume**.
24. On the **Welcome** page, click **Next**.
25. On the **Specify Volume Size** page, click **Next**.
26. On the **Assign Drive Letter or Path** page, click **Next**.
27. On the **Format Partition** page, in the **Volume Label** box, type **Data**. Select the **Perform a quick format** check box, and then click **Next**.

28. Click **Finish**.

Note: If the Microsoft Windows window pops up with a prompt to format the disk, click **Cancel**.

29. Repeat steps 21 through 28 for **Disk 2 and Disk 3**.

Note: Use Data2 and Data3 for Volume Labels.

30. Close the **Computer Management** window.
31. On LON-SVR4, in the **Server Manager**, click **Tools**, and then click **Computer Management**.
32. Expand **Storage**, and then click **Disk Management**.
33. Select and then right-click **Disk Management**, and then click **Refresh**.
34. Right-click **Disk 1**, and then click **Online**.
35. Right-click **Disk 2**, and then click **Online**.
36. Right-click **Disk 3**, and then click **Online**.
37. Close the **Computer Management** window.

▶ Task 2: Install the failover clustering feature

1. On LON-SVR3, if it is not opened, click the **Server Manager** icon to open Server Manager.
2. Click **Add roles and features**.
3. On the **Before You Begin** page, click **Next**.
4. On the **Select installation type** page, click **Next**.
5. On the **Select destination server** page, make sure that **Select server from the server pool** is selected, and then click **Next**.
6. On the **Select server roles** page, click **Next**.
7. On the **Select features** page, in the **Features** list, click **Failover Clustering**. In the **Add features that are required for Failover Clustering?** window, click **Add Features**. Click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. When installation is complete (you receive the message Installation succeeded on LON-SVR3), click **Close**.
10. Repeat steps one through nine on LON-SVR4.

▶ Task 3: Validate the servers for failover clustering

1. On LON-SVR3, in the Server Manager, click **Tools**, and then click **Failover Cluster Manager**.
2. In the **Actions** pane of the Failover Cluster Manager, click **Validate Configuration**.
3. In the **Validate a Configuration Wizard**, click **Next**.
4. In the **Enter Name** box, type **LON-SVR3**, and then click **Add**.
5. In the **Enter Name** box, type **LON-SVR4**.
6. Click **Add**, and then click **Next**.

7. Verify that **Run all tests (recommended)** is selected, and then click **Next**.
8. On the **Confirmation** page, click **Next**.
9. Wait for the validation tests to finish (it might take up to five minutes), and then on the **Summary** page, click **View Report**.
10. Verify that all tests completed without errors. Some warnings are expected.
11. Close **Internet Explorer**.
12. On the **Summary** page, remove the check mark next to **Create the cluster now using the validated nodes,** and click **Finish**.

▶ Task 4: Create the failover cluster

1. On LON-SVR3, in the Failover Cluster Manager, in the center pane, under **Management**, click **Create Cluster**.
2. On the **Before You Begin** page of the Create Cluster Wizard, read the information.
3. Click **Next**, in the **Enter server name** box, type **LON-SVR3**, and then click **Add**. Type **LON-SVR4**, and then click **Add**.
4. Verify the entries, and then click **Next**.
5. In **Access Point for Administering the Cluster**, in the **Cluster Name** box, type **Cluster1**.
6. Under **Address**, type **172.16.0.125**, and then click **Next**.
7. In the **Confirmation** dialog box, verify the information, and then click **Next**.
8. On the **Summary** page, click **Finish** to return to the Failover Cluster Manager.

▶ Task 5: Configuring CSV

1. On LON-SVR3, in the Failover Cluster Manager console, expand **cluster1.Adatum.com**, expand **Storage**, and then click **Disk**.
2. In the right pane, locate a disk that is assigned to **Available Storage** (you can see this in the **Assigned To** column). Right-click that disk, and then click **Add to Cluster Shared Volumes**. If possible, use Cluster Disk 2.
3. Ensure that the disk is assigned to **Cluster Shared Volume**.

Results: After this exercise, you will have installed and configured the failover clustering feature.

Exercise 2: Deploying and Configuring a Highly Available File Server

▶ Task 1: Add the File Server application to the failover cluster

1. On LON-SVR4, in the **Server Manager**, click **Dashboard**, and then click **Add roles and features**.
2. On the **Before You Begin** page, click **Next**.
3. On the **Select installation type** page, click **Next**.
4. On the **Select destination server** page, click **Next**.
5. On the **Select server roles** page, expand **File and Storage Services (1 of 12 installed)**, expand **File and iSCSI services**, and select **File Server**.
6. Click **Next** two times.
7. On the **Confirmation** page, click **Install**.
8. When the **installation succeeded** message appears, click **Close**.
9. On LON-SVR3, in the Failover Cluster Manager, expand **Cluster1.adatum.com**.
10. Expand **Storage**, and click **Disks**.
11. Make sure that three disks are present and online (with the names Cluster Disk 1, Cluster Disk 2, and Cluster Disk 3).
12. Right-click **Roles**, and then select **Configure Role**.
13. On the **Before You Begin** page, click **Next**.
14. On the **Select Role** page, select **File Server**, and then click **Next**.
15. On the **File Server Type** page, click **File Server for general use**, and then click **Next**.
16. On the **Client Access Point** page, in the **Name** box, type **AdatumFS**, and in the **Address** box, type **172.16.0.130**, and then click **Next**.
17. On the **Select Storage** page, select the **Cluster Disk 3** check box, and then click **Next**.
18. On the **Confirmation** page, click **Next**.
19. On the **Summary** page, click **Finish**.

▶ Task 2: Add a shared folder to a highly available file server

1. On LON-SVR4, in the Server Manager console, click **Tools**, and open **Failover Cluster Manager**.
2. Expand **Cluster1.Adatum.com**, and then click **Roles**.
3. Right-click **AdatumFS**, and then select **Add File Share**.
4. In the New Share Wizard, on the **Select the profile for this share** page, click **SMB Share – Quick**, and then click **Next**.
5. On the **Select the server and the path for this share** page, click **Next**.
6. On the **Specify share name** page, in the **Share name** box, type **Docs**, and then click **Next**.
7. On the **Configure share settings** page, review available options, do not make any changes, and then click **Next**.
8. On the **Specify permissions to control access** page, click **Next**.
9. On the **Confirm selections** page, click **Create**.
10. On the View results page, click **Close**.

▶ Task 3: Configure failover and failback settings

1. On LON-SVR4, in the Failover Cluster Manager, click **Roles**, right-click **AdatumFS**, and then click **Properties**.
2. Click the **Failover** tab and then click **Allow failback**.
3. Click **Failback between**, and set values to **4 and 5 hours**.
4. Click the **General** tab.
5. Select both **LON-SVR3** and **LON-SVR4** as preferred owners.
6. Move **LON-SVR4** up.
7. Click **OK**.

Results: After this exercise, you will have configured a highly available file server.

Exercise 3: Validate the Deployment of the Highly Available File Server

▶ Task 1: Validate the highly available file server deployment

1. On LON-DC1, open File Explorer, and in the Address bar, type **\\AdatumFS**, and then press Enter.
2. Verify that you can access the location and that you can open the **Docs** folder. Create a test text document inside this folder.
3. On LON-SVR3, open the Failover Cluster Manager.
4. Expand **Cluster1.adatum.com**, and then click **Roles**. Note the current owner of AdatumFS.

Note: You can view the owner in the Owner node column. It will be either LON-SVR3 or LON-SVR4.

5. Right-click **AdatumFS**, and then click **Move**, and then click **Select Node**.
6. In the **Move Clustered Role** dialog box, select cluster node (it will be either LON-SVR3 or LON-SVR4), and click **OK**.
7. Verify that **AdatumFS** has moved to a new owner.
8. Switch to the LON-DC1 computer, and verify that you can still access the **\\AdatumFS** location.

▶ Task 2: Validate the failover and quorum configuration for the file server role

1. On LON-SVR3, in the Failover Cluster Manager, click **Roles**.
2. Verify the current owner for the AdatumFS role.

Note: You can view the owner in the Owner node column. It will be either LON-SVR3 or LON-SVR4.

3. Click **Nodes**, and then select the node that is the current owner of the AdatumFS role.
4. Right-click the node, select **More Actions**, and then click **Stop Cluster Service**.
5. Verify that **AdatumFS** has moved to another node. To do this, click **Roles** and verify that AdatumFS is running.
6. Switch to the LON-DC1 computer, and verify that you can still access the **\\AdatumFS** location.
7. Switch to the LON-SVR3 computer, and on the Failover Cluster Manager, click **Nodes**. Right-click the stopped node, select **More Actions**, and then click **Start Cluster Service**.
8. Expand **Storage**, and then click **Disks**. In the center pane, right-click the disk that is assigned to **Disk Witness in Quorum** (Note: you can view this in the **Assigned to** column.)
9. Click **Take Offline**, and then click **Yes**.
10. Switch to LON-DC1 and verify that you can still access the **\\AdatumFS** location. By doing this, you verified that the cluster is still running even if the witness disk is offline.
11. Switch to the LON-SVR3 computer, and in Failover Cluster Manager, expand **Storage**, click **Disks**, right-click the disk that is in **Offline** status, and then click **Bring Online**.
12. Right-click **Cluster1.Adatum.com**, select **More Actions**, and then click **Configure Cluster Quorum Settings**...
13. On the **Before You Begin** page, click **Next**.

14. On the **Select Quorum Configuration Option** page, click **Advanced quorum configuration**, and then click **Next**.
15. On the **Select Voting Configuration** page, review available settings. Notice that you can select node or nodes that will or will not have votes in the cluster. Do not make any changes, and click **Next**.
16. On the **Select Quorum Witness** page, make sure that **Configure a disk witness** is selected, and click **Next**.
17. On the **Configure Storage Witness** page, select **Cluster Disk 3**, and click **Next**.
18. On the **Confirmation** page, click **Next**.
19. On the **Summary** page, click **Finish**.

Results: After this exercise, you will have tested the failover scenarios.

Exercise 4: Configuring CAU on the Failover Cluster

▶ Task 1: Configure CAU

1. On LON-DC1, in the Server Manager, click **Add roles and features**.
2. In the **Add Roles and Features** Wizard, on the **Before You Begin** page, click **Next**.
3. On the **Select installation type** page, click **Next**.
4. On the **Select destination server** page, make sure that **Select server from the server pool** is selected, and then click **Next**.
5. On the **Select server roles** page, click **Next**.
6. On the **Select features** page, in the list of features, click **Failover Clustering**. In **Add features that are required for Failover Clustering?** dialog box, click **Add Features**. Click **Next**.
7. On the **Confirm installation selections** page, click **Install**.
8. When installation is complete, click **Close**.
9. Switch to LON-SVR3. Open **Server Manager**, click **Tools**, and then click **Windows Firewall with Advanced Security**.
10. In the Windows Firewall with Advanced Security window, click **Inbound Rules**.
11. In the rules list, find the rule **Inbound Rule for Remote Shutdown (RPC-EP-In).** Verify that rule is enabled. If it is not enabled, right click the rule, and select **Enable Rule**.
12. In the rules list, find the rule **Inbound Rule for Remote Shutdown (TCP-In).** Verify that rule is enabled. If it is not enabled, right click the rule and select **Enable Rule**.
13. Close the Windows Firewall with Advanced Security window.
14. Switch to LON-SVR4, and repeat steps nine through 13.
15. On LON-DC1, in the **Server Manager** dashboard, click **Tools**, and then click **Cluster-Aware Updating**.
16. In the Cluster-Aware Updating window, in the **Connect to a failover cluster** drop-down list box, select **Cluster1**. Click **Connect**.
17. In the Cluster Actions pane, click **Preview updates for this cluster**.
18. In the **Cluster1-Preview Updates** window, click **Generate Update Preview List**. After several minutes, updates will be shown in the list. Review updates, and then click **Close**.

▶ Task 2: Update the failover cluster and configure self-updating

1. On LON-DC1, in the Cluster-Aware Updating console, click **Apply updates to this cluster**.
2. On the **Getting Started** page, click **Next**.
3. On the **Advanced options** page, review the options for updating, and then click **Next**.
4. On the **Additional Update Options** page, click **Next**.
5. On the **Confirmation** page, click **Update**, and then click **Close**.
6. In the Cluster nodes pane, you can review the progress of updating.

 **Note:** Remember that one node of the cluster is in a waiting state, and the other node is restarting after it is updated.

7. Wait until the process is finished.

Note: This may require a restart of both nodes.

The process is finished when both nodes show **Succeeded** in the **Last Run status** column.

8. Sign in to LON-SVR3 with the username **Adatum\Administrator** and the password **Pa$$w0rd**.
9. On LON-SVR3, in the **Server Manager**, click **Tools**, and then click **Cluster-Aware Updating**.
10. In the **Cluster-Aware Updating** dialog box, in the **Connect to a failover cluster** drop-down list box, select **Cluster1**. Click **Connect**.
11. Click the **Configure cluster self-updating options** in the **Cluster Actions** pane.
12. On the **Getting Started** page, click **Next**.
13. On the **Add CAU Clustered Role with Self-Updating Enabled** page, click **Add the CAU clustered role, with self-updating mode enabled, to this cluster**, and then click **Next**.
14. On the **Specify self-updating schedule** page, click **Weekly**, and in the **Time of day** box, select **4:00 AM**, and then in the **Day of the week** box, select **Sunday**. Click **Next**.
15. On the **Advanced Options** page, click **Next**.
16. On the **Additional Update Options** page, click **Next**.
17. On the **Confirmation** page, click **Apply**.
18. After the clustered role is added successfully, click **Close**.

▶ Task 3: Prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. On the **Virtual Machines** list, right-click **20412C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps two and three for 20412C-LON-SVR1, 20412C-LON-SVR3, and 20412C-LON-SVR4.

Results: After this exercise, you will have configured CAU.

Module 11: Implementing Failover Clustering with Hyper-V

Lab: Implementing Failover Clustering with Hyper-V

Exercise 1: Configuring Hyper-V Replicas

▶ Task 1: Import LON-CORE virtual machine on LON-HOST1

1. Sign on to LON-HOST1 as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. On LON-HOST1, open the Hyper-V Manager console.
3. In the Actions pane, click **Import Virtual Machine**.
4. On the **Before You Begin** page of the Import Virtual Machine Wizard, click **Next**.
5. On the **Locate Folder** page, click **Browse**.
6. Browse to folder **E:\Program Files\Microsoft Learning\20412\Drives\20412C-LON-CORE**. Click **Select Folder**, and click **Next**.

 **Note:** The drive letter may be different based upon the number of drives on the physical host machine.

7. On the **Select Virtual Machine** page, select **20412C-LON-CORE**, and then click **Next**.
8. On the **Choose Import Type** page, click **Next**.
9. On the **Connect network** page, ensure that **External Network** is selected and then click **Next.**
10. On the **Summary** page, click **Finish**.

▶ Task 2: Configure a replica on both host machines

1. On LON-HOST2, open the Hyper-V Manager console.
2. In Hyper-V Manager, right-click **LON-HOST2**, and then select **Hyper-V Settings**.
3. In Hyper-V Settings for LON-HOST2, click **Replication Configuration**.
4. In the Replication Configuration pane, click **Enable this computer as a Replica server**.
5. In the Authentication and ports section, select **Use Kerberos (HTTP)**.
6. In the Authorization and storage section, click **Allow replication from any authenticated server**, and then click **Browse**.
7. Click **Computer**, double-click **Local Disk (E)**, and then click **New folder**. Type **VMReplica** for the folder name, and press Enter. Select the **E:\VMReplica** folder, and click **Select Folder**.
8. In Hyper-V Settings for LON-HOST2**,** click **OK**.
9. In the Settings window, read the notice, and then click **OK**.
10. Click the **Start** screen, and click **Control Panel**.
11. In the Control Panel, click **System and Security**, and then click **Windows Firewall**.
12. Click **Advanced settings**.
13. Click **Inbound Rules**.

14. In the right pane, in the rule list, find and right-click the **Hyper-V Replica HTTP Listener (TCP-In)** rule, and then click **Enable Rule**.
15. Close the Windows Firewall with the Advanced Security console, and close the **Windows Firewall**.
16. Repeat steps one through 15 on LON-HOST1.

▶ Task 3: Configure replication for LON-CORE virtual machine

1. On LON-HOST1, open the Hyper-V Manager console. Click **LON-HOST1**, and right-click **20412C-LON-CORE**.
2. Click **Enable Replication**.
3. On the **Before You Begin** page, click **Next**.
4. On the **Specify Replica Server** page, click **Browse**.
5. In the Select Computer window, type **LON-HOST2**, and click **Check Names**, and then click **OK**. Click **Next**.
6. On the **Specify Connection Parameters** page, review the settings, and ensure that **Use Kerberos authentication (HTTP)** is selected, and then click **Next**.
7. On the **Choose Replication VHDs** page, ensure that 20412C-LON-CORE.vhd is selected, and then click **Next**.
8. On the **Configure Replication Frequency** page, select 30 seconds from drop-down list box, and then click **Next**.
9. On the **Configure Additional Recovery Points** page, select **Maintain only the latest recovery point**, and then click **Next**.
10. On the **Choose Initial Replication Method** page, click **Send initial copy over the network**, select **Start replication immediately**, and then click **Next**.
11. On the **Completing the Enable Replication Wizard** page, click **Finish**.
12. Wait five to seven minutes. You can monitor the progress of the initial replication in the **Status** column in the Hyper-V Manager console. When it completes (progress reaches 100 percent), ensure that 20412C-LON-CORE has appeared on LON-HOST2 in Hyper-V Manager.

▶ Task 4: Validate a planned failover to the replica site

1. On LON-HOST2 in Hyper-V Manager, right-click **20412C-LON-CORE**.
2. Select **Replication**, and then click **View Replication Health**.
3. Review content of the window that appears, and ensure that there are no errors.
4. Click **Close**.
5. On LON-HOST1, open Hyper-V Manager, and verify that **20412C-LON-CORE** is turned off.
6. Right-click **20412C-LON-CORE**, select **Replication,** and then click **Planned Failover**.
7. In the Planned Failover window, ensure that the option **Start the Replica virtual machine after failover** is selected, and then click **Fail Over**.
8. On LON-HOST2, in Hyper-V Manager, ensure that **20412C-LON-CORE** is running.
9. On LON-HOST1, right-click **20412C-LON-CORE**, point to **Replication**, and then click **Remove replication**.

10. In the **Remove replication** dialog box, click **Remove Replication**.
11. On LON-HOST2, right-click **20412C-LON-CORE**, and then select **Shut Down**. In the **Shut Down Machine** dialog box, click **Shut Down**.

Results: After completing this exercise, you will have Hyper-V Replica configured.

Exercise 2: Configuring a Failover Cluster for Hyper-V

▶ Task 1: Connect to iSCSI target from both host machines

1. On LON-HOST1, open the **Server Manager**, click **Tools**, and then click **iSCSI Initiator**. At the **Microsoft iSCSI** prompt, click **Yes**.
2. Click the **Discovery** tab.
3. Click **Discover Portal**.
4. In the **IP address or DNS name** box, type **172.16.0.21**, and then click **OK**.
5. Click the **Targets** tab.
6. Click **Refresh**.
7. In the **Targets** list, select **iqn.1991-05.com.microsoft:lon-svr1-target1-target**, and then click **Connect**.
8. Select **Add this connection to the list of Favorite Targets**, and click **OK**.
9. To close iSCSI Initiator Properties, click **OK**.
10. On LON-HOST2, open Server Manager, click **Tools**, and then click **iSCSI Initiator**.
11. In the **Microsoft iSCSI** dialog box, click **Yes**.
12. Click the **Discovery** tab.
13. Click **Discover Portal**.
14. In the IP address or DNS name box, type **172.16.0.21**, and then click **OK**.
15. Click the **Targets** tab.
16. Click **Refresh**.
17. In the **Discovered targets** list, select **iqn.1991-05.com.microsoft:lon-svr1-target1-target**, and then click **Connect**.
18. Select **Add this connection to the list of Favorite Targets**, and click **OK**. To close iSCSI Initiator Properties, click **OK**.
19. On LON-HOST2, in the Server Manager window, click **Tools**, and then click **Computer Management**.
20. Expand **Storage**, and click **Disk Management**.
21. Right-click **Disk 2**, and click **Online**. (Note: The disk letter and number may be different based upon the number of drives on the physical host machine.)
22. Right-click **Disk 2**, and click **Initialize Disk**. In the **Initialize Disk** dialog box, click **OK**.
23. Right-click the unallocated space next to **Disk 2**, and click **New Simple Volume**.
24. On the Welcome page, click **Next**.
25. On the **Specify Volume Size** page, click **Next**.
26. On the **Assign Drive Letter or Path** page, click **Next**.
27. On the **Format Partition** page, in the **Volume label** box, type **ClusterDisk**. Select the **Perform a quick format** check box, and click **Next**.
28. Click **Finish**.
29. Repeat steps 21 through 28 for Disk 3 and Disk 4. In step 27, provide name **ClusterVMs** for Disk 3 and **Quorum** for Disk 4.

30. On LON-HOST1 in the Server Manager, click **Tools**, and then click **Computer Management**.
31. Expand **Storage**, and click **Disk Management**.
32. Right-click **Disk Management**, and click **Refresh**.
33. Right-click **Disk 2**, and click **Online**.
34. Right-click **Disk 3**, and click **Online**.
35. Right-click **Disk 4**, and click **Online**.

▶ Task 2: Configure failover clustering on both host machines

1. On LON-HOST1, on the taskbar, to open the Server Manager, click the **Server Manager** icon.
2. On the Dashboard window**,** click **Add roles and features**.
3. On the **Before You Begin** page, click **Next**.
4. On the **Select installation type** page, click **Next**.
5. On the **Select destination server** page, ensure that **Select server from the server pool** is selected, and then click **Next**.
6. On the **Select server roles** page, click **Next**.
7. On the **Select features** page, in the **Features** list, click **Failover Clustering**. In the **Add features that are required for failover clustering** prompt, click **Add Features**, and then click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. When installation is complete, click **Close**.
10. Repeat steps one through nine on LON-HOST2.
11. On LON-HOST1, in the Server Manager console, click **Tools**, and then click **Failover Cluster Manager**.
12. In Failover Cluster Manager, in the center pane, under Management, click **Create Cluster**.
13. On the **Before You Begin** page of the Create Cluster Wizard, read the information, and then click **Next**.
14. In the **Enter server name** box, type **LON-HOST1**, and then click **Add**. Type **LON-HOST2**, and click **Add**.
15. Verify the entries, and click **Next**.
16. On the **Validation Warning** page, click **No. I don't require support from Microsoft for this cluster**, and then click **Next**.
17. In the **Access Point for Administering the Cluster** page, in the **Cluster Name** box, type **VMCluster**.
18. In the **IP address name** box, under Address, type **172.16.0.126**, and then click **Next**.
19. In the **Confirmation** dialog box, verify the information, clear the check box next to **Add all eligible storage to the cluster**, and then click **Next**.
20. On the **Summary** page, click **Finish**.

▶ Task 3: Configure disks for failover cluster

1. On LON-HOST1, in the Failover Cluster Manager console, expand **VMCluster.Adatum.com**, expand **Storage**, and then right-click **Disks**.
2. Click **Add Disk**.

3. In the **Add Disks to Cluster** dialog box, verify that all disks are selected, and then click **OK**.
4. Verify that all disks appear available for cluster storage in **Failover Cluster Manager**.
5. Select the disk that displays the Volume name of **ClusterVMs**. Right-click the **ClusterVMs** disk, and select **Add to Cluster Shared Volumes**. (Note: Click the disk, and the Volume name will display).
6. Right-click **VMCluster.adatum.com**, select **More Actions**, and then click **Configure Cluster Quorum Settings**. Click **Next**.
7. On the **Select Quorum Configuration Option** page, click **Use default quorum configuration**, and then click **Next**.
8. On the **Confirmation** page, click **Next**.
9. On the **Summary** page, click **Finish**.

Results: After completing this exercise, students will have the failover clustering infrastructure configured for Hyper-V.

Exercise 3: Configuring a Highly Available Virtual Machine

▶ Task 1: Copy virtual machine storage to iSCSI target

1. Ensure that LON-HOST1 is the owner of the ClusterVMs disk in Failover Cluster Manager. If it is not, then move the ClusterVMs resource to LON-HOST1 before doing this procedure.
2. On LON-HOST1, open File Explorer, browse to **E:\Program Files\Microsoft Learning\20412\Drives\20412C-LON-CORE\Virtual Hard Disks**, and then copy the **20412C-LON-CORE.vhd** virtual hard disk file to the **C:\ClusterStorage\Volume1** location.

▶ Task 2: Configure the virtual machine as highly available

1. In the Failover Cluster Manager console, click **Roles**, and then in the Actions pane, click **Virtual Machines**.
2. Click **New Virtual Machine**.
3. Select **LON-HOST1** as the cluster node, and click **OK**.
4. In the New Virtual Machine Wizard, click **Next**.
5. On the **Specify Name and Location** page, type **TestClusterVM** for the Name, click **Store the virtual machine in a different location**, and then click **Browse**.
6. Browse to and select **C:\ClusterStorage\Volume1**, and then click **Select Folder**.
7. On the **Specify Generation** page, select **Next**.
8. Click **Next**.
9. On the **Assign Memory** page, type **1536**, and then click **Next**.
10. On the **Configure Networking** page, click **External Network**, and then click **Next**.
11. On the **Connect Virtual Hard Disk** page, click **Use an existing virtual hard disk**, and then click **Browse**.
12. Locate **C:\ClusterStorage\Volume1**, select **20412C-LON-CORE.vhd**, and then click **Open**.
13. Click **Next**, and click **Finish**.
14. On the **Summary** page click **Finish**.
15. Right-click the **TestClusterVM**, and click **Settings**.
16. In the **Settings for TestClusterVM** on LON-Host1, expand **Processor** in the left navigation pane, and then click **Compatibility**.
17. In the right pane, select the check box before the **Migrate to a physical computer with a different processor version** option.
18. Click **OK**.
19. Right-click **TestClusterVM,** and click **Start**.
20. Ensure that the machine successfully starts.

▶ Task 3: Perform a Live Migration for the virtual machine

1. Open the **Failover Cluster Manager** on LON-HOST2.
2. Expand **VMCluster.Adatum.com**, and click **Roles**.
3. Right-click **TestClusterVM**, select **Move**, select **Live Migration**, and then click **Select Node**.
4. Click **LON-HOST2**, and click **OK**.

5. Right-click **TestClusterVM**, and click **Connect**.
6. Ensure that you can access and operate the virtual machine while it is migrating to another host.
7. Wait until the migration is finished.

▶ Task 4: Perform a Storage Migration for the virtual machine

1. On LON-HOST1, open File Explorer, and browse to **E:\Program Files\Microsoft Learning\20412\Drives** folder.
2. In this folder, create a new folder and name it **LON-GUEST1**.
3. On LON-HOST1, open File Explorer, browse to **E:\Program Files\Microsoft Learning\20412\Drives\20412C-LON-CORE\Virtual Hard Disks**, and then copy the **20412C-LON-CORE.vhd** virtual hard disk file to the **E:\Program Files\Microsoft Learning\20412\Drives\LON-GUEST1** location.
4. On LON-HOST1, open the Hyper-V Manager.
5. In the Hyper-V Manager, on the Actions pane, click **New**, and then click **Virtual Machine**.
6. On the **Before You Begin** page of the New Virtual Machine Wizard, click **Next**.
7. On the **Specify Name and Location** page of the New Virtual Machine Wizard, select **Store the virtual machine in a different location**, enter the following values, and then click **Next**:
   - Name: **LON-GUEST1**
   - Location: **E:\Program Files\Microsoft Learning\20412\Drives\LON-GUEST1**
8. On the **Specify Generation** page, click **Next**.
9. On the **Assign Memory** page of the New Virtual Machine Wizard, enter a value of **1024 MB**, select the **Use Dynamic Memory** for this virtual machine option, and then click **Next**.
10. On the **Configure Networking** page of the New Virtual Machine Wizard, select **External Network**, and then click **Next**.
11. On the **Connect Virtual Hard Disk** page, choose **Use an existing virtual hard disk**. Click **Browse**, and browse to **E:\Program Files\Microsoft Learning\20412\Drives\LON-GUEST1 \20412C-LON-CORE.vhd**. Click **Open**, and click **Finish**.
12. In the central pane of Hyper-V Manager, click **LON-GUEST1**.
13. In the Actions pane, click **Start**. Wait until the virtual machine is fully started.
14. Switch back to the Hyper-V Manager console, and in the Actions pane, click **Move**.
15. On the **Before You Begin** page, click **Next**.
16. On the **Choose Move Type** page, select **Move the virtual machine's storage**, and then click **Next**.
17. On the **Choose Options for Moving Storage** page, select **Move all of the virtual machine's data to a single location**, and then click **Next**.
18. On the **Choose a new location for virtual machine** page, click **Browse**.
19. Locate **C:**, and create a new folder named **Guest1**. Click **Select Folder**.
20. Click **Next**.
21. On the **Summary** page, click **Finish**. Wait for the move process to finish. While the virtual machine is moving, you can connect to it, and verify that it is fully operational.
22. Shut down all running virtual machines.

▶ Task 5: Prepare for the next module

1. Restart LON-HOST1.
2. When you are prompted with the boot menu, select **Windows Server 2012**, and then press Enter.
3. Sign in to the host machine as directed by your instructor.
4. Repeat steps one through three on LON-HOST2.

Results: After completing this exercise, the students will have configured the virtual machine as highly available.

Module 12: Implementing Business Continuity and Disaster Recovery

Lab: Implementing Windows Server Backup and Restore

Exercise 1: Backing Up Data on a Windows Server 2012 R2 Server

▶ Task 1: Install Windows Server Backup

1. Switch on LON-SVR1.
2. In the Server Manager, in the **Welcome** pane, click **Add roles and features**.
3. In the Add Roles and Features Wizard, on the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, click **Next**.
5. On the **Select destination server** page, click **Next**.
6. On the **Select server roles** page, click **Next**.
7. On the **Select features** page, select **Windows Server Backup**, and then click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. On the **Installation progress** page, wait until the **Installation succeeded on LON-SVR1.adatum.com** message displays, and then click **Close**.

▶ Task 2: Configure a scheduled backup

1. On LON-SVR1, in the **Server Manager**, click **Tools**, and then click **Windows Server Backup**.
2. In the navigation pane, click **Local Backup**.
3. Click **Backup Schedule**.
4. In the Backup Schedule Wizard, on the **Getting Started** page, click **Next**.
5. On the **Select Backup Configuration** page, click **Full server (recommended)**, and then click **Next**.
6. On the **Specify Backup Time** page, next to **Select time of day**, select **1:00 AM**, and then click **Next**.
7. On the **Specify Destination Type** page, click **Backup to a shared network folder**, and then click **Next**. Review the warning, and then click **OK**.
8. On the **Specify Remote Shared Folder** page, in the **Location** text box, type **\\LON-DC1\Backup**, and then click **Next**.
9. In the **Register Backup Schedule** dialog box, in the **Username** text box, type **Administrator**, and in the **Password** text box, type **Pa$$w0rd**, and then click **OK**.
10. Click **Finish**, and then click **Close**.

 **Note:** In a production environment, you will not store backup to a domain controller. You do it here for lab purposes only.

▶ Task 3: Complete an on-demand backup

1. On LON-SVR1, switch to **Windows Server Backup**.
2. In the **Actions** pane, click **Backup Once**.
3. In the Backup Once Wizard, on the **Backup Options** page, click **Different options**, and then click **Next**.
4. On the **Select Backup Configuration** page, click **Custom**, and then click **Next**.
5. On the **Select Items for Backup** page, click **Add Items**.
6. Expand **Local disk (C:)**, select the **Financial Data** check box, click **OK**, and then click **Next**.
7. On the **Specify Destination Type** page, click **Remote shared folder**, and then click **Next**.
8. On the **Specify Remote Folder** page, type **\\LON-DC1\Backup**, and then click **Next**.
9. On the **Confirmation** page, click **Backup**.
10. On the **Backup Progress** page, after the backup is complete, click **Close**.

Results: After you complete this exercise, you will have configured the Windows Server Backup feature, scheduled a backup task, and completed an on-demand backup.

Exercise 2: Restoring Files Using Windows Server Backup

▶ Task 1: Delete a file from the server

1. On LON-SVR1, on the taskbar, click the Windows Explorer icon.
2. In File Explorer, click to **Local Disk (C:)**, right-click **Financial Data**, and then click **Delete**.

▶ Task 2: Restore a file from backup

1. In the Windows Server Backup console, in the **Actions** pane, click **Recover**.
2. On the **Getting Started** page, click **A backup stored on another location**, and then click **Next**.
3. On the **Specify Location Type** page, click **Remote shared folder**, and then click **Next**.
4. On the **Specify Remote Folder** page, type **\\LON-DC1\Backup**, and then click **Next**.
5. On the **Select Backup Date** page, click **Next**.
6. On the **Select Recovery Type** page, click **Next**.
7. On the **Select Items to Recover** page, expand **LON-SVR1**, click **Local Disk (C:)**, and in the right pane, select **Financial Data**, and then click **Next**.
8. On the **Specify Recovery Options** page, under **Another Location**, type **C:**, and then click **Next**.
9. On the **Confirmation** page, click **Recover**.
10. On the **Recovery Progress** page, click **Close**.
11. Open drive **C**, and verify that the **Financial Data** folder is restored.

▶ Task 3: Prepare for the next module

When you finish the lab, revert the virtual machines to their initial state.

1. On the host computer, start the Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20412C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps two and three for 20412C-LON-SVR1.

Results: After completing this exercise, you will have tested and validated the procedure for restoring a file from backup.

Notes

Notes

Notes

Notes

Notes

Notes

Notes

Notes

Notes

Notes

Notes

Notes