Microsoft

OFFICIAL MICROSOFT LEARNING PRODUCT

20413C

Designing and Implementing a Server Infrastructure

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

Product Number: 20413C

Part Number: X19-30968

Released: 4/2014

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

Dave Franklyn – Content Developer

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

Vladimir Meloski – Content Developer

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

Marcus Oh – Technical Reviewer

Marcus Oh, System Center Cloud and Datacenter Management MVP, is a senior technical manager for a large telecommunications provider, running directory services and management infrastructure for ~30,000 systems. He has been an MVP since 2004 in System Center, specializing in Configuration Manager, Operations Manager, and Orchestrator. Marcus has written numerous articles for technology websites and blogs on Orchestrator and other System Center components at http://marcusoh.blogspot.com. He coauthored Professional SMS 2003, MOM 2005, and WSUS (Wrox, 2006), was a contributing author to System Center Opalis Integration Server 6.3 Unleashed (2011), and coauthored System Center 2012 Configuration Manager Unleashed (2012). Most recently, Marcus coauthored System Center 2012 Orchestrator Unleashed (2013). Marcus is also the president of the Atlanta Systems Management User Group (http://www.atlsmug.com) and a board member of the Deskside Management Forum.

Telmo Sampaio- Content Developer

Telmo Sampaio is the Chief Geek at MCTrainer.NET and TechKnowLogical in Miami, FL specializing in Windows Server, System Center, SharePoint, SQL and .NET. He is a trainer, consultant, author and speaker at events such as TechEd, MMS, and PASS. Telmo is very active in the MCT community, being one of the first MCT Regional Leads.

David Susemiehl – Content Developer

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 development for Microsoft and Hewlett-Packard, and delivered those courses successfully in Europe, Central America, and across North America. For the last several years, David has been writing courseware for Microsoft Learning, and consulting on infrastructure transitions in Michigan.

Brian Svidergol – Content Developer

Brian Svidergol specializes in Microsoft infrastructure and cloud-based solutions built around Windows, Active Directory, Microsoft Exchange, System Center, virtualization, and MDOP. He holds a bunch of Microsoft and industry certifications. Brian authored the Active Directory Cookbook 4th Edition. He has also worked as an SME and technical reviewer on many Microsoft Official Curriculum courses, Microsoft certification exams, and authored or reviewed related training content.

Orin Thomas – Content Developer

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

Contents

Module 1: Planning Server Upgrade and Migration

Module 2: Planning and Implementing a Server Deployment Strategy

Module 3: Planning and Deploying Servers Using Virtual Machine Manager

Module 4: Designing and Maintaining an IP Configuration and Address Management Solution

Module 5: Designing and Implementing Name Resolution

Module 6: Designing and Implementing an Active Directory Domain Services Forest and Domain Infrastructure

Module 7: Designing and Implementing an AD DS Organizational Unit Infrastructure

Module 8: Designing and Implementing a Group Policy Object Strategy

Module 9: Designing and Implementing an AD DS Physical Topology

Module 10: Planning and Implementing Storage and File Services

Module 11: Designing and Implementing Network Protection

Module 12: Designing and Implementing Remote Access Services

Lab Answer Keys

About This Course

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

Course Description

Get hands-on instruction and practice planning, designing and deploying a physical and logical Windows Server 2012 R2 enterprise infrastructure in this 5-day Microsoft Official course. This course is part one in a series of two courses that provides the skills and knowledge necessary to design and implement a Windows Server 2012 R2 infrastructure in an enterprise environment. The two courses collectively cover designing, planning, deploying, securing, monitoring, automating, and virtualizing an enterprise server infrastructure. This course covers the knowledge and skills needed to provide an enterprise solution that supports manual and automated server installations in a physical and virtual environment including the supporting file and storage services. You will also learn the skills necessary to provide enterprise networking solutions such as DHCP, IPAM, VPN & DirectAccess. You will also learn the skills necessary to design and implement a forest and domain infrastructure including multi domains/forest and branch office scenarios.

Audience

This course is intended for IT professionals who are responsible for planning, designing, and deploying a physical and a logical Windows Server 2012 enterprise AD DS infrastructure, including the necessary network services. They have experience of previous Windows Server operating systems and possess Windows Server 2012 certification Microsoft Certified Solutions Associate (MCSA) or equivalent skills.

The secondary audience for this course includes IT professionals who are looking to take the exam 70-413: *Designing and Implementing a Server Infrastructure*, as a stand-alone, or as part of the requirement for the Microsoft Certified Solutions Expert (MCSE): Server Infrastructure Certification.

Student Prerequisites

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

- A good understanding of TCP/IP fundamentals and networking concepts.
- A good working knowledge of both Windows Server 2012 R2 and AD DS. For example, domain user accounts, domain versus local user accounts, user profiles, and group membership.
- A good understanding of both scripts and batch files.
- A solid understanding of security concepts, such as authentication and authorization.
- Familiarity with deployment, packaging, and imaging tools.
- Ability to work in a team, or as a virtual team.
- Ability to produce good documentation and have the appropriate communication skills to create proposals and make recommendations.
- Knowledge equivalent to Windows 2012 R2 MCSA.

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

- 20410: Installing and Configuring Windows Server 2012

- 20411: Administering Windows Server 2012
- 20412: Configuring Advanced Windows Server 2012 Services, OR
- 20417: Upgrading Your Skills to MCSA Windows Server 2012

Course Objectives

After completing this course, students will be able to:

- Implement server upgrade and migration.
- Design an automated server installation strategy.
- Plan and implement a server deployment infrastructure.
- Plan and implement a System Center 2012 R2 Virtual Machine Manager infrastructure.
- Plan and implement file and storage services.
- Design and implement a Dynamic Host Configuration Protocol (DHCP) solution.
- Design a name resolution solution strategy.
- Design and manage an IP address management solution.
- Design a VPN solution.
- Design a DirectAccess solution.
- Implement a scalable remote access solution.
- Design a network protection solution.
- Implement a network protection solution.
- Design a forest and domain infrastructure.
- Implement a forest and domain infrastructure.
- Design a Group Policy strategy.
- Design an Active Directory permission model.
- Design an Active Directory sites topology.
- Design a domain controller strategy.
- Design and implement a branch office infrastructure.

Course Outline

The course outline is as follows:

Module 1, Planning Server Upgrade and Migration

This module explains how to plan a server upgrade and migration strategy.

Module 2, Planning and Implementing a Server Deployment Strategy

This module explains how to design an automated server installation strategy and plan and implement a server deployment infrastructure.

Module 3, Planning and Deploying Servers Using Virtual Machine Manager

This module explains how to plan and deploy a Virtual Machine Manager (VMM) infrastructure for deploying servers.

Module 4, Designing and Maintaining an IP Configuration and Address Management Solution

This module explains how to design and maintain IP address management (IPAM) and a Dynamic Host Configuration Protocol (DHCP) solution.

Module 5, Designing and Implementing Name Resolution

This module explains how to design a name resolution strategy.

Module 6, Designing and Implementing an AD DS Forest and Domain Infrastructure

This module explains how to design and implement an AD DS forest and domain infrastructure.

Module 7, Designing and Implementing an AD DS Organizational Unit Infrastructure

This module explains how to design and implement an OU infrastructure and an AD DS permissions model.

Module 8, Designing and Implementing a Group Policy Object Strategy

This module explains how to design and implement a Group Policy Object (GPO) strategy.

Module 9, Designing and Implementing an AD DS Physical Topology

This module explains how to design an AD DS sites topology and a domain controller placement strategy.

Module 10, Planning and Implementing Storage and File Services

This module explains how to plan and implement storage and file services.

Module 11, Designing and Implementing Network Protection

This module explains how to design and implement network protection.

Module 12, Designing and Implementing Remote Access Services

This module explains how to design and implement remote access services.

Exam/Course Mapping

This course *20413C: Designing and Implementing a Server Infrastructure*, maps directly to, and is the preferred choice for, hands-on preparation for Microsoft Certified Solutions Expert (MCSE): Exam 413: *Designing and Implementing and Server Infrastructure*, which is the fourth of five exams required for MCSE: Server Infrastructure certification.

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 use the unique experience and skills of your qualified Microsoft Certified Trainer (MCT).

Note: The exam objectives are available online at: http://www.microsoft.com/learning/en/us/exam.aspx?id=70-413.

Exam Objective Domain: 70-413: Designing and Implementing a Server Infrastructure		**Course Content**		
1. Plan and deploy a server infrastructure (20 – 25%)		Module	Lesson	Lab
1.1. Design an automated server installation strategy	This objective may include but is not limited to: design considerations, including images and bare metal/virtual deployment; design a server implementation using Windows Assessment and Deployment Kit (ADK); design a virtual server deployment	Mod 2	Lesson 2	Mod 2 Lab
1.2. Plan and implement a server deployment infrastructure	This objective may include but is not limited to: configure multicast deployment; configure multi-site topology and distribution points; configure a multi-server topology; configure autonomous and replica Windows Deployment Services (WDS) servers	Mod 2	Lesson 1	Mod 2 Lab
1.3. Plan and implement server upgrade and migration	This objective may include but is not limited to: plan for role migration; migrate server roles; migrate servers across domains and forests; design a server consolidation strategy; plan for capacity and resource optimization	Mod 1	Lessons 1/2/3	Mod 1 Lab
		Mod 2	Lesson 1	Mod 2 Lab
1.4 Plan and deploy Virtual Machine Manager services	This objective may include but is not limited to: design Virtual Machine Manager service templates; define operating system profiles; configure hardware and capability profiles; manage services; configure image and template libraries; manage logical networks	Mod 3	Lessons 1/2/3	Mod 3 Lab
1.5 Plan and implement file and storage services	This objective may include but is not limited to: planning considerations include iSCSI SANs, Fibre Channel SANs, Virtual Fibre Channel, storage spaces, storage pools, and data de-duplication; configure the iSCSI Target server; configure the Internet Storage Name server (iSNS); configure Network File System (NFS); install Device Specific Modules (DSMs)	Mod 10	Lessons 1/2/3	Mod 10 Lab
2. Design and implement network infrastructure services (20 – 25%)				
2.1. Design and maintain a Dynamic Host Configuration Protocol (DHCP) solution	This objective may include but is not limited to: design considerations, including a highly available DHCP solution including split scope, DHCP failover, and DHCP failover clustering, DHCP interoperability, and DHCPv6; implement DHCP filtering; implement and configure a DHCP management pack; maintain a DHCP database	Mod 4	Lessons 1/2	Mod 4 Lab
2.2 Design a name resolution solution strategy	This objective may include but is not limited to: design considerations, including secure name resolution, DNSSEC, DNS Socket Pool, cache locking, disjoint namespaces, DNS interoperability, migration to application partitions, IPv6, Single-Label DNS Name Resolution, zone hierarchy, and zone delegation	Mod 5	Lessons 1/2/3/4/5/6	Mod 5 Lab

Exam Objective Domain: 70-413: Designing and Implementing a Server Infrastructure		Course Content		
2.3. Design and manage an IP address management solution	This objective may include but is not limited to: design considerations, including IP address management technologies including IPAM, Group Policy based, manual provisioning, and distributed vs. centralized placement; configure role-based access control; configure IPAM auditing; migrate IPs; manage and monitor multiple DHCP and DNS servers; configure data collection for IPAM	Mod 4	Lessons 3/4	Mod 4 Lab
3. Design and implement network access services (15 – 20%)				
3.1. Design a VPN solution	This objective may include but is not limited to: Design considerations including certificate deployment, firewall configuration, client/site to site, bandwidth, protocol implications, and VPN deployment configurations using Connection Manager Administration Kit (CMAK).	Mod 12	Lesson 2	Mod 12 Lab
3.2 Design a DirectAccess solution	This objective may include but is not limited to: design considerations, including topology, migration from Forefront UAG, DirectAccess deployment, and enterprise certificates	Mod 12	Lesson 1	Mod 12 Lab
3.3 Implement a scalable remote access solution	This objective may include but is not limited to: Configure site-to-site VPN; configure packet filters; implement packet tracing; implement multi-site Remote Access; configure Remote Access clustered with Network Load Balancing (NLB); configure DirectAccess	Mod 12	Lesson 3	Mod 12 Lab
3.4 Design a network protection solution	This objective may include but is not limited to: Design considerations including Network Access Protection (NAP) enforcement methods for DHCP, IPSec, VPN, and 802.1x, capacity, placement of servers, firewall, Network Policy Server (NPS), and remediation network	Mod 11	Lessons 1/2/3	Mod 11 Lab
3.5 Implement a network protection solution	This objective may include but is not limited to: Implement multi-RADIUS deployment; configure NAP enforcement for IPSec and 802.1x; deploy and configure the Endpoint Protection client; create anti-malware and firewall policies; monitor for compliance	Mod 11	Lessons 1/2/3	Mod 11 Lab
4. Design and implement an Active Directory infrastructure (logical) (20 – 25%)				
4.1 Design a forest and domain infrastructure	This objective may include but is not limited to: design considerations, including multi-forest architecture, trusts, functional levels, domain upgrade, domain migration, forest restructure, and hybrid cloud service	Mod 6	Lessons 1/2/3/4/5/6	Mod 6 Labs A/B

Exam Objective Domain: 70-413: Designing and Implementing a Server Infrastructure		Course Content		
4.2 Implement a forest and domain infrastructure	This objective may include but is not limited to: configure domain rename; configure Kerberos realm trusts; implement a domain upgrade; implement a domain migration; implement a forest restructure; deploy and manage a test forest including synchronization with production forests	Mod 6	Lessons 1/2/3/4/5/6	Mod 6 Labs A/B
4.3 Design a Group Policy strategy	This objective may include but is not limited to: design considerations, including inheritance blocking, enforced policies, loopback processing, security, and WMI filtering, site-linked Group Policy Objects (GPOs), slow-link processing, group strategies, organizational unit (OU) hierarchy, and Advanced Group Policy Management (AGPM)	Mod 8	Lessons 1/2/3/4	Mod 8 Lab
4.4 Design an Active Directory permission model	This objective may include but is not limited to: design considerations, including Active Directory object security and Active Directory quotas; customize tasks to delegate in Delegate of Control Wizard; deploy administrative tools on the client computer; delegate permissions on administrative users (AdminSDHolder); configure Kerberos delegation	Mod 7	Lessons 1/2/3	Mod 7 Lab
5. Design and implement an Active Directory infrastructure (physical) (20 – 25%)				
5.1 Design an Active Directory sites topology	This objective may include but is not limited to: design considerations, including proximity of domain controllers, replication optimization, and site link; monitor and resolve Active Directory replication conflicts	Mod 9	Lessons 1/2/3	Mod 9 Lab
5.2 Design a domain controller strategy	This objective may include but is not limited to: design considerations, including global catalog, operations master roles, Read-Only Domain Controllers (RODCs), partial attribute set, and domain controller cloning	Mod 9	Lessons 3/4/5	Mod 9 Lab
5.3 Design and implement a branch office infrastructure	This objective may include but is not limited to: design considerations, including RODC, Universal Group Membership Caching (UGMC), global catalog, DNS, DHCP, and BranchCache; implement confidential attributes; delegate administration; modify filtered attributes set; configure password replication policy; configure hash publication	Mod 9	Lessons 1/2/3/4/5	Mod 9 Lab

Important Attending this course in itself does not guarantee that you will pass any associated certification exams.

In addition to attendance at this course, you should also have the following:

- A good understanding of TCP/IP fundamentals and networking concepts.
- A good working knowledge of both Windows Server 2012 and AD DS. For example, domain user accounts, domain versus local user accounts, user profiles, and group membership.
- A good understanding of both scripts and batch files.
- A solid understanding of security concepts, such as authentication and authorization.
- Familiarity with deployment, packaging, and imaging tools.
- Ability to work in a team and on a virtual team.
- Ability to produce good documentation and have the appropriate communication skills to create proposals and make budget recommendations.
- Knowledge equivalent to Windows 2012 R2 MCSA.

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 http://www.microsoft.com/learning/en/us/course.aspx?id=20413C, 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 http://www.microsoft.com/learning/en/us/course.aspx?id=20413C under Overview, Audience Profile.

The exam/course mapping table previously outlined 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/companionmoc* 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, Microsoft Developer Network (MSDN®), or Microsoft Press®.

Student Course files. On the http://www.microsoft.com/learning/companionmoc site.

- ***Course evaluation.*** At the end of the course, you will have the opportunity to complete an online evaluation to provide feedback on the course, training facility, and instructor.

Virtual Machine Environment

This section provides the information about the lab scenario that is used in this course.

Virtual Machine Configuration

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

Important At the end of each lab, you must revert the virtual machines to a snapshot. You can find the instructions for this procedure at the end of each lab.

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

Virtual machine	Role
20413C-LON-DC1/-B	A domain controller running Windows Server 2012 R2 in the Adatum.com domain.
20413C-LON-SVR1	A member server running Windows Server 2012 R2 in the Adatum.com domain.
20413C-LON-SVR2	A member server running Windows Server 2012 R2 in the Adatum.com domain
20413C-LON-SVR3	A blank virtual machine on which you will install Windows Server 2012 R2.
20413C-LON-SVR4	A member server running Windows Server 2012 R2 in the Adatum.com domain. This server is located on a second subnet.
20413C-LON-RTR	A router that is used for network activities requiring a separate subnet.
20413C-LON-Host1	A boot-to-VHD Windows 2012 R2 host machine that is used for the Virtual Machine Manager lab.
20413C-LON-VMM1	A server with Virtual Machine Manager deployed.
20413C-TREY-DC1	A domain controller running Windows Server 2012 R2 in the Treyresearch.net domain. This server is used in a variety of labs, principally those where multiple domains are required.
20413C-CON-SVR	A stand-alone server running Windows Server 2012 R2 that you will use for joining domains and initial configuration. It is part of the Contoso Ltd organization.
20413C-LON-CL1 20413C-LON-CL2	Client computers running Windows 8.1 and Microsoft Office 2013 in the Adatum.com domain. You will use these computers primarily to test server configurations.

Software Configuration

The following software is installed on each virtual machine:

- Windows Server 2012 R2
- Windows 8.1 Enterprise
- Microsoft Office 2013
- Solution accelerators: MAP 8.5, MAP sample database
- System Center 2012 R2 Virtual Machine Manager
- SQL Server 2012

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.

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

Course Hardware Level

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

- The minimum equipment configuration for this course is hardware level 7 with 16 gigabytes (GB) of random access memory (RAM)

Hardware Level 7

- Intel Virtualization Technology (Intel VT) or AMD Virtualization (AMD-V) processor
- Dual 120 gigabyte (GB) hard disks 7200 RM SATA or better. The hard disks should be configured with a separate volume (Drive C: and Drive D:) on each hard disk.
- 16 GB random access memory (RAM) or higher
- DVD drive
- Network adapter
- Super VGA (SVGA) 17-inch monitor
- Microsoft Mouse or compatible pointing device
- Sound card with amplified speakers

Module 1

Planning Server Upgrade and Migration

Contents:

Module Overview

Planning an operating-system deployment can be one of your organization's most important activities. Planning must begin with your organization's business requirements and goals. The information technology (IT) department is responsible for determining an appropriate solution that meets an organization's business requirements, and then an organization typically spends significant time on design and planning the operating-system deployment. A well-designed solution can result in an IT infrastructure that is cost-effective and yields positive return on investment (ROI). Planning should produce detailed documentation and checklists for the steps that the deployment will include. Additionally, documentation should include major decisions about the new solution, including the operating-system edition that you are deploying, the licensing model you will use, and whether you will deploy the solution in a physical or virtual environment.

Because Windows Server® 2012 is a cloud-ready operating system, one of an organization's most important decisions is whether to use virtualization technology or physical servers. Organizations also must create a compatibility plan in which they check all current infrastructure and application solutions for compatibility with Windows Server 2012, and ascertain whether an upgrade or migration is necessary.

Objectives

In this module, you will learn on how to plan a server upgrade and migration strategy for Windows Server 2012 by:

- Analyzing upgrade and migration considerations.
- Creating a server upgrade and migration plan.
- Planning for virtualization.

Lesson 1
Considerations for Upgrades and Migrations

When planning your Windows Server 2012 operating-system deployment, you must determine which edition of the operating system best suits your organization. To do this, you must consider your organization’s business needs, the solution’s cost, and the ROI.

You must have a firm understanding of your organization's requirements to select and then deploy the appropriate Windows Server 2012 edition. You also must understand which hardware configuration is appropriate for Windows Server 2012, whether a virtual deployment is more suitable than a physical deployment, and which installation method enables you to deploy Windows Server 2012 efficiently. This lesson provides an overview of the different Windows Server 2012 editions, hardware requirements, deployment options, and installation processes.

Lesson Objectives

At the end of this lesson, you will be able to:

- Describe the different Windows Server 2012 editions.
- Describe the recommended minimum requirements for installing Windows Server 2012.
- Differentiate between an in-place upgrade and server migration.
- Describe the supported in-place upgrade scenarios.
- Describe the benefits of migrating to Windows Server 2012.
- Describe the tools that are available to help plan for an upgrade and migration.
- Plan for server consolidation.
- Plan for cloud server deployments.

Windows Server 2012 Editions

There are four editions of the Windows Server 2012 operating system. Organizations should select the Windows Server 2012 edition that best meets their needs. Systems administrators can save costs by selecting the appropriate Windows Server 2012 edition when deploying a server for a specific role. The following table details the four Windows Server 2012 editions.

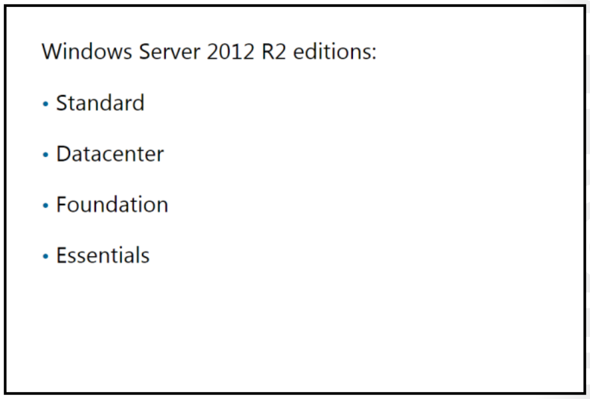

Windows Server 2012 R2 editions:

- Standard
- Datacenter
- Foundation
- Essentials

Edition	Features
Windows Server 2012 R2 Standard	• Provides all roles and features available on the Windows Server 2012 platform. • Supports up to 64 sockets and up to 4 terabytes (TB) of random access memory (RAM).

Edition	Features
	• Includes two virtual machine licenses for a server that has up to two processors. An additional license is necessary for each additional two processors.
Windows Server 2012 R2 Datacenter	• Provides all roles and features that are available on the Windows Server 2012 platform. • Supports up to 64 sockets and up to 4 terabytes (TB) of random access memory (RAM). • Includes unlimited virtual machine licenses for virtual machines that are run on the same hardware for a server that has up to two processors. An additional license is necessary for each additional two processors.
Windows Server 2012 R2 Foundation	• Allows only 15 users, and cannot join to a domain. • Supports one processor core and up to 32 gigabytes (GB) of RAM. • Includes limited server roles. • Does not include Active Directory® Domain Services (AD DS). • Offered through original equipment manufacturer (OEM) program.
Windows Server 2012 R2 Essentials	• Serves as the next edition of Small Business Server. • Operates as a single, multipurpose server in small organizations. • Provides AD DS and Active Directory Certificate Services (AD CS). • Does not support Microsoft Hyper-V® Server failover clustering server, or Remote Desktop Services. • Cannot install Server Core. • Supports up to 25 users and 50 devices. • Supports two processor cores and 64 GB of RAM. • Must be the only domain controller in the domain. • Can be installed as a stand-alone or as a role within Windows Server 2012 R2 Standard or Windows Server 2012 R2 Datacenter editions.

The first consideration in choosing the appropriate Windows Server 2012 edition is the number of users that connect to a server. If that number is greater than 25, then you should choose either the Windows Server 2012 Standard edition or the Windows Server 2012 Datacenter edition.

Choosing between Windows Server 2012 R2 Datacenter and Windows Server Standard

In earlier editions of Windows Server, organizations had to base their choice on the different capabilities of the Standard, Enterprise, or Datacenter editions. Now, organizations have a simple and economic choice between Standard and Datacenter editions, based on only one consideration—virtualization. Windows Server 2012 Standard and Datacenter editions have the same set of capabilities, except for virtualization. Although the Windows Server 2012 Standard operating system includes two virtual machine licenses, the Windows Server 2012 Datacenter operating system includes unlimited virtual machine licenses.

The number of processors per physical service also determines the number of necessary licenses. Organizations that use physical servers that have up to two processors will need one license regardless of which edition they use—Windows Server 2012 Standard or Datacenter. If the physical server has more than two processors, then an additional license is necessary for each additional two processors.

If your organization’s strategy is to deploy servers and applications in a virtual environment, then the Windows Server 2012 Datacenter operating system is the preferable choice. If your organization’s strategy is to deploy servers and applications in mostly nonvirtual environments, then you should select the Windows Server 2012 Standard operating system.

Using Windows Server 2012 R2 Foundation

The Windows Server 2012 Foundation operating system is suitable for small organizations that do not require AD DS and that have fewer than 15 users. The OEM program makes this edition available.

Using Windows Server 2012 R2 Essentials

The Windows Server 2012 Essentials operating system is suitable for an organization with fewer than 25 users. This edition does not have enterprise features, such as virtualization or high availability, and is not available for server core deployment.

The Windows Server 2012 Essentials operating system can also be installed as a role within Windows Server 2012 R2 Standard edition or Windows Server 2012 R2 Datacenter edition. The role that can be installed in Windows Server 2012 R2 Standard or Datacenter Edition is called Windows Server Essentials Experience. The Essentials Experience role includes functionalities of Windows Server 2012 R2 Essentials, such as Dashboard and client computer backups, and it does not have the functionality limits and locks that exist in the stand-alone deployment of Windows Server 2012 R2 Essentials.

Preinstallation Requirements

The following table lists the minimum hardware requirements for Windows Server 2012.

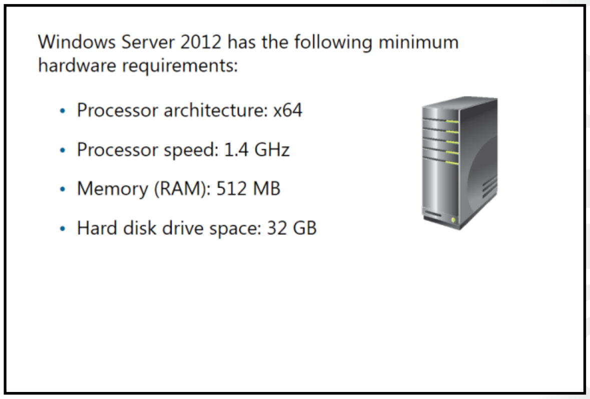

Windows Server 2012 has the following minimum hardware requirements:

- Processor architecture: x64
- Processor speed: 1.4 GHz
- Memory (RAM): 512 MB
- Hard disk drive space: 32 GB

Component	Requirement
Processor architecture	x64
Processor speed	1.4 gigahertz (GHz)
Memory (RAM)	512 megabytes (MB)
Hard disk drive space	32 GB

The hardware requirements that the previous table lists define the absolute minimum requirements to run the server software. Because each service and feature or server role places a unique load on the network, and the resources for disk input/output (I/O), the processor, and memory, the actual hardware requirements depend on the following:

- The applications and the services that the server is running.
- The number of users who are connecting to the server.
- Whether the solution is running in a physical or virtual environment.

Furthermore, when estimating hardware requirements, you should consider whether you will implement the solution in a high availability configuration, where you distribute application load among multiple servers, or if you will run it on a single server. If you implement it in a high availability configuration, the solution might require less powerful hardware because it may distribute server utilization between multiple servers.

Additionally, when planning for hardware requirements, you should consider best practices or recommendations for the specific products that you are installing, such as Microsoft Exchange Server, Microsoft SQL Server®, or Microsoft System Center.

Virtualized deployments of Windows Server 2012 must match the same hardware specifications as physical deployments. Hyper-V and certain non-Microsoft virtualization platforms support Windows Server 2012.

Additional Reading: For more information about the Windows Server Virtualization Validation Program, see Welcome to the Windows Server Virtualization Validation Program at http://go.microsoft.com/fwlink/?linkid=279917.

In-Place Upgrade vs. Server Migration

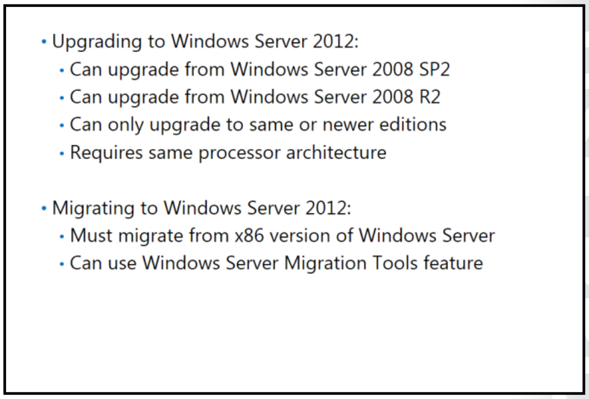

- Upgrading to Windows Server 2012:
  - Can upgrade from Windows Server 2008 SP2
  - Can upgrade from Windows Server 2008 R2
  - Can only upgrade to same or newer editions
  - Requires same processor architecture
- Migrating to Windows Server 2012:
  - Must migrate from x86 version of Windows Server
  - Can use Windows Server Migration Tools feature

When deploying Windows Server 2012, organizations must make the following choice:

- Use existing hardware and upgrade from supported editions of Windows Server 2008 or Windows Server 2008 R2 to Windows Server 2012 or Windows Server 2012 R2.
- Install Windows Server 2012 on new hardware, and, if required, migrate the roles, features, and settings from servers that are running from supported earlier Windows Server editions.

When planning whether to upgrade or migrate a server to Windows Server 2012, consider the options that the following table shows.

Installation option	Description
Upgrade	An upgrade preserves the files, settings, and applications that are installed on the original server. You perform an upgrade when you want to keep all these items and want to continue using the same server hardware. An upgrade requires x64 processor architecture and an x64 edition of the Windows Server operating system. If you are upgrading from Windows Server 2008, you must install Service

<table>
<tr><th>Installation option</th><th>Description</th></tr>
<tr><td></td><td>Pack 2 (SP2). If you are upgrading from Windows Server 2008 R2, you must install Service Pack 1 (SP1).
You start an upgrade by running Setup.exe from the original Windows Server operating system.
You can perform the upgrades to Windows Server 2012 that the following table lists.
<table>
<tr><th>Original operating system and edition</th><th>Upgrade edition</th></tr>
<tr><td>Windows Server 2008 Standard or Windows Server 2008 Enterprise</td><td>Windows Server 2012 Standard, Windows Server 2012 Datacenter</td></tr>
<tr><td>Windows Server 2008 Datacenter</td><td>Windows Server 2012 Datacenter</td></tr>
<tr><td>Windows Web Server 2008</td><td>Windows Server 2012 Standard</td></tr>
<tr><td>Windows Server 2008 R2 Standard or Windows Server 2008 R2 Enterprise</td><td>Windows Server 2012 Standard, Windows Server 2012 Datacenter</td></tr>
<tr><td>Windows Server 2008 R2 Datacenter</td><td>Windows Server 2012 Datacenter</td></tr>
<tr><td>Windows® Web Server 2008 R2</td><td>Windows Server 2012 Standard</td></tr>
<tr><td>Windows Server 2008 R2 Datacenter with SP1</td><td>Windows Server 2012 R2 Datacenter</td></tr>
<tr><td>Windows Server 2008 R2 Enterprise with SP1</td><td>Windows Server 2012 R2 Standard or Windows Server 2012 R2 Datacenter</td></tr>
<tr><td>Windows Server 2008 R2 Standard with SP1</td><td>Windows Server 2012 R2 Standard or Windows Server 2012 R2 Datacenter</td></tr>
<tr><td>Windows Web Server 2008 R2 with SP1</td><td>Windows Server 2012 R2 Standard</td></tr>
<tr><td>Windows Server 2012 Datacenter</td><td>Windows Server 2012 R2 Datacenter</td></tr>
<tr><td>Windows Server 2012 Standard</td><td>Windows Server 2012 R2 Standard or Windows Server 2012 R2 Datacenter</td></tr>
</table>
</td></tr>
<tr><td>Migration</td><td>Use migration when you migrate from an x86 edition of Windows Server 2003, Windows Server 2003 R2, or Windows Server 2008. You can use the Windows Server Migration Tools feature in Windows Server 2012 to transfer</td></tr>
</table>

Installation option	Description
	files and settings from computers that are running the following editions: • Windows Server 2003 • Windows Server 2003 R2 • Windows Server 2008 • Windows Server 2008 R2

Additional Reading: For more information on migration, see Install, Use, and Remove Windows Server Migration Tools at http://go.microsoft.com/fwlink/?linkid=280376.

In-Place Upgrade Scenarios

An in-place upgrade involves upgrading a Windows Server operating system on the server that is running an earlier Windows Server edition. A benefit of an in-place upgrade is that you avoid hardware expenses, because when you select this method, you install Windows Server 2012 on the existing hardware. You would choose an in-place upgrade of the Windows Server operating system in the following scenarios:

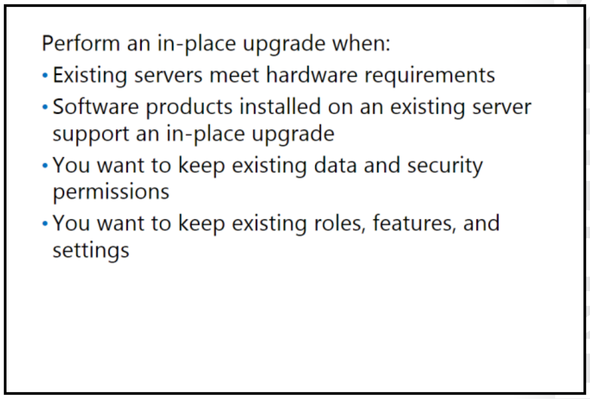

Perform an in-place upgrade when:

- Existing servers meet hardware requirements
- Software products installed on an existing server support an in-place upgrade
- You want to keep existing data and security permissions
- You want to keep existing roles, features, and settings

- When the hardware configuration of the existing servers meets the requirements for Windows Server 2012. Because the hardware requirements for Windows Server 2012 do not differ significantly from those for Windows Server 2008 and Windows Server 2008 R2, you can perform an in-place upgrade on those servers.
- When the software products that run on the existing servers support in-place upgrade of Windows Server 2012. Before performing an in-place upgrade, you must list all of the software products that are running on the server, such as SQL Server, Exchange Server, non-Microsoft software, and antivirus software. Next, verify that these products support an in-place upgrade of Windows Server 2012. If so, refer to the specific product's documentation to determine how to perform an in-place upgrade, including any issues or risks that might occur.
- When you want to keep all user data that is on the existing servers, such as data stored on file servers, and security permissions for accessing those data. When performing an in-place upgrade, user data and security permissions for accessing the data remain unchanged. This scenario is convenient, because after the in-place upgrade, users can continue to access their data that on the same file servers.
- When you want to install Windows Server 2012, but you want to keep all roles, features, and settings of the existing server. Before performing an in- place upgrade on a server that has specific roles, features, or settings—such as Dynamic Host Configuration Protocol (DHCP), Domain Name System (DNS), or AD DS—list those configurations. Then, check if those configurations support an in-place upgrade of Windows Server 2012. If so, refer to the detailed instructions for the specific roles, features, or settings on how to perform the in-place upgrade, including any issues or risks that might occur.

If any of these scenarios do not meet your organization’s requirements, then you should perform a migration to Windows Server 2012.

Benefits of Migrating

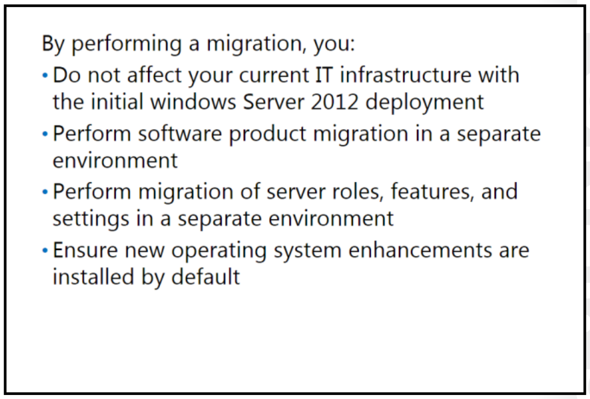

By performing a migration, you:

- Do not affect your current IT infrastructure with the initial windows Server 2012 deployment
- Perform software product migration in a separate environment
- Perform migration of server roles, features, and settings in a separate environment
- Ensure new operating system enhancements are installed by default

When deploying Windows Server 2012, some organizations should consider migration instead of an in-place upgrade. There can be risks that arise from an in-place upgrade, such as server unavailability or data being inaccessible. Therefore, your organization might choose to perform a migration because of the following benefits:

- You will deploy servers with the Windows Server 2012 operating system installed, and they will not affect the current IT infrastructure. Once you install Windows Server 2012, you can perform tests, such as drivers or system performance tests, before you introduce that server to the domain. In this way, the process of installation and testing is less likely to affect your current IT infrastructure.
- You will perform software product migration in a separate environment. For any software solution with an earlier Windows Server edition, you must refer to the product documentation for information about how to migrate that solution to Windows Server 2012. In some scenarios, software products that you are using are not supported for installation on Windows Server 2012, and you will require newer editions of those software products. In this case, by using migration, you can perform systematic installation of the operating system and the software products, in a separate environment. This ensures that the migration does not affect the availability of current services that the software provides.
- You will perform migration of server roles, features, and settings in a separate environment. As with the migration of software products, refer to the documentation on how to migrate the specific roles, features, or settings, such as DHCP, DNS, or AD DS, to Windows Server 2012. Again, migration enables you to perform systematic configuration in a separate environment, which means that the migration should not affect availability of server roles, features, and settings.
- New operating system enhancements are installed by default. When performing an in-place upgrade, for compatibility reasons, Windows Server 2012 is configured with settings for Windows Server 2008 or Windows Server 2008 R2. This means that many enhancements that Windows Server 2012 introduces, such as security, functionality, or performance enhancements, are not enabled by default. When performing migration, Windows Server 2012 deploys as a clean installation, with all new enhancements installed. This ensures that the operating system is more secure and has new functionalities installed by default.

Using MAP Toolkit to Plan for Upgrades and Migrations

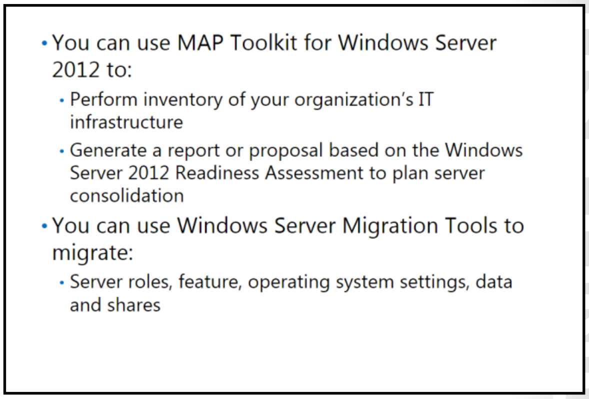

- You can use MAP Toolkit for Windows Server 2012 to:
  - Perform inventory of your organization's IT infrastructure
  - Generate a report or proposal based on the Windows Server 2012 Readiness Assessment to plan server consolidation
- You can use Windows Server Migration Tools to migrate:
  - Server roles, feature, operating system settings, data and shares

Organizations should consider using software tools to help them plan their upgrade and migration to Windows Server 2012.

The Microsoft Assessment and Planning Toolkit (MAP) analyzes the inventory of an organization's server infrastructure, performs an assessment, and then creates reports that you can use for upgrade and migration plans. MAP is available for Windows 8 and Windows Server 2012, and for other products, such as SQL Server 2012 and Microsoft Office 365™.

Use MAP to perform the following tasks:

- Inventory your organization's IT infrastructure. Based on the inventory, MAP displays a detailed report about which machines are capable of running Windows Server 2012, which machines are capable of running Windows Server 2012 with minimum system requirements, and which machines are not capable of running Windows Server 2012. MAP also recommends specific upgrades that ensure computers are capable of running Windows Server 2012.
- Generate a report or proposal based on the Windows Server 2012 Readiness Assessment. The report or proposal is a document that contains an Executive Overview, Assessment Results, Next Steps, and a worksheet summarizing Windows Server 2012 readiness for computers that are running Windows Server.
- Capture the performance metrics of the current IT infrastructure, to help plan consolidation and server virtualization. The performance assessment generates reports on performance and presents the server consolidation recommendations.
- Estimate server utilization based on that metric before and after the virtualization. You also can choose which current physical servers are the best candidates for virtualization, and the hosts on which you should place those virtual machines.

Reference Links: For more information about the Microsoft Assessment and Planning (MAP) Toolkit for Windows Server 2012, see http://go.microsoft.com/fwlink/?linkid=279918.

Planning For Server Consolidation

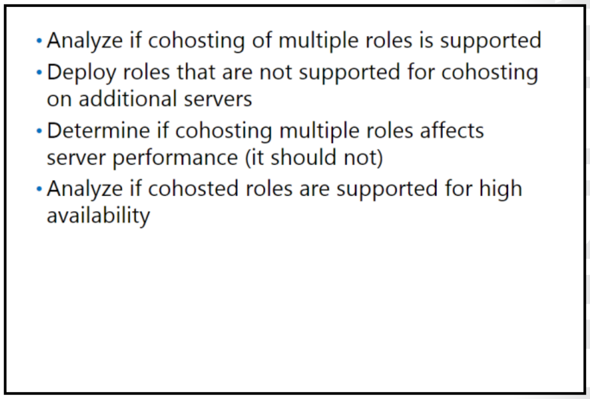

- Analyze if cohosting of multiple roles is supported
- Deploy roles that are not supported for cohosting on additional servers
- Determine if cohosting multiple roles affects server performance (it should not)
- Analyze if cohosted roles are supported for high availability

When Deploying Windows Server 2012, you should plan your placements of server roles, such as AD DS, DNS, and DHCP. Organizations should consider cohosting multiple roles, where possible, to achieve the most economical solution. Virtualization is also considered as a consolidation of the server roles. You should not implement cohosting if it affects server performance or available disk space. Therefore, organizations should evaluate and test whether installing multiple server roles on a server would result in lower overall performance and disk usage.

Furthermore, organizations should evaluate the security risks of collocating server roles. For example, the server that hosts the root Active Directory Certificate Services role should not be collocated with other server roles and should be offline most of the time.

Smaller organizations should consider the following best practices:

- Plan which server roles you need. If the operating system supports cohosting of those roles on one server, then multiple roles can be installed and cohosted on a single server. If cohosting multiple server roles on one physical server affects the performance of the physical server, then administrators should not cohost the server roles, and should install server roles on different physical servers.
- If the operating system on a physical host does not support that multiple server roles are cohosted, then administrators should deploy server roles on multiple physical servers.

Medium and large organizations should consider the following performance and high-availability issues when cohosting:

- If you are cohosting multiple roles on a single server, there might be performance issues because of the large number of client computers that are connecting to that server. In this situation, organizations should consider adding multiple servers that cohosts the same multiple roles. They also should consider relocating some of the roles from the first server to the other physical servers.
- High availability configurations of roles have specific requirements and setting, which might not support cohosting of multiple roles. In this situation, organizations should have a high availability solution for one server role, but must locate remaining roles on other servers.

Planning For Cloud Server Deployments

Organizations have the option to deploy servers with the Windows Server 2012 operating system, and then host them on a cloud platform, such as Windows Azure™. Windows Azure is a cloud platform that enables organizations to build, deploy, host, and manage applications and servers that are in Microsoft data centers. Windows Azure offers infrastructure as a service (IaaS). IaaS enables Microsoft to provide virtual-machine hosting, networking, and storage, and enables organizations create virtual machines, deploy an operating system, and create custom applications.

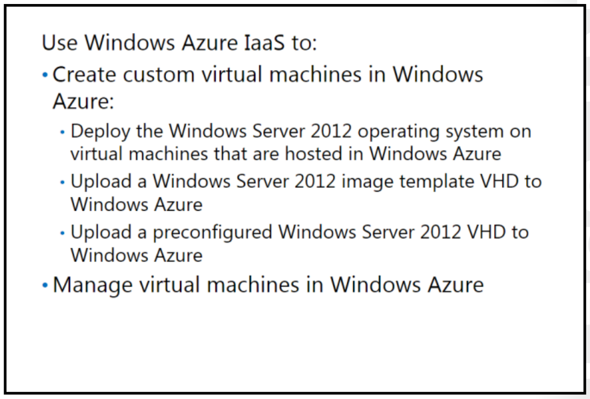

Use Windows Azure IaaS to:

- Create custom virtual machines in Windows Azure:
  - Deploy the Windows Server 2012 operating system on virtual machines that are hosted in Windows Azure
  - Upload a Windows Server 2012 image template VHD to Windows Azure
  - Upload a preconfigured Windows Server 2012 VHD to Windows Azure
- Manage virtual machines in Windows Azure

You can perform the following tasks by using IaaS:

- Create custom virtual machines. You can create different virtual machines that are running Windows Server 2012 and that you host on Windows Azure by using following deployment methods:
  - Deploy the Windows Server 2012 operating system on virtual machines that Windows Azure hosts. You can use a platform image that is available from the Windows Azure Management Portal.
  - Upload a Windows Server 2012 image template virtual hard disk (VHD) to Windows Azure. You can upload a VHD file on which there is an operating-system image template. This file would contain the Windows Server 2012 operating system. The template file has no preconfigured settings, and you can create it by using the System Preparation Tool (Sysprep). Once you upload the VHD image to Windows Azure, you can start configuring settings, such as computer name and network settings.

  - Upload a preconfigured Windows Server 2012 VHD to Windows Azure. You can use your own preconfigured image in a VHD file format, and upload the image to create a virtual machine in Windows Azure

- Manage virtual machines in Windows Azure. Once a virtual machine is running on the Windows Azure platform, you can use management tools from the Windows Azure portal to manage the hosted virtual machines.

Demonstration: Using the Microsoft Assessment and Planning Toolkit

In this demonstration you will see how to:

- Review the MAP options.
- Perform an inventory assessment by using MAP.
- Review the inventory from a sample database.

Demonstration Steps

Review the MAP options

1. On LON-CL1, run the **Microsoft Assessment and Planning Toolkit**.
2. In the **Microsoft Assessment and Planning Toolkit** console, review the default window that displays the **Overview** page.
3. In the Microsoft Assessment and Planning Toolkit console, in the left pane, select **Cloud**, and then review the readiness information for the different cloud scenarios.
4. In the Microsoft Assessment and Planning Toolkit console, in the left pane, click **Desktop**, and review the readiness information for the different desktop scenarios.
5. Repeat step 4 for all remaining items in the left pane: **Server**, **Desktop Virtualization**, **Server Virtualization**, **Database**, **Usage Tracking**, and **Environment**.

Perform inventory

1. On LON-CL1, in the Microsoft Assessment and Planning Toolkit console, in the left pane, select **Overview**, and then in the **Overview** page, create an inventory database named **INVENTORY**.
2. On **Overview** page, select **Perform an inventory**.
3. In **Inventory and Assessment Wizard** window, perform the following steps:
   a. On the **Inventory Scenarios** page, select the following check boxes:
      - **Windows computers**
      - **Exchange Server**
      - **Lync Server**
      - **SQL Server**
      - **Windows Azure Platform Migration**

   b. On the **Discovery Methods** page, select **Use Active Directory Domain Services**, **Use Windows networking protocols**, and **Scan an IP address range**.

   c. On the **Active Directory Credentials** page, in the **Domain** field, enter **Adatum.com**. In the **Domain Account** field, enter **Adatum\Administrator**, and then in the **Password** field, type **Pa$$w0rd**, and on the next two pages accept the default settings.

d. On the **Scan an IP Address Range** page, enter the range from **172.16.0.1**, to **172.16.0.100**.

e. On the **All Computers Credentials** page, accept the default settings.

f. On the **Summary** page, review the inventory options, and then cancel the wizard.

Note: You cancel the inventory procedure because the lab does not contain an environment with older operating systems for MAP to discover. In the next step, you review the test inventory that you import from the sample database in MAP.

Review the MAP inventory from a sample database

1. In the Microsoft Assessment and Planning Toolkit console, from the **File** menu, select **Manage Databases**.
2. In Microsoft Assessment and Planning Toolkit window, import the sample database using the following steps:

a. Select **Manage**.

b. Import the sample database located in following path: In the **File name** field, type **C:\Program Files\ Microsoft Assessment and Planning Toolkit\Sample \MAP_SampleDB.bak**.

c. In the **Database Name** field, type **MAPDEMO**.

d. In the **Microsoft Assessment and Planning Toolkit** window, choose an option **Use an existing database**, and select **MAPDEMO** database.

3. In the Microsoft Assessment and Planning Toolkit console, review the default window that displays the **Overview** page that includes inventory information from the sample database. Refresh the window in **Overview** page, if necessary.
4. In the Microsoft Assessment and Planning Toolkit console, in the left pane, click **Cloud**, and then review the readiness information on the different cloud scenarios that displays with inventory information from the sample database.
5. In the Microsoft Assessment and Planning Toolkit console, in the left pane, click **Desktop**, and review the readiness information on the different desktop scenarios that displays with inventory information from the sample database.
6. Repeat step 4 for all remaining items in the left pane: **Server**, **Desktop Virtualization**, **Server Virtualization**, **Database**, **Usage Tracking**, and **Environment**.

Lesson 2
Creating a Server Upgrade and Migration Plan

Organizations should plan to spend time creating a server upgrade and migration plan. Planning is critical for organizations that are considering new operating-system deployments. There are different elements that affect the planning for a new operating-system deployment, such as analyzing current IT infrastructure, choosing an operating-system edition, creating an upgrade or migration strategy, and creating a strategy for backup, restoring, monitoring, and maintaining the operating system.

There also are additional steps that you should consider as part of the planning process, such as operating-system licensing and activation, and determining which roles you can migrate, which you can cohost, and which you can consolidate into a virtual environment.

Lesson Objectives

At the end of this lesson, you will be able to:

- Explain why it is important to develop a proper deployment plan.
- Describe Windows Server 2012 volume licensing and activation considerations.
- Explain how to plan a suitable volume-activation mechanism.
- Explain how to implement server migrations.
- Explain how to migrate servers across domains.
- Explain how to deploy cloud servers.

Developing a Deployment Plan

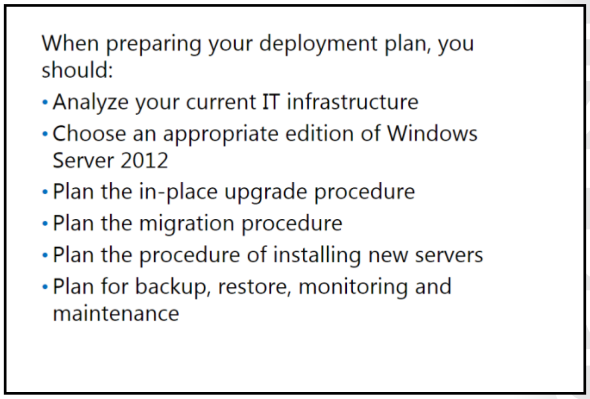

When preparing your deployment plan, you should:

- Analyze your current IT infrastructure
- Choose an appropriate edition of Windows Server 2012
- Plan the in-place upgrade procedure
- Plan the migration procedure
- Plan the procedure of installing new servers
- Plan for backup, restore, monitoring and maintenance

Before you introduce Windows Server 2012 into your organization, you should develop a deployment plan. A deployment plan helps ensure that when you upgrade, migrate, or install new servers, the process runs smoothly and without service interruption for your users. The output of your deployment plan should include detailed documentation and a checklist of analysis, preparation, installation, and post-installation activities.

When preparing your deployment plan, you should:

- Analyze your current IT infrastructure. By analyzing your current IT infrastructure, you can assemble detailed information on current IT resources. This enables you to analyze how your IT environment can better align to your organization’s business requirements.
- Select the appropriate edition of Windows Server 2012. Developing a deployment plan helps you choose the appropriate edition of Windows Server 2012, and helps you determine whether to perform a full graphical user interface (GUI) installation or server-core deployment. You also can plan whether to deploy your Windows Server 2012 solutions in a physical or virtual environment.

- Plan the in-place upgrade procedure. The in-place upgrade requires preparation, testing, and post-installation steps that you must include in your planning, so that you experience a minimal amount of downtime.
- Plan the migration procedure. Migration enables you to perform many installation steps in a separate environment. However, this process also requires preparing for planning, testing, and post-installation steps.
- Plan to install new servers. Some organizations plan to deploy Windows Server 2012 on new servers. This requires preparing a detailed checklist of the activities that a successful deployment requires.
- Plan for backup, restore, monitoring, and maintenance. After deployment, every software product or server role that runs on Windows Server 2012 must have a plan and strategy for backup, restore, monitoring, and maintenance.

The deployment plan should not be a static document. You should evaluate and update your deployment plan every time you initiate a new deployment process, and include steps and activities from previous deployment experiences. Improving your deployment plan regularly streamlines and improves your deployment process.

Windows Server 2012 Licensing and Activation

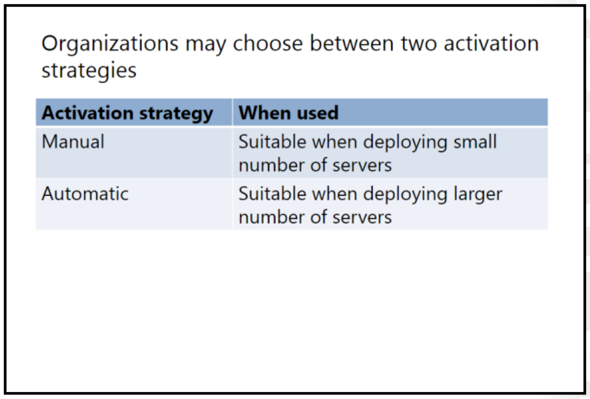

Organizations may choose between two activation strategies

Activation strategy	When used
Manual	Suitable when deploying small number of servers
Automatic	Suitable when deploying larger number of servers

To ensure that your organization has the proper licenses, and to receive notices for product updates, you must activate every copy of Windows Server 2012 that you install. Windows Server 2012 requires that you activate the operating system after installation. This ensures that the products are licensed and that you receive important update information. Unlike earlier Windows Server versions, there no longer is an activation grace period. If you do not activate Windows Server 2012, you cannot customize your operating system. There are two general activation strategies:

- Manual activation. This strategy is suitable when you deploy a small number of servers.
- Automatic activation. This strategy is suitable when you deploy a larger numbers of servers.

Manual Activation

When you use manual activation, you must enter the product key. Microsoft or an administrator performs the activation over the phone or through a special clearinghouse website.

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

OEM keys are a special type of activation key that a manufacturer receives, and which enable automatic activation when a computer turns on. You typically use this type of activation key with computers that are running Windows client operating systems, such as Windows 7 and Windows 8. You rarely use OEM keys with computers that are running Windows Server operating systems.

Automatic Activation

Performing activation manually in large-scale server deployments can be cumbersome. Microsoft provides a method of activating large numbers of computers automatically, without having to enter product keys manually on each system. In earlier editions of the Windows Server operating system, you could use the Key Management Service (KMS) to perform centralized activation of multiple clients. In Windows Server 2012, the Volume Activation Services server role enables you to manage a KMS server through a new interface. This simplifies installing a KMS key on the KMS server.

When you install Volume Activation Services, you also can configure Active Directory–based activation. Active Directory–based activation enables automatic activation of domain-joined computers. When you use Volume Activation Services, each activated computer must contact the KMS server periodically to renew its activation status.

To activate multiple computers on networks that do not connect directly to the Internet, you use the Volume Activation Management Tool (VAMT) 3.0 in conjunction with Volume Activation Services. You can use VAMT to generate license reports and manage client and server activation on enterprise networks.

Reference Links: For more information on VAMT, see Introduction to VAMT at http://go.microsoft.com/fwlink/?LinkID=391881.

Discussion: Planning Volume Activation

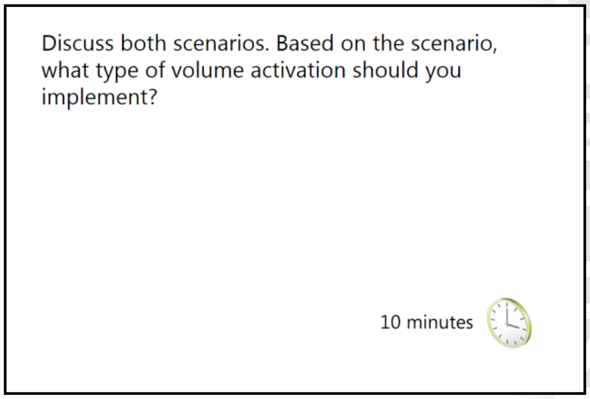

Discuss both scenarios. Based on the scenario, what type of volume activation should you implement?

10 minutes

To implement a volume-activation process, you must consider which activation type is most suitable for your organization. Not all companies have the same IT infrastructure, and therefore scenarios differ for each company. You should consider the two scenarios that are shown on the slide when planning your organization's volume activation process.

Question: Your organization's IT infrastructure consists of personal computers and servers that are running different editions of Windows client operating systems and Windows Server operating systems. Next month, your organization plans to deploy 500 Windows 8 client computers and 20 Windows Server 2012 servers. Because of a legacy application in the finance department, you must deploy 10 client computers that are running Windows 7 and two servers that are running Windows Server 2008 R2. What type of volume activation should you implement?

Question: Your organization's IT infrastructure was upgraded from different editions of Windows client operating systems and Windows Server operating systems to Windows 8 and Windows Server 2012, respectively. What type of volume activation should you implement?

Implementing Server Migrations

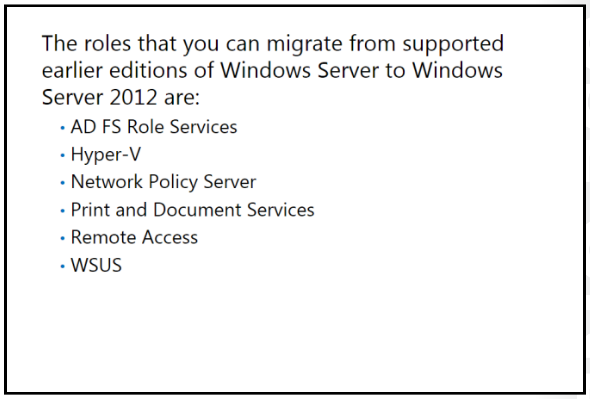

The roles that you can migrate from supported earlier editions of Windows Server to Windows Server 2012 are:

- AD FS Role Services
- Hyper-V
- Network Policy Server
- Print and Document Services
- Remote Access
- WSUS

When planning to migrate servers, you must create a list of the server roles that you want to migrate and the steps that each involves. For each server role that you plan to migrate, you should refer to the technical documentation and migration guides about how to perform the migration. When performing migration, you also may use the Windows Server Migration Tools, which are available with Windows Server 2012.

The roles that you can migrate include:

- Active Directory Certificate Services
- Active Directory Federation Services (AD FS) Role Services
- File and Storage Services
- DHCP
- DNS
- Hyper-V
- Network Policy Server
- Print and Document Services
- Remote Access
- Remote Desktop Services
- Cluster Role Services
- Windows Server Update Services (WSUS)

Note: You can migrate roles only from supported earlier Windows Server editions to Windows Server 2012.

Additional Reading: For more information about Windows Server Migration Tools, see Install, Use, and Remove Windows Server Migration Tools at http://go.microsoft.com/fwlink/?LinkID=391879.

Additional Reading: For more information about determining which roles and features to migrate, see the migration guides for both Windows Server 2012 and Windows Server 2012 R2 on following web page - Migrate Roles and Features to Windows Server 2012 at http://go.microsoft.com/fwlink/?LinkID=391880.

Implementing Server Migrations Across Domains

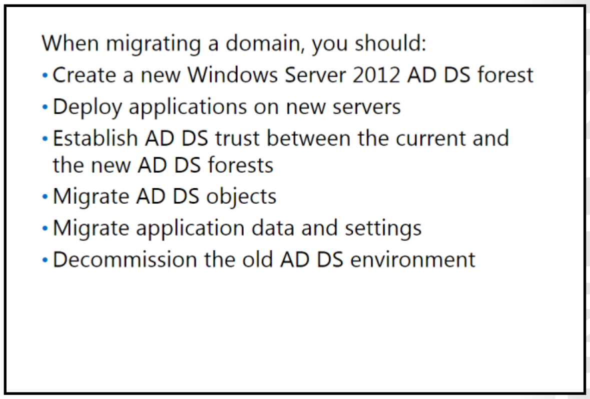

When migrating a domain, you should:

- Create a new Windows Server 2012 AD DS forest
- Deploy applications on new servers
- Establish AD DS trust between the current and the new AD DS forests
- Migrate AD DS objects
- Migrate application data and settings
- Decommission the old AD DS environment

Organizations may choose to deploy Windows Server 2012 in a new AD DS forest. In that scenario, administrators should plan the migration steps carefully to provide users with seamless access to data and services during the migration process. Once the migration is complete, administrators should begin the process of decommissioning and removing the infrastructure of the previous operating-system environment.

The process of migrating a server across domains includes:

- Creating a new Windows Server 2012 AD DS forest that is independent from the forest that is running a previous operating-system version.
- Deploying new servers that are running the Windows Server 2012 operating system.
- Deploying Microsoft applications, such as Exchange Server, SQL Server, and Microsoft SharePoint® server in the new AD DS forest.
- Deploying corporate custom applications or third-party applications in the new AD DS forest that the previous infrastructure environment used.
- Configuring DNS infrastructure in both forests.
- Establishing AD DS trust between the current and the new AD DS forests.
- Migrating AD DS objects, such as users, computers, groups, and mailboxes.
- Migrating application data and settings for Microsoft applications, corporate custom applications, and third-party applications.
- Ensuring that users can connect to corporate IT resources in the new AD DS forest.
- Decommissioning and removing the environment, based on previous operating system's AD DS forest.

Note: For each product and application that you plan to migrate to Windows Server 2012 AD DS forest, read the product documentation and best practices, including the supported migration procedures.
You will find this information on the web site of each of the product.

Deploying Servers in the Cloud

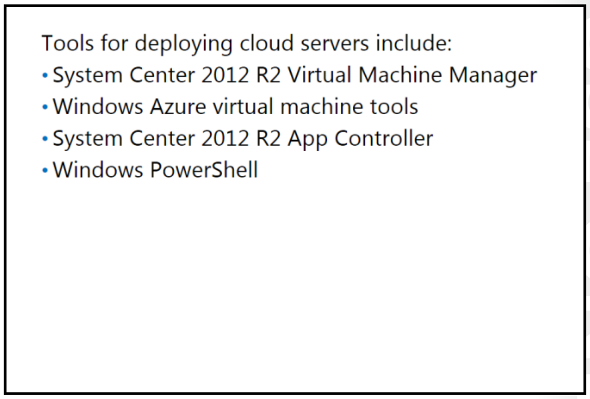
Tools for deploying cloud servers include:

- System Center 2012 R2 Virtual Machine Manager
- Windows Azure virtual machine tools
- System Center 2012 R2 App Controller
- Windows PowerShell

Organizations that decide to deploy Windows Server 2012 on public or private clouds require tools for deploying and managing Windows Server 2012 operating systems. Administrators use various tools depending on the type of cloud solution.

When deploying cloud servers, administrators may use the following tools:

- System Center 2012 R2 Virtual Machine Manager. System Center 2012 R2 Virtual Machine Manager (VMM) provides administrators with a single administrative tool for deploying virtual servers and managing a virtualization infrastructure. A virtualization infrastructure can include multiple resources, such as hosts, virtual machines, storage, networks, and libraries. You also can use VMM to update virtual servers.
- Windows Azure virtual machine tools. The Windows Azure web portal includes multiple tools for creating and managing virtual machines that are hosted on the Windows Azure cloud platform. Windows Azure virtual machine tools enable you to perform the following tasks:
    - Create custom virtual machines.
    - Create a virtual machine quickly.
    - Attach a disk to, or detach a disk from, a virtual machine.
    - Upload a Windows Server VHD.
    - Load balancing virtual machines.
    - Manage availability of virtual machines.
- System Center 2012 R2 App Controller. System Center 2012 R2 App Controller is an application that administrators can use to deploy and manage services across the Microsoft private cloud services and the Microsoft public cloud services, such as Windows Azure. App Controller has a web-based interface that enables administrators to manage services rather than servers.
- Windows PowerShell®. Administrators can use Windows PowerShell for automating tasks when deploying cloud servers in a private cloud environment or in public cloud environments based on Windows Azure.

Note: System Center 2012 R2 Virtual Machine Manager and System Center 2012 R2 App Controller are discussed in more detail in Module 3, "Planning and Deploying Servers Using Virtual Machine Manager."

Organizations that plan to migrate their services to the Windows Azure should consider the following steps:

- Analyze business requirements, define applications that will be migrated to Windows Azure, and define performance and scalability needed.
- Migrate applications to Windows Azure and perform detailed testing on the applications.
- Migrate data to Windows Azure, perform detailed testing of the application and data functionality, and perform optimization if needed.
- Once you complete optimization, create procedures for monitoring and managing applications and servers that are hosted in Windows Azure.

Lesson 3
Planning for Virtualization

Windows Server 2012 is a cloud-ready operating system. Organizations can benefit from Windows Server 2012 by using a physical server to create their own private cloud, or by consolidating part of their IT infrastructure in a virtual environment. Organizations should consider the ROI from virtualization, including reduced power consumption and licensing costs, and more efficient server utilization and the flexibility of a virtual-machine deployment.

In this lesson, you will learn how to plan the deployment of your operating-system infrastructure in a virtual environment. You will learn what your company should consider when implementing virtualization, and about the Windows Server 2012 licensing model. Finally, you will learn how to plan your deployment, using the guidelines for configuring Hyper-V hosts and for designing virtual machines for different types of applications.

Lesson Objectives

At the end of this lesson, you will be able to:

- Explain the considerations for implementing virtualization in your organization.
- Describe the virtual licenses that Windows Server 2012 includes.
- Explain how to apply guidelines for configuring Hyper-V host computers.
- Describe the considerations for virtualizing common applications.
- Explain how to apply guidelines for designing virtual machines for applications.
- Explain how to choose between virtualization and physical deployments.

Considerations for Implementing Virtualization

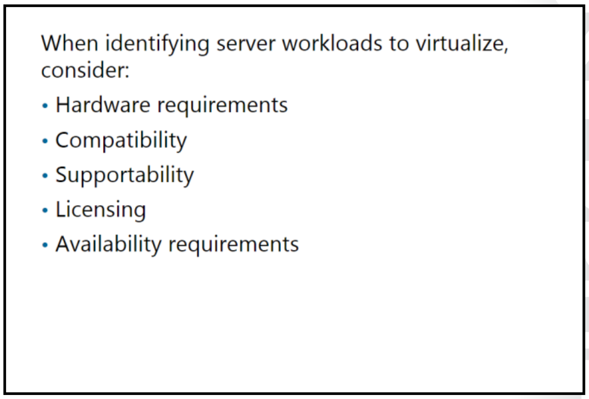

When identifying server workloads to virtualize, consider:

- Hardware requirements
- Compatibility
- Supportability
- Licensing
- Availability requirements

You can use virtualization to address many business and IT requirements, but you cannot virtualize all servers and applications. Before implementing virtualization, you need to identify those applications and servers within your organization that are the best virtualization candidates.

You must consider several factors when choosing whether to virtualize server workloads, including:

- Hardware requirements. Typically, virtual machines require approximately the same resources as a physical server. For example, say you have a physical server that utilizes 1 GB of RAM. You should expect the virtual machine to use the same amount of RAM, assuming that it runs the same operating system and applications as the physical server.

Note: When you plan resource utilization on your host computer, remember that the host computer requires additional resources for running the virtual machine. For each additional GB of memory that you assign to the virtual machine, there is a potential overhead of 8 MB on the host machine. For example, if the virtual machine requires 1GB of RAM, there is a potential overhead on the host computer of 32 MB for the virtual machine.

In some cases, a server workload may require hardware resources that make it impractical to deploy the workload on a virtual machine. For example, if the server workload requires more than half of the hardware resources that are available on a virtualization host, there may not be any server consolidation benefit.

Note: Ensure that you are using the actual hardware utilization rather than the actual physical hardware when evaluating the virtual machine's hardware requirement. You can deploy a physical server that is using only 5 percent of its current hardware resources on a virtual machine that has much lower hardware resources.

- Compatibility. You also must determine whether the application can run in a virtualization environment. Business applications range from simple executables to complex, distributed, multitier applications. You need to consider the requirements for the specific components of distributed applications, such as the need for communication with other infrastructure components, or requirements for direct access to the system hardware.

Applications and services that have specific hardware or driver requirements generally are not good virtualization candidates. An application may not be a good candidate for application virtualization if it contains low-level drivers that require direct access to the system hardware. This may not be possible through a virtualization interface and it can affect performance negatively.

- Supportability. You need to evaluate whether a virtualized environment will support the operating system and the application. You can verify vendor support policies for operating-system and application deployments that use the virtualization technologies.
- Licensing. You need to evaluate whether you can license the application for use in a virtual environment. Reduced licensing costs of multiple applications or operating systems could result in significant savings, which supports a strong financial reason for using virtualization.
- Availability requirements. Most organizations have some applications that must be made highly available for users. Some applications provide built-in options for enabling high availability, while you cannot easily make other applications highly available outside of a virtual machine environment. When considering whether to virtualize a server, evaluate whether the application has high availability options, whether the virtual-machine environment supports those options, and whether you can use failover clustering to make the virtual machine highly available.

The goal in most organizations is to utilize all servers adequately, regardless of whether they are physical or virtual servers. You can utilize some server roles fully, such as SQL Server or Exchange Server Mailbox servers, by deploying additional SQL Server instances or moving more mailboxes to the server.

Which Virtual Licenses are Included?

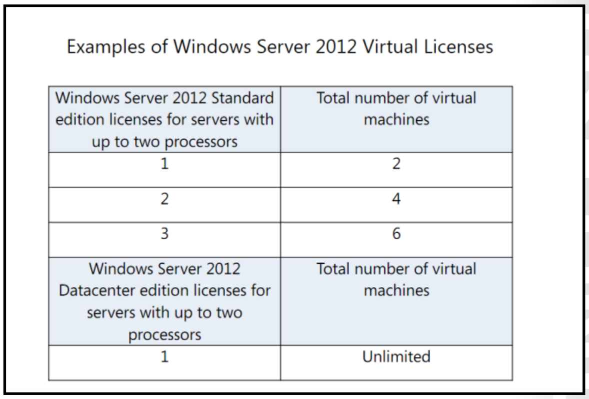

Examples of Windows Server 2012 Virtual Licenses

Windows Server 2012 Standard edition licenses for servers with up to two processors	Total number of virtual machines
1	2
2	4
3	6
Windows Server 2012 Datacenter edition licenses for servers with up to two processors	Total number of virtual machines
1	Unlimited

Windows Server 2012 provides organizations with capabilities for running their IT infrastructure based on virtualization technologies or for connecting and interoperating their IT infrastructure with external cloud services. The Windows Server operating-system edition that you deploy depends on the level of virtualization that you plan to implement.

There are two Windows Server 2012 editions that include virtual licenses:

- The Windows Server 2012 Standard operating system includes licenses for two virtual machines for a server that has up to two processors. An additional license is necessary for each additional two processors.
- The Windows Server 2012 Datacenter operating system includes licenses for an unlimited number of virtual machines for a server that has up to two processors. An additional license is necessary for each additional two processors.

Windows Server 2012 Standard edition is the choice for organizations that require an environment based mostly on physical servers, and that uses virtualization technology less frequently. Organizations with an IT infrastructure that predominantly virtual should use the Windows Server 2012 Datacenter edition.

If organizations must run more than two virtual machines on the same hardware, but still do not require an unlimited number of virtual licenses, they might choose to install a second instance of Windows Server 2012 Standard edition on the server. In this scenario, they would be licensed to run four virtual machines on the same hardware.

The following table lists some of the virtual licenses that are available.

Windows Server 2012 Standard edition licenses for a server with up to two processors	**Total number of virtual machines**
1	2
2	4
3	6
Windows Server 2012 Datacenter edition licenses for a server with up to two processors	Total number of virtual machines
1	Unlimited

Windows Server 2012 R2 introduces Automatic Virtual Machine Activation. It enables organizations to install virtual machines on activated Windows servers without entering product keys for virtual machines. Automatic Virtual Machine Activation activates the virtual machine when the virtual machine starts up by using the virtualization server’s license. Automatic Virtual Machine Activation also provides reporting and tracking data for the license usage of virtual machines.

Note: This topic provides information only on server licensing and not on client-access licenses. For most server-based applications, client computers that connect must have an appropriate client-access license.

Guidelines for Configuring Hyper-V Hosts

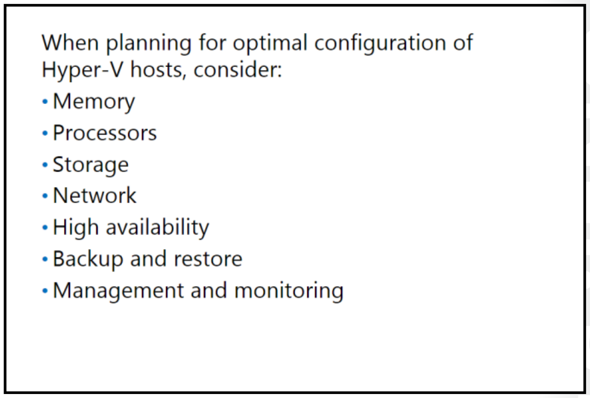

When planning for optimal configuration of Hyper-V hosts, consider:

- Memory
- Processors
- Storage
- Network
- High availability
- Backup and restore
- Management and monitoring

To optimize usage of Hyper-V hosts, you must create a detailed plan for the virtual machines that you plan to deploy. When you plan your Hyper-V host configuration, you should consider the number of virtual machines that you want to host, the type of the application or solution that you will deploy on each virtual machine, and the performance necessary for the applications that you are deploying on the virtual machines.

When you are planning for an optimal Hyper-V host configuration, you should consider:

- Memory. Plan the amount of memory for a Hyper-V host based on each virtual machine's requirements. If you plan to allocate more memory for some virtual machines, we recommend that you ensure that the Hyper-V host has sufficient memory.
- Processors. You should plan how many processors you require, and the type you need, according to the number and type of virtual machines.
- Storage. The storage solution should consist of disk drives that are fast enough to support the performance requirements of the applications that you deploy in the virtual environment. The available storage space also should be sufficient to support adding additional logical drives to the virtual machines, and extending the current logical drives, as necessary.
- Network. For better throughput of network communication, we recommend that you install more than one network adapter on the Hyper-V host. The number of network adapters depends on the installed applications and their requirements. We also recommend that you dedicate a separate network adapter if you have Internet Small Computer System Interface (iSCSI)–based storage.
- High availability. Hyper-V hosts manage multiple virtual machines, so any issue and risk that threatens Hyper-V host availability threatens all virtual machines. Therefore, you should plan a strategy for high availability of the Hyper-V host.
- Backup and restore. Even if you configure the Hyper-V host to be highly available, you must develop a backup and restore strategy to address any issues that the deletion or corruption of data cause.
- Management and monitoring. For an effective and optimized virtualization environment, we recommend that you use management and monitoring solutions, such as Microsoft System Center 2012 Operations Manager (Operations Manager) and System Center 2012 Virtual Machine Manager (VMM).

Guidelines for Designing Virtual Machines

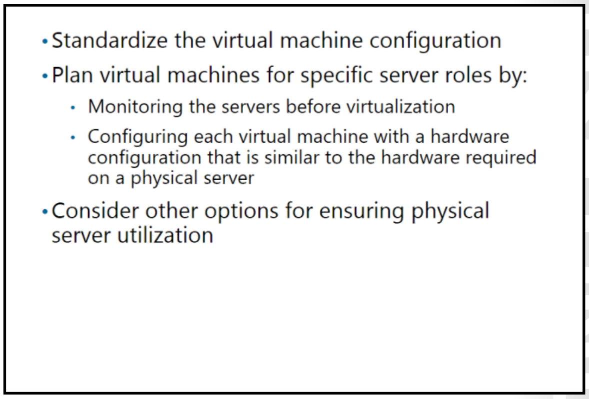

- Standardize the virtual machine configuration
- Plan virtual machines for specific server roles by:
  - Monitoring the servers before virtualization
  - Configuring each virtual machine with a hardware configuration that is similar to the hardware required on a physical server
- Consider other options for ensuring physical server utilization

One important goal when developing a virtualization strategy is to simplify and standardize the host computer and virtual machine configuration as much as possible. Consider the following general guidelines, which apply to all of your virtual machines:

- Develop a small number of standard virtual-machine builds. To streamline virtual-machine deployment and management, develop a set of standard virtual-machine builds. For example, consider creating a standard low-end server build, a medium server build, and a high-end server build. Assign a standard central processing unit (CPU) and memory configuration for each role. You also should consider configuring each virtual machine with a standard 50 GB system partition and providing additional disks on which to store data or install applications. Consider using SCSI controllers for all hard disks, other than the disk that contain the boot and system partition. Windows Server 2012 enables you to add new VHDs that you connect to a SCSI controller, without having to restart the server.

- Plan virtual machines for specific server roles. Although you should be able to configure most virtual machines with the same basic disk and operating-system configuration, the actual physical requirements for each virtual machine varies. For example, some virtual machines require significantly more RAM or CPU resources than others. To design the actual physical requirements for a virtual machine, you should:

  - Monitor the servers before virtualizing them. Collect performance data on the servers to evaluate how specific applications perform on physical servers. If an application uses a low percentage of a physical server's hardware resources, deploy a virtual server with significantly less capacity to run the same application.

  - Configure each virtual server with a hardware configuration that is similar to the hardware that the application requires on physical servers. Virtualizing a server does not change the hardware resources that the server requires.

- Consider other options for ensuring physical server utilization. One of the goals of server virtualization is to ensure that you utilize all servers adequately, regardless of whether they are physical or virtual. For example, you can utilize the SQL Server or Exchange Server Mailbox servers more fully by deploying additional SQL Server instances or by moving more mailboxes on to the server. When considering virtualization, also consider other options for fully utilizing the hardware.

Designing Virtual Machines for Applications

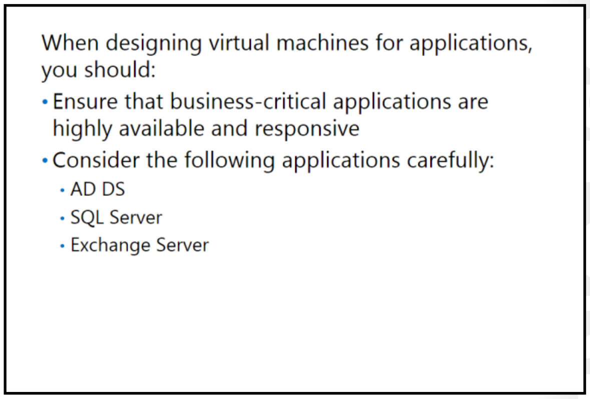

When designing virtual machines for applications, you should:

- Ensure that business-critical applications are highly available and responsive
- Consider the following applications carefully:
  - AD DS
  - SQL Server
  - Exchange Server

One of the most important Windows Server 2012 server virtualization benefits is that it provides an enhanced option for server consolidation. However, many organizations run business-critical applications, so it is essential that these applications are highly available and responsive.

The following are best practices for implementing virtualization for business-critical applications in AD DS, SQL Server, and Exchange Server.

Virtualizing AD DS

You should consider the following guidelines when you are virtualizing AD DS:

- You can install a Windows Server domain controller as a virtual machine, along with other application servers on a single physical Windows Server 2012 server.
- Windows Server 2012 introduces virtualized domain-controller cloning. In earlier Windows Server editions, domain controllers that were running within a virtual machine were unaware of their virtual state. This made performing processes such as cloning and restoring virtual-machine snapshots potentially dangerous, because changes could occur to the operating-system environment that the domain controller did not expect.
- Safe backup and restore. Rolling back to a previous snapshot of a virtualized domain controller is problematic because AD DS uses multimaster replication that relies on transactions being assigned numeric values, or update sequence numbers (USNs). The virtualized domain controller tries to assign USNs to prior transactions that have been assigned to valid transactions, which causes inconsistencies in the AD DS database. Windows Server 2003 and newer Windows Server operating systems implement a process that is known as USN rollback protection. It ensures that the virtualized domain controller does not replicate, and you must demote it forcibly or restore it manually. Windows Server 2012 now detects the snapshot state of a domain controller, and synchronizes or replicates the delta of changes, between a domain controller and its partners for AD DS and the SYSVOL. Furthermore, you can use snapshots without risk of permanently disabling domain controllers and requiring manual forced demotion, metadata cleanup, and repromotion.
- Ensure the maintenance of physical security for your host computer. The VHD that contains the virtualized domain controller stores very sensitive information. Compromising this data can create significant additional work for your organization.

Virtualizing SQL Server

You should consider the following recommendations when you are configuring virtual machines that run SQL Server:

- Plan to configure the hardware settings for the virtual machines the same as the hardware settings that you would configure on a physical server with the same workload.
- Plan virtual machine storage. One of the most critical components to ensure optimal performance for any SQL Server instance is to ensure that your storage system's size and configuration is correct. The storage hardware should provide sufficient I/O throughput and storage capacity to meet the current

and future needs of the planned virtual machines. Additionally, you should follow the recommended best practices for configuring disks for transaction logs and database storage.

- Provide adequate CPU capacity. In Windows Server 2012, you can achieve the same throughput on a virtual machine as on physical hardware, with only slightly increased CPU utilization.

Virtualizing Exchange Server

You can use a virtualization environment to run all Exchange Server 2010 or Exchange Server 2013 server roles.

You should consider the following guidelines when virtualizing Exchange Server servers:

- Use standard server sizing. From an application perspective, running Exchange Server on a guest virtual machine does not change the Exchange Server design requirements. The Exchange Server guest virtual machine must be the appropriate size to handle the workload.
- Configure appropriate storage. The storage that the Exchange Server virtual machine uses can be fixed virtual hard disk drives, SCSI pass-through storage, or iSCSI storage. As with SQL Server–based servers, pass-through storage provides the best performance. Dynamically expanding virtual disks and differencing drives are not supported for Exchange Server servers.
- Separate logical unit numbers (LUNs).You should use separate LUNs on redundant array of independent disks (RAID) arrays for the host operating system, each guest operating system disk, and all virtual machine storage. You should create separate LUNs for each database and set of transaction log files, as you would when working with physical servers.
- Configure adequate CPU resources. Exchange Server supports a ratio of virtual processors to logical processors that is no greater than two-to-one. For example, a dual-processor system that uses quad-core processors contains eight logical processors in the host system. On a system with this configuration, do not allocate more than 16 virtual processors to all guest virtual machines combined.
- High availability for servers that are running Exchange Server. Exchange Server 2010 provides several options for high availability. You can combine virtualization high-availability solutions, such as the Live Migration feature, with Exchange Server high-availability solutions, such as Database Availability Groups. However, in most cases, we recommend that you use the Exchange Server solution. High-availability virtualization solutions are not application-aware, while Exchange Server solutions are. For example, servers that are running Exchange Server can detect when a single database dismounts, and can mount the database automatically on another server. Hyper-V Live Migration cannot detect the status of individual databases or other services.
- Mailbox server performance. The most common performance issues for mailbox servers are disk and network I/O. Running mailbox servers in a virtual environment means that the virtual machines have to share this I/O bandwidth with the host machine and with other virtual machine servers that are deployed on the same host. If a single virtual machine is running on the physical server, the disk I/O and network I/O that are available to the virtual machine are almost equivalent to the I/O that is available to a physical server. However, a heavily utilized mailbox server can consume all the available I/O bandwidth, which will make it impractical to host additional virtual machines on that physical server.

Best Practices for Planning Application Virtualization

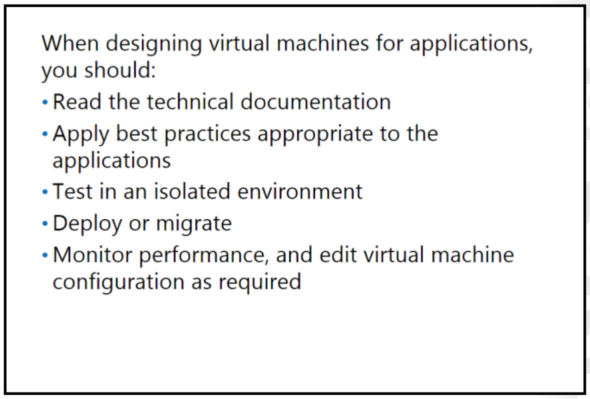

When designing virtual machines for applications, you should:

- Read the technical documentation
- Apply best practices appropriate to the applications
- Test in an isolated environment
- Deploy or migrate
- Monitor performance, and edit virtual machine configuration as required

When deploying applications on virtual machines, follow the technical documentation for that specific application. There may be different deployment procedures, depending on whether you are installing the application on a physical or virtual server.

The following is a list of best practices for deploying applications on virtual machines:

- Read the technical documentation. Each product has detailed technical documentation on supported deployment scenarios. Not every product that is recommended for physical environments is supported in a virtual environment. If it is, it might require additional configuration for virtual environments.
- Apply best practices for specific applications. Configuration of applications that are running in a virtual environment depends on that organization’s specific business requirements and the application’s instructions. Application vendors often publish common best practices, and usually update their own best practices during the product life cycle. Therefore, check regularly for updated documentation.
- Test in an isolated environment. Before deploying the application in production, we highly recommend that you test it in an isolated virtual environment. You then can resolve any potential issues without disrupting current services. You also can optimize virtual-machine and application settings based on the tests that you perform.
- Migrate or convert to a virtual environment. If your software is deployed in a physical environment, follow the technical documentation on how to migrate it to a virtual environment. In some scenarios, using virtualization software solutions, such as VMM, can help you convert the physical machine to virtual machine, and then move it to another physical host. However, be sure that you verify that the software vendor supports this scenario.
- Monitor performance, and edit virtual machine configuration, as necessary. Once you deploy your application successfully in a virtual environment, you should monitor its performance and utilization. This enables you to provide additional hardware resources when necessary or resolve any potential issues. By using specialized software solutions that provide a centralized monitoring dashboard, such as Operations Manager, you can monitor your virtual environment, and troubleshoot any potential warnings or alert messages.

Discussion: Choosing Between Virtual and Physical Deployments

Organizations might choose to run their business application or infrastructure services in physical or virtual environments. The choice of physical or virtual deployment depends on corporate business requirements and on corporate strategy for development of current and future IT infrastructure. For example, medium and enterprise organizations that want to consolidate hardware resource, and save on power consumption, might choose to invest in virtualization. Smaller organizations might only migrate to Windows Server 2012 deployed in a physical environment or use cloud-based solutions.

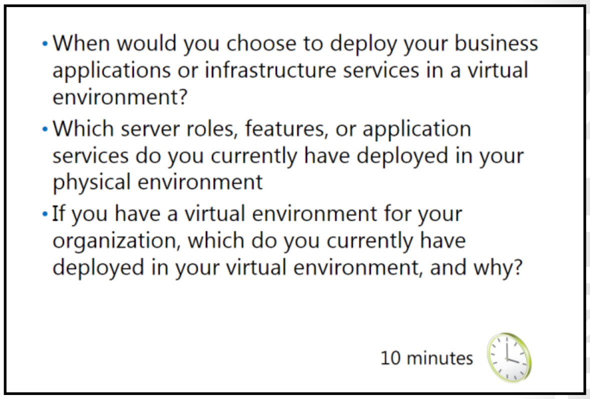

- When would you choose to deploy your business applications or infrastructure services in a virtual environment?
- Which server roles, features, or application services do you currently have deployed in your physical environment
- If you have a virtual environment for your organization, which do you currently have deployed in your virtual environment, and why?

10 minutes

Question: When would you choose to deploy your business applications or infrastructure services in a virtual environment?

Question: Which server roles, features, or application services do you deploy currently in your physical environment?

Question: If your organization has a virtual environment, what do you deploy currently in your virtual environment, and why?

Lab: Planning a Server Upgrade and Migration

Scenario

The head office for A. Datum Corporation, based in London, England, has a mix of Windows Server 2008 and Windows Server 2008 R2 servers. The AD DS environment is based around Windows Server 2008 R2 domain controllers.

The regional hub sites (Toronto and Sydney) predominantly have a Windows Server 2008–based infrastructure, while the smaller, regional branches and distribution centers have Windows Server 2008 servers that support line-of-business (LOB) applications.

The management team at A. Datum has stipulated that the new IT infrastructure must be more efficient with respect to hardware utilization and power consumption. This suggests that consolidating server roles and virtualizing workloads is an important part of the server upgrade and migration plan.

You have been involved with early discussions surrounding A. Datum's likely acquisition of two companies. This has not been formally announced, but is likely to occur in the coming months. You understand that one of these organizations has no Windows Server infrastructure in place, and you may be required to develop a strategy for how best to deploy Windows Server 2012 in that environment. The second potential acquisition has a Windows Server 2008 forest.

However, before the server deployment can begin, you must plan a server upgrade and migration strategy. This plan will involve determining which server roles will migrate, which you can consolidate onto a single server, and whether a given workload is suitable for virtualization. Initially, you will perform this analysis on the servers in London and its perimeter network only. You plan to analyze Toronto and Sydney later.

During this process, you must remain focused on optimizing server resources and capacity. In addition to these mainly technical considerations, you must consider what implications there may be with respect to server licensing.

Objectives

After completing this lab, students will be able to plan a server upgrade and migration strategy.

Lab Setup

Estimated Time: 50 minutes

Virtual machines	20413C-LON-DC1 20413C-LON-CL1
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, on the **Start** screen, click **Hyper-V Manager**.
2. In Hyper-V Manager, click **20413C-LON-DC1**, and then 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 20413C-LON-CL1.

Exercise 1: Planning a Strategy for Server Upgrade and Migration

Scenario

All current servers are installed with either single or dual processors. The table below identifies these servers and their current configuration.

Server name	Operating-system version	Roles installed	Processor details/processor utilization	Memory installed/average utilization
Head Office				
LON-EX1	Windows Server 2008 R2	Exchange Server 2010: Client Access and Hub Transport	Dual processor/75%	32 GB/16 GB
LON-EX2	Windows Server 2008 R2	Exchange Server 2010: Mailbox	Dual processor/75%	32 GB/24 GB
LON-EX3	Windows Server 2008 R2	Exchange Server 2010: Mailbox	Dual processor/75%	24 GB/18 GB
LON-SQL1	Windows Server 2008	SQL Server (clustered)	Dual processor/80%	32 GB/28 GB
LON-CA	Windows Server 2008 R2	Active Directory Certificate Services (AD CS): Enterprise Root CA	Single processor/25%	16 GB/4 GB
LON-DC1	Windows Server 2008 R2	AD DS	Single processor/60%	8 GB/3 GB
LON-DC2	Windows Server 2008 R2	AD DS	Single processor/60%	8 GB/3 GB
LON-DC3	Windows Server 2008 R2	AD DS	Single processor/60%	8 GB/3 GB
LON-INF1	Windows Server 2008	DHCP, DNS	Single processor/50%	8 GB/4 GB
LON-INF2	Windows Server	DHCP, DNS	Single processor/50%	8 GB/4 GB

Server name	Operating-system version	Roles installed	Processor details/processor utilization	Memory installed/average utilization
	2008			
Perimeter network				
LON-PER-NS1	Windows Server 2008	DNS	Single processor/25%	8 GB/2 GB
LON-PER-NS2	Windows Server 2008	DNS	Single processor/25%	8 GB/2 GB
LON-RADIUS	Windows Server 2008 R2	Network Policy Server	Single processor/35%	8 GB/4 GB
LON-WEB	Windows Server 2008	Multiple websites hosted with Internet Information Services (IIS)	Dual processor/75%	16 GB/10 GB
LON-VPN1 LON-VPN2	Windows Server 2008 R2	Routing and Remote Access Service (RRAS)	Single processor/25%	6 GB/3 GB
LON-EX-EDGE1	Windows Server 2008 R2	Exchange Server 2010 Edge Transport server	Dual processor/40%	16 GB/8 GB

The following diagram shows the placement of these servers:

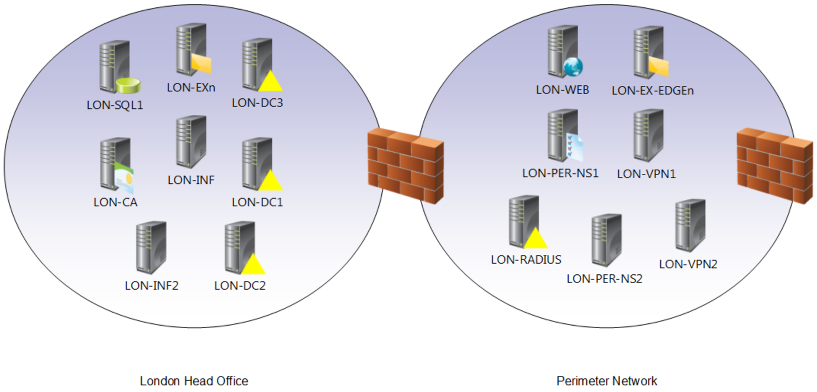

Network Access Services Strategy	
Document Reference Number: BS0901/1	
Document Author Date	Brad Sutton 7th September

Network Access Services Strategy

Requirements Overview

Design a server upgrade and migration strategy to support the following objectives:

- Ensuring that the new IT infrastructure is more efficient, in terms of hardware utilization and power consumption.
- Consolidating server roles and virtualizing workloads is an important part of the server upgrade and migration plan.
- Developing a strategy for deploying Windows Server 2012 in a scenario in which we acquire another company that has a Windows Server 2008 forest.
- Determining which server roles will migrate.
- Determining which server roles to consolidate onto a single server, and whether a given workload is suitable for virtualization.
- Determining the host server operating system that you will use to support the virtualized workloads.
  - Initially, you plan to perform this analysis only on the servers located in London and its perimeter network. You plan to analyze Toronto and Sydney later.
- Considering the implications of server licensing.

Additional Information

- A. Datum Corporation has a head office based in London, England, and has a mix of Windows Server 2008 R2 and Windows Server 2008 servers.
- The AD DS environment is based around Windows Server 2008 R2 domain controllers.
- The regional hub sites (Toronto and Sydney) predominantly have a Windows Server 2008–based infrastructure, while the smaller, regional branches and distribution centers have Windows Server 2008 servers that support line-of-business (LOB) applications.

Proposals

1. You plan to run the Microsoft Assessment and Planning Toolkit (MAP) to help you decide on a server consolidation strategy. What result do you expect to get from this tool?

2. Besides using the MAP, what would help you determine which machines you would move to the virtual environment?

3. What is your decision regarding virtualization of domain controllers?

4. What is your decision regarding virtualization of the LON-IF1 and LON-IF2 infrastructure servers?

Network Access Services Strategy

5. Do the virtualized machines require high availability?

6. What system resources, such as processors, memory, or disk space should you allocate to server roles?

7. What are the best virtualization candidates on the internal network, considering current physical server utilization and high availability requirements?

8. What are your plan's licensing considerations for internal network servers? Do these considerations have an impact on the host operating system?

9. What are the best virtualization candidates on the perimeter network, considering current physical server utilization and high availability requirements?

10. What are your plan's licensing considerations for perimeter network servers? How does this impact the selection of the host operating system?

11. How would you manage licensing and activation?

12. Are there any servers that you should not cohost?

13. Sketch your plan's relevant network areas.

The main tasks for this exercise are as follows:

1. Read the supporting documentation

2. Update the proposal document with your planned course of action

3. Examine the suggested proposals in the Lab Answer Key

4. Discuss your proposed solution with the class, as guided by your instructor.

▶ Task 1: Read the supporting documentation

- Read the supporting documentation in the lab exercise scenario.

▶ Task 2: Update the proposal document with your planned course of action

- Analyze the internal and perimeter networks separately. Do this because of the different security configurations and settings of each of those networks.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Examine the suggested proposals in the Lab Answer Key.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor.

1. What was your approach to the design plan?
2. Did your design plan differ from the suggested solution?

Results: After completing this exercise, you will have planned a server upgrade and migration successfully.

Exercise 2: Evaluating Candidates for Server Virtualization

Scenario

To gain familiarity with the MAP toolkit, you will install the toolkit and use a sample database to evaluate which servers are virtualization candidates.

The main tasks for this exercise are as follows:

1. Evaluate server virtualization candidates

2. To prepare for the next module

▶ Task 1: Evaluate server virtualization candidates

1. On LON-CL1, start the **Microsoft Assessment and Planning Toolkit**.
2. In Microsoft Assessment and Planning Toolkit window, import the sample database using the following steps:
   a. In the **Microsoft Assessment and Planning Toolkit** dialog box, select **Manage**.
   b. Import the sample database located in following path: In the **File name** field, type **C:\Program Files\ Microsoft Assessment and Planning Toolkit\Sample \MAP_SampleDB.bak**.
   c. In the **Database Name** field, type **MAPDEMO**.
   d. In the **Microsoft Assessment and Planning Toolkit** window, choose an option **Use an existing database**, and select **MAPDEMO** database.
3. In the Microsoft Assessment and Planning Toolkit window, run the Server Consolidation Wizard. Select **Windows Server 2012 Hyper-V**, and then click **Sample host**.
4. On the **Utilization Settings** page, type **75** in each field.
5. On the **Computer List** page, select all of the computers, and then complete the assessment.
6. On the **Summary** page, review the settings, and then click **Finish**. When the assessment process completes, click **Close**.

7. In the MAP console, on the **Server Virtualization** page, review the server consolidation information, and then run the **Server Virtualization Report**.
8. In **File Explorer**, locate and open the report.
9. At the bottom of the Microsoft Excel® workbook, click each tab, and then review the information in the report.
10. When finished, close **Excel**, and then close **File Explorer**.

▶ Task 2: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20413C-LON-CL1**.

Results: After completing this exercise, you will have installed MAP and used a sample database to evaluate which servers are virtualization candidates.

Question: Why would you want to use MAP when planning your upgrade and migration strategy?

Question: Why would you choose Windows Server 2012 Datacenter edition for virtualization and consolidation of both the A. Datum internal and perimeter networks?

Module Review and Takeaways

Best Practice:

When planning to deploy Windows Server 2012 in a physical or virtual environment, always consider high availability and backup/restore strategy for services or applications that run on that operating system. If you are running solutions in the private cloud, always ensure that you use management and monitoring tools, such as System Center 2012, to help the IT environment run efficiently. Additionally, ensure that you have a properly designed storage solution with appropriate size and performance for the virtual machines.

Review Question

Question: What are the key considerations that should guide your organization's strategy regarding different scenarios for Windows Server 2012 operating system deployment?

Real-world Issues and Scenarios

Question: Your organization has low usage of virtualization technologies. You have deployed the Windows Server 2012 Standard edition operating system that supports two instances of virtual machines. The management is concerned about future plans that require you to deploy new products in a virtual environment. They would like to have scalable and extensible solution without having to purchase additional licenses when deploying new products.

What strategy should the IT department suggest to the management?

Answer: The IT department should create a server deployment strategy that includes a hardware solution that is running on the Windows Server 2012 Datacenter edition. This enables the organization to deploy applications in a virtual environment and scale flexibly without requiring additional licenses.

Tools

Tool	Used for	Where to find it
Microsoft Assessment and Planning Toolkit (MAP)	Analyzing the inventory of an organization's server infrastructure, performs an assessment, and creates reports that you can then use when planning upgrades and migration.	Microsoft website: http://go.microsoft.com/fwlink/?linkid=279918

Module 2

Planning and Implementing a Server Deployment Strategy

Contents:

Module Overview

With the increase in the number of IT solutions in organizations, the number of physical and virtual servers has also increased. Due to this situation, operating system deployments take longer to complete and require valuable organizational resources. As a result, companies are looking for new ways to automate the server deployment process.

Before beginning to implement your automation process, you should design a proper deployment strategy and deployment method. A proper design will increase productivity and lower the time required for server deployment.

Objectives

After completing this module, you will be able to:

- Explain how to select an appropriate server deployment strategy.
- Explain how to implement an automated deployment strategy.

Lesson 1
Selecting an Appropriate Server Deployment Strategy

When developing a server deployment strategy, organizations base their strategies on different parameters such as business requirements, number of users, and size of their IT infrastructure. Smaller organizations will most likely perform a manual deployment of Windows Server® 2012.

There are two manual deployment methods. One method is to perform a retail installation by using a Windows Server 2012 DVD. The other method is to prepare a custom captured image from a *reference computer* (a computer that has operating system and applications already installed and configured by the IT administrator), and then manually configure the settings on each computer. For companies that have a dedicated IT staff, most deployments will require limited to no human interaction. Although these deployments make use of several external tools, such as the Microsoft Deployment Toolkit (MDT) 2013 and Microsoft® System Center 2012 R2 Configuration Manager (Configuration Manager), these lite-touch and zero-touch deployment methods provides return on investment in a very short time.

In addition, most organizations' server deployment needs differ, depending mostly on the organization size and the number of technologies that they will be running on Windows Server 2012. For example, some organizations will use automated deployment procedures for physical servers because most of their server infrastructure is installed on the physical environment. However, other organizations may look for a solution that will automate server deployment in a virtual environment, because their server infrastructure is virtualized. Therefore, you should propose a deployment automation strategy that will be the most appropriate for your organization.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the Windows® image file format and its benefits.
- Perform High Touch with Retail Media deployments.
- Perform High Touch with Standard Image deployments.
- Perform lite-touch high-volume deployments.
- Perform zero-touch high-volume deployments.
- Describe your current deployment strategy.
- Describe the advantages of automated deployment methods.
- Select an image-based deployment strategy.

What Is the Windows Image File Format?

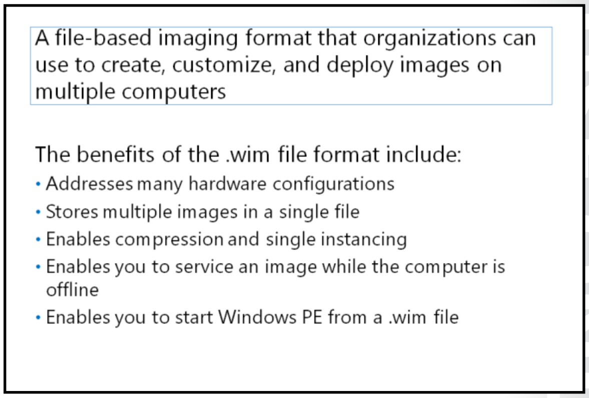

A file-based imaging format that organizations can use to create, customize, and deploy images on multiple computers

The benefits of the .wim file format include:

- Addresses many hardware configurations
- Stores multiple images in a single file
- Enables compression and single instancing
- Enables you to service an image while the computer is offline
- Enables you to start Windows PE from a .wim file

When you make an image of a hard drive, the image is an exact replica of your hard drive at the time the image was taken. Organizations use different technologies for imaging computers' hard drives to make the deployment process easier. The Windows image file format (or the .wim file) is a file-based imaging format that was first introduced in Windows Vista® and the Windows Server 2008 operating system. Organizations can use .wim images to create, customize, and deploy images on multiple computers, thus simplifying operating system deployments on multiple computers. The benefits of the .wim imaging format include the following:

- A single .wim file can address many hardware configurations. The .wim file does not require that the destination hardware match the source hardware. This helps to reduce the number of images greatly, and provides the advantage of having only one image that addresses several hardware configurations.
- The .wim file can store multiple images in a single file. This is useful because you can store images with and without core applications, in a single image file. Another benefit is that you can mark one of the images as bootable, which means that you can start a machine from a disk image that a .wim file contains.
- The .wim file enables compression and single instancing. This reduces the size of image files significantly. Single instancing is a technique that enables multiple images to share a single copy of files that are common between the instances.
- The .wim file enables you to service an image while the computer is offline. You can add or remove operating system elements, files, updates, and drivers without creating a new image. For example, to add an update to a Windows 8 image, you must start the master image, add the update, and then prepare the image again. With Windows Server 2012 and Windows 8, you can mount the image file, and then slipstream the update into the image file without the need to start or recapture the master image.
- The .wim file enables you to install an image on a partition that is smaller, equal to, or larger than the original partition that was captured, provided the target partition has sufficient space to store the image content. This is unlike sector-based image formats that require you to deploy a disk image to a partition that is the same size or larger than the source disk.
- Windows Server 2012 provides an application programming interface (API) for the .wim image format, named Windows Imaging API, which developers can use to work with .wim image files.
- The .wim file allows for nondestructive image deployment. This means that you can leave data on the volume where you apply the image, because when you apply the image, it does not delete the disk's existing contents.
- The .wim file format enables you to start the Windows Preinstallation Environment (Windows PE) from a .wim file. Windows 8 and Windows Server 2012 setup processes uses Windows PE. The .wim file is loaded into a random access memory (RAM) disk, and run directly from memory.
- The .wim file format supports media (disk) spanning, which provides you with the option to use large customizable images that contain operating systems and applications. You can span them across multiple pieces of media for their own installation and recovery solutions.

The .wim File Structure

A .wim file structure contains up to six types of resources:

- The .wim header. The .wim header defines the .wim file content, such as memory location of key resources (metadata resource, lookup table, and XML data), and .wim file attributes (version, size, and compression type).
- File resource. The file resource is a series of packages that contain captured data, such as source files.
- Metadata resource. The metadata resource stores information on how captured data is organized in the .wim file, including directory structure and file attributes. There is one metadata resource for each image in a .wim file.
- Lookup table resource. The lookup table resource contains the memory location of resource files in the .wim file.
- XML data resource. The XML data resource contains additional miscellaneous data about the .wim image, such as directory and file counts, total bytes, creation and modification times, and description information.
- Integrity table resource. The integrity table resource contains security hash information used to verify the integrity of the image during an apply operation.

Additional Reading: For more information on Windows Imaging File Format (WIM), visit the following link: http://go.microsoft.com/fwlink/?LinkID=391886

Performing High-Touch with Retail Media Deployments

The High Touch with Retail Media deployment method is the most common deployment strategy for small organizations that do not intend to deploy many servers. Those organizations typically do not implement a strategy for unmanaged networks and distributed locations. In most deployment scenarios, small organizations order new servers with the operating system installed already, or perform installation by using local media such as a Windows Server 2012 product DVD.

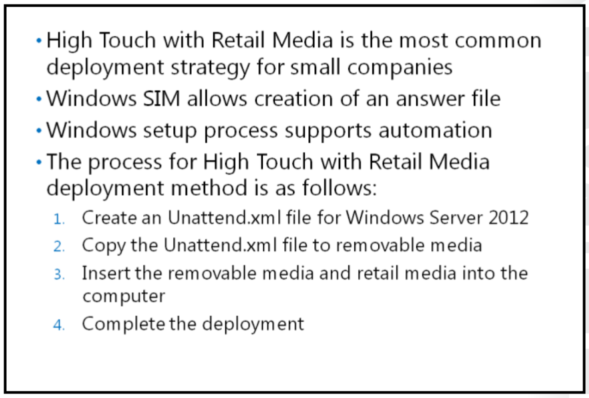

- High Touch with Retail Media is the most common deployment strategy for small companies
- Windows SIM allows creation of an answer file
- Windows setup process supports automation
- The process for High Touch with Retail Media deployment method is as follows:
  1. Create an Unattend.xml file for Windows Server 2012
  2. Copy the Unattend.xml file to removable media
  3. Insert the removable media and retail media into the computer
  4. Complete the deployment

Some organizations might consider using technologies for automation to speed up the deployment process. By doing this, they eliminate the need to enter the same information and perform the same tasks repeatedly for each new deployment. This strategy combines the functionalities of the Windows System Image Manager (Windows SIM) tool with Windows Server 2012 retail media to create an answer file (Unattend.xml) that is stored on a USB drive. After you create an answer file, you can enter all entries in the answer file during the Windows Server 2012 installation. The operating system installation process then uses the answer file as a basis for steps that require user input.

Note: Windows SIM is included in the Windows Assessment and Deployment Kit (Windows ADK).

The Windows Setup program supports automation in the following areas:

- Hard disk partitioning
- Device driver installation
- Application installation
- Applying updates
- Configuring settings
- Adding roles and features
- User interface setup suppression

The following steps describe the process involved in the High Touch with Retail Media deployment method:

1. Create an Unattend.xml file for Windows Server 2012 by using Windows SIM.
2. Copy the Unattend.xml file to your removable media, such as a USB flash drive.
3. Insert the removable media and the retail media into the computer.
4. Complete the deployment by installing applications and configuring the computer as required.

The High Touch with Retail Media deployment method does not scale well because it always produces the same configuration settings and features on newly-installed servers. Therefore, if organizations need to install several servers with different server roles, manual configuration tasks should be performed on each server.

Note: You need not perform this deployment method with local media. You can use any media, including media that is hosted on a network share for ease of distribution.

Question: What are the limitations of the High Touch with Retail Media deployment method?

Performing High-Touch with Standard Image Deployments

The High Touch with Standard Image deployment method is very similar to the High Touch with Retail Media deployment method. However, with this method, rather than using retail media, you use a standard server image captured from a reference computer. (In this scenario, a *reference computer* is the computer that contains the master image.) The High Touch with Standard Image deployment method is faster, because the standard server image includes settings and applications, and provides consistency across all servers. Hence, the validation and testing times are reduced. This method also decreases maintenance costs because administrators can perform many updates to the standard images offline.

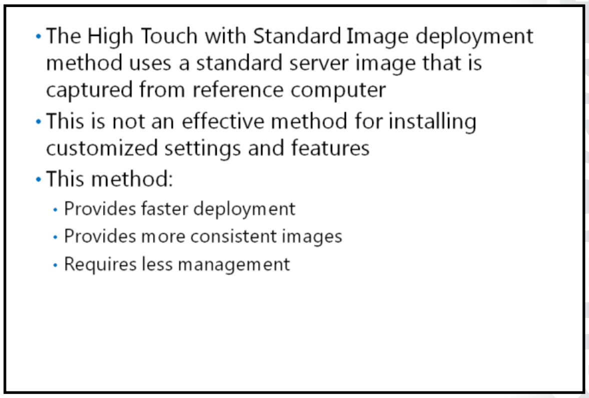

- The High Touch with Standard Image deployment method uses a standard server image that is captured from reference computer
- This is not an effective method for installing customized settings and features
- This method:
  - Provides faster deployment
  - Provides more consistent images
  - Requires less management

The deployment process for the High Touch with Standard Image deployment method involves the following steps:

1. Set up a server machine as the template for creating your first image.
2. Install the Windows Server 2012 operating system from the retail media or volume license media. We recommend that you use an answer file (Unattend.xml) to install Windows Server 2012 on the reference computer to make this process consistent and reproducible.
3. Install any applications, drivers, settings, and updates that you want to include in the final image.
4. Run the System Preparation Tool (Sysprep) to generalize the image for deployment. This removes the unique and identifying characteristics of the reference computer.
5. Boot into the computer by using Windows PE, and capture the image by using the Deployment Image Servicing and Management (DISM) command-line tool.
6. Copy the image to removable media, external hard drive, or network share.
7. Prepare the installation media. You can do this in one of two ways:
   - Create an answer file (Unattend.xml) and point it to the image that you copied to the drive or network share.
   - Create a new installation media by replacing the Install.wim file with the image file you captured previously.
8. Start deploying the image onto each server computer, either by using the answer file (Unattend.xml) or by using the setup media that you created.
9. Activate the machines online.

Requirements for High Touch with Standard Image Deployment

The High Touch with Standard Image deployment method has the following requirements:

- Windows Server 2012 retail or volume license media
- Windows ADK
- Removable storage device
- Reference computer on which to create and configure the source image

Performing Lite-Touch High-Volume Deployments

The lite-touch high-volume deployment method involves only limited manual intervention, primarily at the beginning of the deployment. The rest of the process is automated and well suited for organizations with a dedicated IT staff.

The lite-touch high-volume deployment method requires MDT 2013. For Window PE images used for Windows Server 2012 and Windows 8, MDT 2013 requires Windows ADK, and Windows Deployment Services (Windows DS) for Pre-Boot Execution Environment (PXE) boot. MDT 2013 also requires at least one volume license media provided by Microsoft, and a file server on which to store the distribution share that is accessed by computers during deployments.

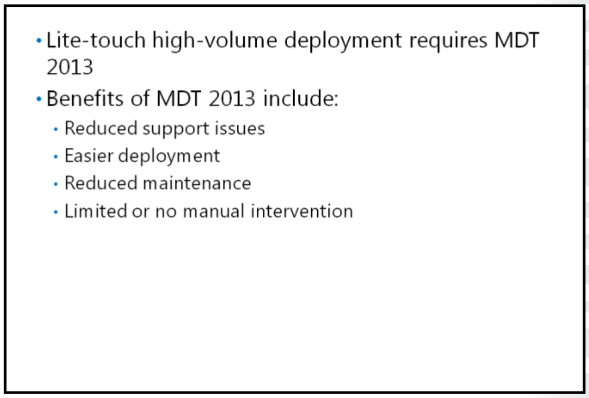

- Lite-touch high-volume deployment requires MDT 2013
- Benefits of MDT 2013 include:
  - Reduced support issues
  - Easier deployment
  - Reduced maintenance
  - Limited or no manual intervention

MDT 2013 provides the following benefits when supporting lite-touch high-volume deployments:

- Reduced support issues as a result of consistent images.
- Easier deployment. MDT 2013 supports device drivers, updates, and applications.
- Reduced maintenance effort. Tasks such as updating drivers, applications, and the core operating system are simplified.

The lite-touch high-volume deployment method may require initial effort from the administrators. This includes receiving training about the technology, and testing the deployment process. After the deployment is in place, administrators can use the lite-touch method to deliver complex migration scenarios. As organizations grow, administrators can provide a nearly zero-touch experience if they configure the MDT 2013 database and deploy the Windows DS role.

When planning MDT 2013 deployments, there are factors that you must consider. The deployment process can consume a significant amount of space on both networks and hard drives, because lite-touch installation (LTI) deployment files are stored in the MDT 2013 deployment shares. Most LTI files are deployed across the network, which requires sufficient bandwidth between deployment shares and computers. You can create .wim files containing a single image or multiple images. Additionally, you may decide to include applications in your images or install them after deploying the base image.

Additional Reading: For more information on advanced deployment usage scenario by using Microsoft Deployment Toolkit 2013, go to http://go.microsoft.com/fwlink/?LinkID=391887.

Performing Zero-Touch High-Volume Deployments

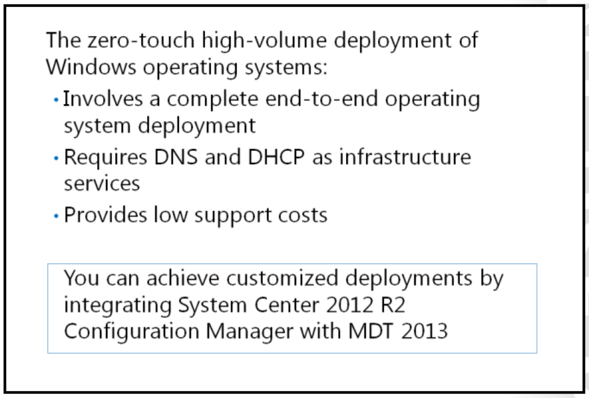

The zero-touch high-volume deployment of Windows operating systems:

- Involves a complete end-to-end operating system deployment
- Requires DNS and DHCP as infrastructure services
- Provides low support costs

You can achieve customized deployments by integrating System Center 2012 R2 Configuration Manager with MDT 2013

The zero-touch high-volume deployment method combines all the technologies that focus on complete end-to-end deployment of Windows operating systems. You can configure the operating system deployment to occur with no user interaction on computers that have no operating system installed, or on machines that have an installation of the Windows Server 2012 operating system.

In all zero-touch high-volume deployment methods, the primary requirement is Configuration Manager, which requires Active Directory® Domain Services (AD DS), Windows DS, and Windows ADK. However, the most effective and customized deployments are achieved when you integrate Configuration Manager with MDT 2013. This integration uses customizability through MDT task sequences, and Configuration Manager infrastructure services, such as a management point and a distribution point.

Note: Zero-touch high-volume deployments also require the Domain Name System (DNS) and Dynamic Host Configuration Protocol (DHCP) infrastructure services.

The benefits of the zero-touch high-volume deployment method include:

- Streamlined deployments with consistent configurations.
- Lower support costs for ongoing management of deployed servers, as a result of spending less time in post-deployment configurations.

However, administrators must meet the prerequisites of knowledge and skills to manage zero-touch installation (ZTI), because ZTI requires configuration of a relatively complex infrastructure service.

Compared to the LTI approach that uses MDT 2013 with Windows DS, the ZTI approach that uses Configuration Manager provides the following additional benefits:

- Bandwidth management of image transfers
- Reporting on driver availability for devices across your organization
- Tolerance of poor or intermittent network connectivity
- Fully unattended deployment
- Offline deployment from media and CD or DVD spanning
- Encryption and password protection

Discussion: What Is Your Current Deployment Strategy?

Review the discussion questions and participate in a discussion about your current deployment strategy.

Question: How do you currently deploy operating systems within your organization?

Question: Discuss the different scenarios and deployment strategies used by various organizations.

Question: Is your deployment strategy based on files or binary images?

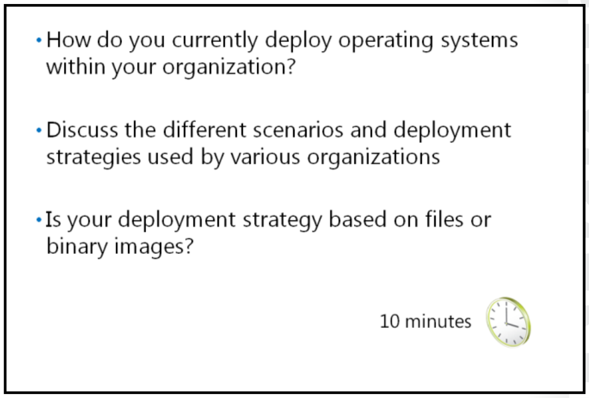

- How do you currently deploy operating systems within your organization?
- Discuss the different scenarios and deployment strategies used by various organizations
- Is your deployment strategy based on files or binary images?

10 minutes

Why Use Automated Deployment Methods?

Most companies are searching for a deployment process that can provide consistent configuration and faster deployment. An *automated deployment* is an installation in which user input during software installation is either limited or not required at all. The benefits of automating deployments are enhanced business efficiency, effectiveness, and capability. Using automated deployment also helps prevent user errors that can occur when manually deploying operating systems and additional applications.

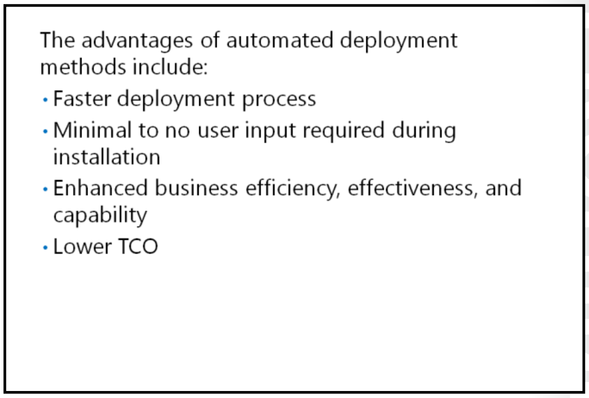

The advantages of automated deployment methods include:

- Faster deployment process
- Minimal to no user input required during installation
- Enhanced business efficiency, effectiveness, and capability
- Lower TCO

Initially, moving from manual to automated deployments can seem daunting. In the beginning, there is an increased burden on deployment teams to learn the system and create new processes to manage the automation. In some cases, the standard configuration created by an automated deployment process may not be suitable for all servers. However, once you establish the automation process, the benefits that you gain in future deployments outweigh the initial effort of repeated testing. The benefits include increased overall productivity and lowered total cost of ownership (TCO) of the system.

Selecting an Image-Based Deployment Strategy

In medium-to-large organizations, the common practice for administrators is to put initial effort into preparing the image strategy, and then add optional applications, including customizations and updates. By having one standard image for all deployments, you can reduce the cost of maintaining, updating, and modifying the image. However, in reality, most enterprises will have a number of different images. For example, organizations may have one image for laptops, one for desktops, and possibly one for each department's desktops. In a worst-case scenario, enterprises could have one image for each hardware configuration, which will increase the TCO in maintaining deployment strategies.

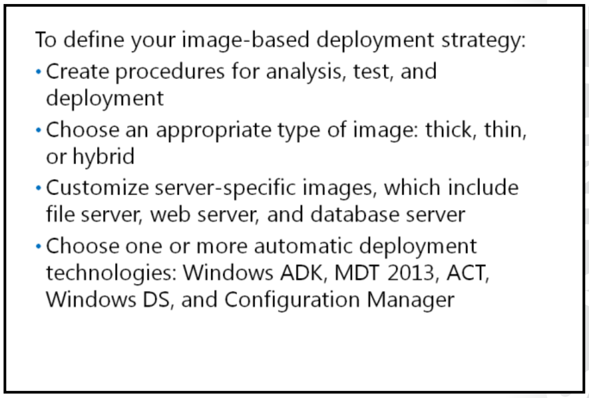

To define your image-based deployment strategy:

- Create procedures for analysis, test, and deployment
- Choose an appropriate type of image: thick, thin, or hybrid
- Customize server-specific images, which include file server, web server, and database server
- Choose one or more automatic deployment technologies: Windows ADK, MDT 2013, ACT, Windows DS, and Configuration Manager

In almost any deployment method, to reduce the cost of ownership, you must perform the following steps:

1. Planning. In this step, you select a strategy based on your organizations' business requirements.
2. Testing. Initial time and effort that you spend in testing the deployment will reduce the number of support and maintenance calls, and result in a proper deployment scenario.
3. Storing images. In this step, you select a file server on which to store the images. During the deployment process, these images will be applied to your servers.
4. Distributing. This is the last phase in the deployment strategy. In this step, you apply images across the network to the target servers, or you apply images from removable media such as a USB flash drive or a DVD.

These steps ensure that you spend less administrative time and effort on updating and modifying deployment images, and particularly in supporting failed or improperly deployed images. In addition, a unified task sequence provides a more customized and flexible deployment solution, thereby avoiding changes in several images. Furthermore, a proper distribution scenario is critical for avoiding network saturation.

As part of the planning process, you determine the types of images that you will create. The types of images from which you can choose to create are thick images, thin images, and hybrid images.

Thick Images

A *thick image* is an image that contains everything inside the image—that is, the operating system and all required applications and updates. Thick images are popular in standardized environments where most of the users require the same set of applications. These types of images decrease the time for image preparation and deployment. To prepare and deploy this type of image, you build a reference machine and install all possible applications to ensure that users have access to all applications that they could possibly require. Next, you apply software updates to the operating system and to all the applications. Finally, you use Sysprep to capture the image.

The disadvantage of thick images is the lack of customizability, because all computers receive the same set of applications. This results in your organization paying for several applications that may not be necessary for all users. In addition, images are larger, and as such, multiple applications can affect operating system performance. Finally, images are more difficult to maintain, and flexibility is greatly reduced.

Thin Images

A *thin image* is an image that installs only the operating system and a few core applications or language packs. This approach uses the opposite strategy of thick images by keeping the image relatively small, or *thin*. The advantages of thin images are that they provide for a more flexible deployment by combining different applications for different departments, thereby reducing the costs related to image development, testing, storage, and distribution. Thin images are also smaller than thick images, and thereby have less of an effect on operating system performance.

However, thin images introduce complexity because you must deploy additional applications in combination with thin images as needed for each user. Each application that you have to deploy requires additional customization. More importantly, these types of images can take considerably longer to deploy on target systems, because after the initial image deployment, the target computer has to be further customized with additional applications and settings.

Hybrid Images

Hybrid images combine the benefits of both thin and thick image strategies. As such, organizations most commonly use the hybrid approach. Because most organizations have a common application that they need to deploy on almost every system, the remaining applications are deployed after the initial image, and in custom order for different deployment scenarios.

Server-specific images, such as file server, web server, or database server images require additional preparation and testing. For example, most of the file server images will have common operating system partitions and installed roles, but they will need additional storage area network (SAN) connections to provide unique file system resources. Furthermore, web server images are commonly used for deploying a large number of similar web servers. However, after initially deploying the base image containing the Windows Server operating system and web application role, web server images require additional customization. This is because not all departments require all of the same common services. In addition, database server images will have common Windows Server operating system installations and Microsoft SQL Server® versions. After deploying the prepared image, you will create additional databases.

These types of workloads are good candidates for creating service templates by using System Center 2012 R2 Virtual Machine Manager (VMM). This is because service templates lower the time required to provision and set up a virtualized environment.

Technologies for Automated Deployment

Regardless of the deployment strategy that you use, you can combine some of the following technologies for automated deployment:

- Windows ADK
- MDT 2013
- Microsoft Application Compatibility Toolkit (ACT)
- Windows DS
- Configuration Manager

Note: For more information on the technologies for automated deployment, refer to Lesson 2 in this module: Implementing an Automated Deployment Strategy.

Lesson 2
Implementing an Automated Deployment Strategy

Organizations that choose to perform an automated server deployment should first decide on the deployment strategy. This strategy should include the methods and tools necessary to perform the deployment, based on their organizational requirements. Smaller organizations might use free, automatic deployment tools that are available for download from the Microsoft website. Medium and large enterprise organizations may choose to use Configuration Manager as a solution for automated deployment.

Regardless of the tool used, organizations should train their IT administrators on these technologies. Moreover, we recommended that the automated deployment strategy include thorough testing of the deployment process prior to implementing it in the production environment.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the tools that you can use for image-based installations of Windows operating systems and Windows Server operating systems.
- Describe the components included in the Windows ADK.
- Describe the Windows DS architecture, and its enhancements in Windows Server 2012.
- Plan for Windows DS.
- Explain how MDT 2013 is useful for deploying Windows 8 and Windows Server 2012 operating systems.
- Deploy Windows Server by using Configuration Manager.
- Select a deployment topology.
- Prepare the Windows Server 2012 image.

Overview of Image-Based Installation Tools

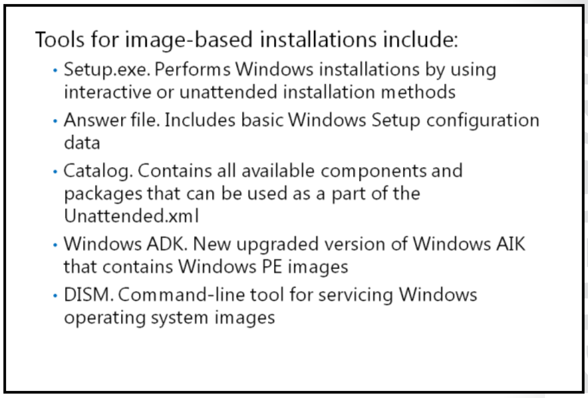

Tools for image-based installations include:

- Setup.exe. Performs Windows installations by using interactive or unattended installation methods
- Answer file. Includes basic Windows Setup configuration data
- Catalog. Contains all available components and packages that can be used as a part of the Unattended.xml
- Windows ADK. New upgraded version of Windows AIK that contains Windows PE images
- DISM. Command-line tool for servicing Windows operating system images

The .wim files are compressed packages that contain several related files. All Windows Server 2012 installations use the .wim format. When installing Windows Server 2012, you apply an image to a hard disk. This process occurs at a file-level instead of at a hard-disk sector level.

Tools for Performing Image-Based Windows Installation

You can choose from several tools and technologies to perform an image-based Windows operating system installation. You must be familiar with these tools and know where and when to use them in deployment situations:

- Windows setup command-line options (Setup.exe). This tool performs Windows installations by using interactive or unattended installation methods.

- Answer file (Unattend.xml). A simple answer file includes basic Windows Setup configuration data, and minimum Windows Welcome customizations, which starts after the Windows Setup program runs.
- Catalog. This tool contains all available components and packages that can be used as a part of the Unattend.xml answer file, and can be modified through Windows SIM.
- Windows ADK. This is a new upgraded version of Windows Automated Installation Kit (Windows AIK) that contains Windows PE images, which are necessary for customized deployment of Windows Server 2012 and Windows 8.

DISM

Deployment Image Servicing and Management (DISM) is a command-line tool that combines separate Windows platform technologies into a single, cohesive tool for servicing Windows operating system images. You can use DISM to perform the following tasks:

- View the contents of a .wim file. DISM provides you with the ability to view the contents of a .wim file. This is useful to see which images are available to deploy from within the .wim file.
- Capture and apply images. You can capture an image of a source computer and save it as a .wim file format. You can save the image to a distribution share, from which users can use Windows 8 Setup to install the image, or you can push the image out to the desktop by using various deployment techniques. You also can use DISM to apply the image to the destination computer.
- Store multiple images in a single file. You can use DISM to store multiple images in a single .wim file to utilize single instancing, which minimizes the size of the image file. This also simplifies the process for an administrator deploying multiple images either by using removable media, or by deploying across a slower network connection. When you install Windows 8 by using a file with multiple images, end users can select which image to apply. For example, you can have a .wim file that contains several role-based configurations or images before and after certain updates.
- Compress the image files. DISM supports two different compression algorithms—fast, and maximum—to further reduce the image size.
- Implement scripts for image creation. You can use scripting tools to create and edit images.

DISM uses the following technologies:

- Unattended installation answer file. When you use DISM to apply an Unattend.xml answer file, the updates that the answer file specifies are implemented either on the Windows operating system image, or on the running operating system. You can configure default Windows settings and add drivers, packages, software updates, and other applications by using the settings in an answer file.
- Windows SIM. You can use Windows SIM to create unattended installation answer files, create distribution shares, and modify the files that are in a configuration set.
- Mount images for offline image editing. A common scenario for DISM is customizing an existing image, which includes updating files and folders, and adding drivers and additional server roles or features.
- OCSetup. OCSetup is a command-line tool that you can use when you are applying updates to an online Windows operating system image. OCSetup installs or removes Component-Based Servicing (CBS) packages online by passing packages to DISM for installation or removal. You can also use OCSetup to install Windows Installer package (.msi) files by calling the Windows Installer service (MSIExec.exe) and passing Windows Installer components to it for installation or removal. Additionally, you can use OCSetup to install packages that have custom installers, such as .exe files.

Note: In Windows AIK, the ImageX tool is deprecated, and all imaging functionality has been added to the DISM tool.

Windows ADK for Windows 8.1

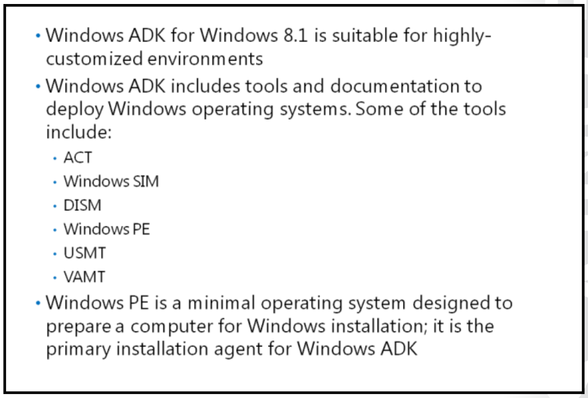

- Windows ADK for Windows 8.1 is suitable for highly-customized environments
- Windows ADK includes tools and documentation to deploy Windows operating systems. Some of the tools include:
  - ACT
  - Windows SIM
  - DISM
  - Windows PE
  - USMT
  - VAMT
- Windows PE is a minimal operating system designed to prepare a computer for Windows installation; it is the primary installation agent for Windows ADK

Windows ADK for Windows 8.1 is a collection of tools and documentation designed to help IT professionals deploy the Windows 8.1 and Windows Server 2012 R2 operating systems. Windows ADK is ideal for using with highly customized environments because you can use the Windows ADK tools to configure many deployment options. Depending on your business needs, you can choose which tools to use from Windows ADK for Windows 8.1.

By default, Windows ADK is installed to the C:\Program Files (x86)\Windows Kits directory. This directory contains all the tools and documentation included in the Windows ADK.

Documentation in Windows ADK

Windows ADK consists of the documentation components listed in the following table.

Documentation	Description
Getting Started with ADK (ADK_GetStarted.chm)	Getting Started with ADK provides the conceptual and procedural information required for unattended installation of Windows operating systems. This user's guide includes the following information: • Planning • Preparing the deployment environment • Creating and customizing an image • Capturing, modifying, and testing the image • Deploying, maintaining, and servicing the image
Component Platform Interface (CPI) Reference (Cpiapi.chm)	The CPI Reference documents the APIs that are used in Windows SIM.
Application Compatibility Toolkit User’s Guide (ACT.chm)	The ACT helps you determine whether the applications, devices, and computers in your organization are compatible with versions of the Windows operating system.
Windows Assessment Services User’s Guide (ASMT.chm)	Windows Assessment Services is a framework that you can use to automate quality measurements such as performance, reliability, and functionality, on multiple computers in a lab environment.
Unattended Windows Setup Reference (Unattend.chm)	The Unattended Windows Setup Reference provides a complete listing of all the settings that you can use to automate the configuration and the deployment of Windows 8, Windows 7, Windows Vista, Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008.

Documentation	Description
Volume Activation Management Tool (Vamt.chm)	The Volume Activation Management Tool (VAMT) is designed to manage volume activation for Windows Vista, Windows 7, Windows 8, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, and Microsoft Office 2010.
Windows Performance Toolkit Technical Reference (WPT.chm)	The Windows Performance Toolkit consists of performance monitoring tools that produce in-depth performance profiles of Windows operating systems and applications.

Tools in Windows ADK

The Windows ADK tools that are used in most Windows deployment scenarios are as follows:

- ACT
- Windows SIM
- Windows Recovery Environment (Windows RE)
- DISM
- Windows PE
- User State Migration Tool (USMT)
- VAMT
- Command-line tools such as BCDBoot and Oscdimg
- Windows Imaging API
- DISM API

ACT

ACT is a management tool that you can use in managing your overall application portfolio. This tool helps reduces the cost and time involved in resolving application compatibility issues, and helps deploy and update Windows operating systems. You can perform the following key functions by using ACT:

- Analyze overall portfolio of applications, websites, and computers.
- Evaluate operating system deployments and the impact of operating system updates.
- Centrally manage compatibility evaluators and configuration settings.
- Deploy automated mitigations to known compatibility issues.
- Participate in exchange of compatibility information within the Microsoft Compatibility Exchange web service.

Windows SIM

Windows SIM is a tool that you can use to create an unattended Windows Setup answer file. *Answer files* are XML-based files that you create by using information from a .wim file and a catalog (.clg) file, which can contain configuration settings used during Windows operating system deployments. Some common usage scenarios for the Windows SIM tool include adding additional device drivers during Windows setup, adding applications that you can install during setup, or adding offline updates to Windows operating system image. After you create the answer file, you can use this file either by specifying explicitly the location of the answer file by using the **Setup.exe /unnatend:*filename.xml*** command, or by implicitly searching in several different locations (such as the root of the drive or registry defined locations).

Additional Reading: For a complete list of valid search paths, look for the Implicit Answer File Search Order section on the Methods for Running Windows Setup webpage at http://go.microsoft.com/fwlink/?LinkID=277144.

Windows PE

Windows PE is a compact, special-purpose Windows operating system that prepares and initiates a computer for Windows operating system setup, maintenance, or imaging tasks. Windows PE also recovers Windows operating systems such as Windows 8 and Windows Server 2012, and is the core deployment foundation for Windows 8.

With Windows PE, you can start a subset of the Windows 8 operating system from a network location or from removable media, which provides network and other resources necessary to install and troubleshoot Windows 8. Although Windows PE is not a general-purpose operating system, you can use it to start a computer that has no functioning operating system installed, and it can act as a replacement for MS-DOS–based boot disks that were utilized in previous Windows operating system versions.

USMT

USMT is a tool that you can use to migrate user data from a previous Windows operating system to Windows 8. USMT scans the computer for existing user accounts, user files, operating system settings, and application settings. USMT then migrates them to a new Windows Server installation. You can customize the migration settings by using the migration rule (.xml) files, and you can control which files and settings will be migrated.

VAMT

VAMT is a tool that enables network administrators and other IT professionals to automate and centrally manage the Windows volume activation process for computers in their organization. VAMT can manage volume activation by using multiple activation keys (MAKs) or the Windows Key Management Service (KMS).

Windows Deployment Command-Line Tools

Windows ADK includes a large set of command-line tools for use in different deployment scenarios. Some of the most commonly used command-line tools are:

- BCDBoot. This tool initializes the Boot Configuration Data (BCD) store, and copies boot environment files to the system partition.
- Oscdimg. This tool creates .iso images.
- Powercfg. This tool controls power settings.
- Tzutil. This tool manages available time zones.

DISM API

Developers use DISM API to perform different management tasks for Windows operating system images. These tasks can include installing, uninstalling, configuring, and updating Windows operating system features, packages, and drivers in a Windows operating system image.

The DISM API is designed for use by C or C++ programmers.

Additional Reading: For more information about Windows deployment command-line tools, see http://go.microsoft.com/fwlink/?LinkID=391888.

Windows DS

Windows DS is an update to Remote Installation Services (RIS). You can use Windows DS to provide a rapid deployment of Windows operating systems by using PXE over a network. Alternatively, you can use Windows DS to start a remote computer by using Windows PE. You can also use Windows DS to help support both lite-touch and zero-touch high-volume deployments.

Windows DS in Windows Server 2012 enables support for deploying Windows client and Windows Server operating systems, including Windows Server 2012, Windows Server 2008, Windows Server 2008 R2, Windows 8, and Windows 7.

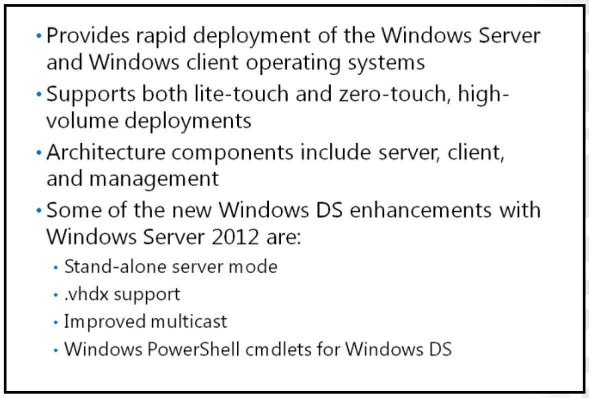

- Provides rapid deployment of the Windows Server and Windows client operating systems
- Supports both lite-touch and zero-touch, high-volume deployments
- Architecture components include server, client, and management
- Some of the new Windows DS enhancements with Windows Server 2012 are:
  - Stand-alone server mode
  - .vhdx support
  - Improved multicast
  - Windows PowerShell cmdlets for Windows DS

Note: If you combine Windows DS with MDT 2013, you can deploy highly-customized installation images by using network-based installation.

Windows DS architecture contains the components listed in the following table.

Component	Description
Server	The server component contains: • A PXE server. • Trivial File Transfer Protocol (TFTP) for starting a client from the network to install a Windows operating system. • A shared folder and image repository that contains boot images, installation images, and files that are required specifically for the network boot functionality.
Client	A GUI that: • Runs within Windows PE. • Communicates with server components to select and install a Windows operating system image.
Management	A set of tools to manage: • The server. • Windows operating system images and bootable images. • Client computer accounts. • Driver groups.

Benefits of Windows DS

Windows DS provides the following installation and deployment benefits:

- Reduces the complexity of deployments and the costs associated with inefficient manual installation processes.
- Provides the ability for a user to perform a network-based installation of Windows operating systems.
- Deploys Windows operating system images to computers that do not have operating systems installed.
- Provides an end-to-end solution for deploying Windows operating systems to client computers and servers.
- Uses standard setup technologies, including Windows PE, .wim files, and image-based setup.

New Windows DS Features and Improvements in Windows Server 2012

Windows DS in Windows Server 2012 contains the following new features and improvements:

- Stand-alone server mode. This mode of installation removes AD DS dependency by using a local store for pre-staged devices, but still requires DHCP and DNS services.
- Virtual hard disk file support. Windows Server 2012 has a new virtual hard disk file format (.vhdx) that supports larger storage capacity than the older .vhd format in Windows Server 2008. The .vhdx format also provides data corruption protection.
- Improved multicast. Overall performance of multicast is improved compared to previous multicast transmission rates. This is because Windows DS reduces the default block size within an Ethernet Maximum Transmission Unit (MTU). In addition, Windows DS is compatible with hardware that does not support IP fragmentation.
- Expected Deployment Result Wizard. This wizard provides you with an option to view what boot, install image, and driver groups are offered to specific (pre-staged) computers.
- TFTP enhancements. TFTP enhancements include:
  - Performance improvements regarding scalable buffer and port management.
  - Usage of variable-size transmission windows.
  - Option to specify Maximum TFTP block size directly from the Microsoft Management Console (MMC) in Windows DS, or through the command-line tool for Windows DS, Wdsutil.exe.
- Boot image and install image priorities. This feature provides control over which images will be provided to computers.
- Windows PowerShell® cmdlets support for Windows DS. Windows Server 2012 R2 adds support for Windows PowerShell cmdlets for Windows DS. This means that you can use Windows PowerShell cmdlets to add driver packaging, add client and server images, enable and disable bootable and install images, and perform additional common Windows DS tasks.

Additional Reading: For more information about Windows Deployment Services Cmdlets in Windows PowerShell, go to http://go.microsoft.com/fwlink/?LinkID=391889.

Planning for Windows DS

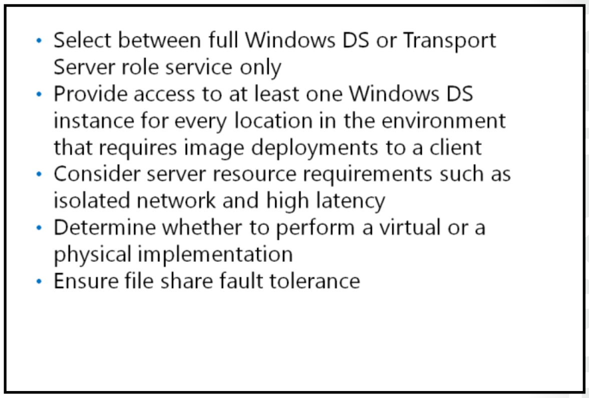

- Select between full Windows DS or Transport Server role service only
- Provide access to at least one Windows DS instance for every location in the environment that requires image deployments to a client
- Consider server resource requirements such as isolated network and high latency
- Determine whether to perform a virtual or a physical implementation
- Ensure file share fault tolerance

While installing the Windows DS server, you can choose one of the following role services:

- Transport Server role service. This role provides only the core networking components required for creating and managing a multicast stream. A multicast stream allows multiple clients to tune into a stream of data without requiring that the data be sent individually to each client on a separate unicast stream.
- Deployment Server role service. This role provides full functionality of Windows DS, including the following:
  - PXE boot services
  - MMC tools
  - The ability for the client to select which image to install from a presented list
  - Both unicast and multicast deployments

For every location in the environment that requires image deployments to a client, you should deploy access to at least one Windows DS instance. If the clients are separated by a wide area network (WAN) from the planned Windows DS instance, ensure that the WAN provides low latency and enough available bandwidth for Windows DS to function properly.

Although a single Windows DS instance may be sufficient to meet the image deployment requirements of a location, additional requirements may require you to plan for multiple Windows DS instances within a single physical location.

You may require additional Windows DS instances for the following reasons:

- Isolated network. You may have isolated networks that require image deployments, such as training labs that you must keep separate from the organization's network.
- Low bandwidth availability or high latency. If clients are separated from the Windows DS servers by segments that have low available bandwidth or high latency, another Windows DS deployment may be necessary to manage those clients.

The factors that determine the number of Windows DS servers to be implemented and the appropriate hardware configuration include:

- Total number of computers. Identify the total number of computers in the organization in which the operating system will be deployed by using Windows DS. This information sets the peak number of imaging requests that the infrastructure may need to manage simultaneously. The worst-case scenario is that all clients are imaged simultaneously.
- Image deployment speed. Identify the targeted amount of time that an image deployment should take.
- Size and number of images. Identify the total number of images that are available in the location, and the size of each image. This will help to identify disk capacity requirements for the server.

To increase the availability of the infrastructure, you can make the share through which the .wim-based images are accessed, fault tolerant by using the following methods:

- Distributed File System (DFS). You can use DFS to provide a fault tolerant method for accessing file shares. You can also use DFS to define a file namespace and provide multiple targets for folders contained within the namespace.

  When a client attempts to access a DFS-enabled share, the request is handled by the nearest DFS server that is hosting that particular share. By placing .wim-based images on DFS shares, administrators can control which server will provide the client with the install image. Furthermore, administrators can control bandwidth usage on WANs between clients in branch office sites and Windows DS servers in the organization headquarters.
  One important consideration that you need to plan for is IP addressing in the remote location. This is because locating the closest DFS depends on a proper Active Directory site design, which means correctly configured Active Directory sites.

- Server clustering. Server clustering can increase the fault tolerance of a single content storage system file share. The file share becomes a clustered resource that is running on a cluster with two or more computers. If the computer that hosts the file share fails, the file share moves to a remaining active node.

MDT 2013

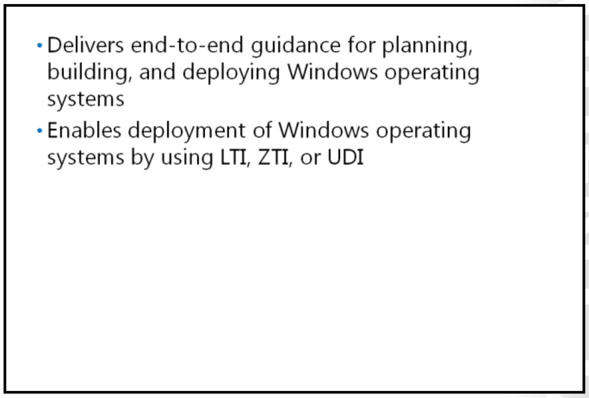

- Delivers end-to-end guidance for planning, building, and deploying Windows operating systems
- Enables deployment of Windows operating systems by using LTI, ZTI, or UDI

Microsoft Deployment Toolkit MDT 2013 provides end-to-end guidance for planning, building, and deploying the Windows 8.1 and Windows Server 2012 R2 operating systems. MDT 2013, together with several related technologies, allows you to deploy Windows Server 2012 by using a LTI or ZTI methodology, or user-driven installation (UDI).

All three types of deployment require that you install MDT 2013. MDT includes deployment benchmarks as a starting tool, and requires that you create a deployment share containing additional scripts and task sequences to customize and deploy the installation process. MDT has dependencies on Windows ADK, but also provides links for other useful tools that can produce customized Windows operating system installations.

MDT 2013 supports:

- Windows ADK for Windows 8.1.
- Deployment of Windows 8.1 and Windows Server 2012 R2, in addition to the Windows 7 and Windows 8 operating systems.
- ZTI with Configuration Manager.

Note: You can use the earlier MDT version, MDT 2012, to deploy Windows Server 2012. However, to deploy Windows Server 2012 R2, you must implement MDT 2013.

Lite-Touch Deployments

LTI commonly requires that an administrator or user start the installation on a client computer. Based on configuration, the user might also be required to interact with the installation process by providing some input. In a typical LTI deployment, you need to complete the following high-level steps:

1. Create a deployment share.
2. Import either a default install.wim image or a captured image from the reference computer.
3. Import drivers for specific hardware configuration.
4. Create an LTI task sequence to control and customize deployment.
5. Update the deployment share.

This final step will create a bootable image that you can use either to boot the computer from media, or to import in Windows DS for PXE deployment.

Zero-Touch Deployments

ZTI deployments can provide complete control over the installation process without the need for users to start the installation process locally. The ZTI process starts first with integrating Configuration Manager and MDT 2013. This extends the default Configuration Manager console with task sequences from MDT which provides:

- Input about which custom image file will be used for installation.
- An option to choose a custom boot.wim image.
- An option to select a driver package.
- An option to install updates and additional applications after the Windows operating system deployment.

MDT 2013 utilizes the Configuration Manager infrastructure components. Boot and install images are imported as packages into the Configuration Manager console, and computers can access these files through distribution points.

User-Driven Deployments

One of the main improvements in MDT 2013 is the UDI Wizard Designer, which is the primary tool for customizing wizard pages for the different deployment scenarios. Changes made in the UDI Wizard Designer are saved in the UDI Wizard configuration file, and are ultimately reflected in the user experience in the UDI Wizard. The user performing the deployment will see only the wizard pages in the UDI Wizard that you have selected and configured by using the UDI Wizard Designer. The remaining infrastructure requirements are the same based on whether you are using LTI or ZTI deployments.

Deploying Windows Server by Using Configuration Manager

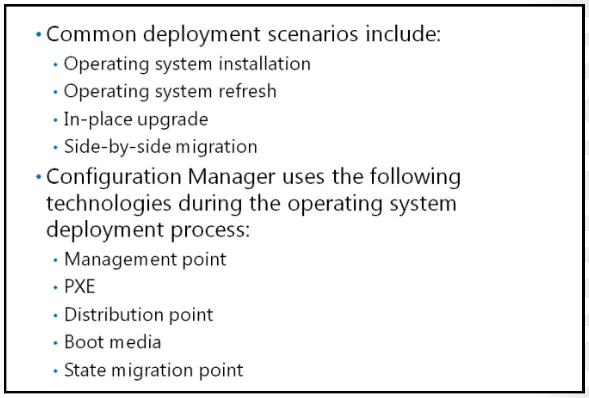

- Common deployment scenarios include:
  - Operating system installation
  - Operating system refresh
  - In-place upgrade
  - Side-by-side migration
- Configuration Manager uses the following technologies during the operating system deployment process:
  - Management point
  - PXE
  - Distribution point
  - Boot media
  - State migration point

Enterprise organizations that have a complex IT infrastructure can benefit from the automated capabilities of ZTI. A zero-touch solution is targeted primarily toward enterprise-class organizations that have large numbers of computers deployed in the network infrastructure. These organizations can utilize deployment automation capabilities, and can choose whether any end-user involvement is required.

Configuration Manager provides change and configuration management by providing functionalities such as the ability to deploy operating systems, software applications, and software updates. You can also use Configuration Manager for monitoring hardware and software inventory and remote computer administration.

Common scenarios in which you use the operating system deployment feature in Configuration Manager include:

- Operating system installation. You can use Configuration Manager to install a supported operating system on computer hardware that does not currently have an operating system.
- Operating system refresh. You can use Configuration Manager to install a supported operating system on a computer system with an existing operating system. In an operating system refresh scenario, you are not saving any data on the client system. You are only installing a new operating system.
- In-place upgrade. Sometimes, when performing an operating system refresh, you need to save user data on the system that you are refreshing. An in-place upgrade provides you with the tools necessary to automate saving data from the client system prior to the operating system refresh, and then restoring it after the operating system refresh completes.
- Side-by-side migration. When you replace a user's computer with a new computer, you can use side-by-side migration to save the data off the old system, install an operating system on the new system, and then restore the data to the new system. This method requires that the old computer is a Configuration Manager client, and that the new computer is linked to the old computer via computer association in Configuration Manager.

Operating system deployment uses many Configuration Manager components to deploy operating systems. Depending on the decisions you make, some of the components become optional. The Configuration Manager components that the operating system deployment uses are:

- Management point. When deploying to existing clients, the instructions that the administrator creates are copied to the management point.
- PXE. The PXE functionality that is enabled on distribution points will install the Windows DS role on a server. When deploying to computers that do not have an operating system installed, you can use PXE to boot the computer.
- Distribution point. The packages that the administrator creates are copied to the distribution point.
- Boot media. If deploying to computers that are offline, you can create a USB or optical boot media, and then use the boot media for the complete installation.
- State migration point. When deploying to existing clients, you can use a state migration point to store user state information.

Depending upon your deployment scenario, a Configuration Manager task sequence may reference one or more packages during installation. You can preconfigure these packages before you create a task sequence, or you can use the Import Microsoft Deployment Task Sequence command to create the packages that you need automatically. The following table lists the packages within a task sequence.

Package or image	Contains
Boot image package	Boot image that is used to initiate the ZTI deployment process.
Microsoft Deployment Files package	Contents of the Microsoft Deployment distribution share directory. (The files used from the distribution share directory are the scripts and control files.)
Windows Server operating system image	Windows Server operating system image to be deployed to the target computer.
Client package	Configuration Manager client installation files.

Package or image	Contains
USMT package	USMT files that are used to capture and restore the user state.
Custom settings package	Unattended files and customsettings.ini.
Sysprep files package (optional)	Specific Sysprep files that are defined for a package. (This package is required only for legacy operating systems such as the Windows XP operating system.)
Application packages (optional)	An option that you can use to customize deployment with different application installations.
Driver packages (optional)	Driver packages contain drivers that you import into the Configuration Manager distribution point, and allow the task sequence to choose Plug and Play drivers from the catalog

Designing the ZTI Environment

Designing the ZTI environment is a planning process. The design process consists of the following high-level steps:

1. Select the appropriate deployment scenarios.
2. Select the deployment methods.
3. Ensure that the required infrastructure exists.
4. Determine the appropriate processing rules. (These are required only if you are using MDT 2013 or Configuration Manager.)
5. Determine a monitoring plan.
6. Train team members.

Configuring and Deploying a Task Sequence

You can create task sequences that allow you to install an existing image package, or build and capture a reference operating system image. Alternatively, you can create a custom task sequence to perform a customized task using variables. You create a task sequence to deploy an existing Windows operating system image to a target computer by using the New Task Sequence Wizard in the Configuration Manager console.

You deploy task sequences to collections by using the Deployment Wizard in the Configuration Manager console. Before you run the Deployment Wizard, you need to know what target collections and desired run-time behavior you want for the deployment. Read access to the task sequence is required to deploy the task sequence, and the task sequence must exist prior to creating the deployment.

Choosing a Deployment Scenario

Operating system deployments can help you by deploying or upgrading systems on your network. Deployments can occur as a result of many different scenarios. The following table summarizes these scenarios.

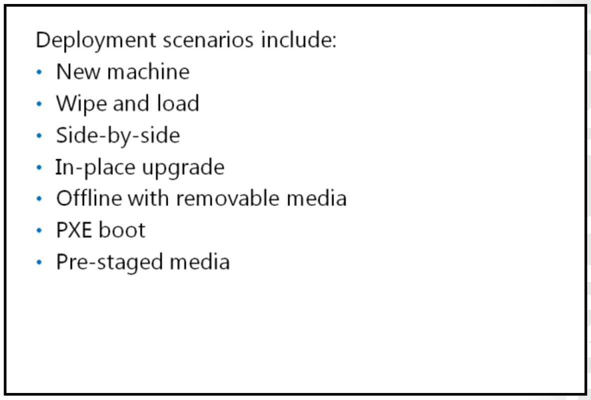

Deployment scenarios include:

- New machine
- Wipe and load
- Side-by-side
- In-place upgrade
- Offline with removable media
- PXE boot
- Pre-staged media

Scenario	Description
New machine	In this scenario, you can install a new Windows operating system on a client or server computer. The systems can be new, or have repurposed hardware. Use this scenario when the system has no operating system currently installed, or when the existing operating system is not the desired type for the system.
Wipe and load	Although similar to the new machine scenario, you use this scenario when there is an existing operating system already installed. Using this scenario, you utilize USMT to save the user's profile to a secure network share, install a new operating system on existing client or server hardware, and then restore the user state information.
Side-by-side	This is similar to the new machine scenario, except that you use this scenario to install a new operating system on new client hardware for an existing user. You then use Configuration Manager to reinstall applications on the new computer. Finally, use the USMT to move the user state from the old computer to the new computer through a secure network share.
In-place upgrade	Unlike the previous scenarios, do not use this scenario to install an operating system on a machine that has no operating system installed. Instead, use this scenario to upgrade one supported operating system to another supported operating system on an existing client or server. Any applications that are installed are migrated in place.
Offline with removable media	When there is little or no network bandwidth available, use this scenario. You use a Configuration Manager deployment package with removable media such as a CD, DVD, or USB flash drive. In a no-connectivity scenario, there is no status reporting.
PXE boot	This is similar to the new machine scenario, except that you use this scenario when you want to deploy with no or limited end-user interaction. This scenario uses Windows DS PXE technology. You can control the deployment actions with Configuration Manager advertisements.
Pre-staged media	Pre-staged media generally is copied to the hard disk drive of a new computer as a part of the computer manufacturing process, or at an enterprise-staging center before the computer is sent to the end user. When the computer starts for the first time after you load the pre-staged media, the computer will boot to Windows PE and connect to the site management point to check for available task sequences.

A clean installation is the most straightforward and simple installation method for Windows Server 2012. A clean installation involves the least number of variables in the installation process, and it results in a new, default installation of Windows Server 2012.

The Windows Server 2012 installation can be robust and usually trouble-free if your hardware meets the minimum requirements. However, a variety of problems can occur during an installation, and a methodical approach helps solve them.

Troubleshooting Approach

You can use the following four-step approach in any troubleshooting environment:

1. Determine what has changed.
2. Eliminate the possible causes to determine the probable cause.
3. Identify a solution.
4. Test the solution.

If the problem persists, go back to step 3, and repeat the process.

The following table describes several installation problems, and the solutions that you can use to identify and solve specific problems.

Problem	Solution
Installation media is damaged	Test the installation media on another system.
BIOS upgrade is required	Check your computer supplier's Internet site to determine whether a BIOS upgrade is available for Windows Server 2012.
Hardware is installed improperly	Review any messages that appear during the boot phase. Install add-on hardware (such as video cards and memory modules) properly. Choose pre-created drivers packages instead of allowing setup to choose best-matched plug and play drivers.
Hardware fails to meet minimum requirements	Use Windows Catalog to locate products designed for Windows operating systems, and ensure that your hardware meets the minimum requirements for the Windows Server 2012 edition that you want to install.
Error messages display during setup	Carefully note any messages, and search the Microsoft Knowledge Base for an explanation.

Question: When would you typically perform a clean installation of Windows Server 2012?

Question: What potential issues might you encounter when installing Windows Server 2012?

Demonstration: Preparing the Windows Server 2012 Image

In this demonstration, you will see how to:

- Create an image store, and map a network drive to the image store.
- Use the DISM tool to mount the relevant image.
- Use the DISM tool to modify the default installation image to include the Web Server Internet Information Services (IIS) role.

Demonstration Steps

1. Switch to LON-SVR1.
2. Create a new folder named **Images** on the **Allfiles (E:)** drive.
3. Create a new folder named **Custom Images** in the **E:\Images** folder.
4. On your host, in the 20413C-LON-SVR1 window, on the toolbar, click **Media**, point to **DVD Drive**, and then click **Insert Disk**.
5. In the **Open** dialog box, in the **File name** text box, type the following address, and then click **Open**:

   D:\Program Files\Microsoft Learning\20413\Drives\Windows2012R2.iso

6. Copy **D:\Sources\install.wim** (which is the default installation .wim image) from the installation media to the newly created folder **E:\Images\Custom Images**.
7. Share the **E:\Images** folder. Grant the **Administrator** user **Full Control** on the shared folder.
8. Map the network drive **Z:** to **\\lon-svr1\Images** so that drive Z is the default share for all images.
9. Open an elevated command prompt.
10. At the command prompt, type the following command, and then press Enter:

```
Mkdir c:\mounted
```

11. At the command prompt, type the following command, and then press Enter:

```
Dism /get-imageinfo /imagefile:”z:\Custom Images\install.wim”
```

Note: This command lists all images that are contained in install.wim. Notice the index for the Windows Server 2012 Datacenter edition.

12. At the command prompt, type the following command, and then press **Enter**:

```
Dism /mount-wim /wimfile:”z:\Custom Images\install.wim” /index:4 /mountdir:c:\mounted
```

Note: This command mounts the install.wim image for offline servicing. After you mount the image, you can add drivers, add packages, or enable features. This step will take approximately five minutes for the mounting of the image finish.

13. At the command prompt, type the following command, and then press Enter:

```
Dism /image:c:\mounted /get-features
```

Note: This command lists all available features and their state in the image file. If you want to view more details, redirect the output in the text file by using the following command:

```
Dism /image:c:\mounted /get-features > c:\All Features.txt
```

14. At the command prompt, type the following command, and then press Enter:

```
Dism /image:c:\mounted /get-featureinfo /featurename:IIS-WebServerRole
```

Note: The state of Web Server (IIS) role is disabled.

15. At the command prompt, type the following command, and then press Enter:

```
Dism /image:c:\mounted /enable-feature /featurename:IIS-WebServerRole –all
```

Note: This command installs the Web Server (IIS) role with all the depended features. Note that the name of the role is case sensitive.

16. At the command prompt, type the following command, and then press Enter:

```
Dism /unmount-wim /mountdir:c:\mounted /commit
```

Note: This command commits the changes in the install.wim image, which you will use later for deploying to a machine that has no operating system installed. This step to commit the changes in install.wim can take approximately five minutes.

Lab: Planning and Implementing a Server Deployment Infrastructure

Scenario

A. Datum Corporation has opted to implement Windows Server 2012 R2 throughout their organization. The head office, based in London, England, has been using a mix of Windows Server 2008 and Windows Server 2003 servers. The AD DS environment is based on Windows Server 2008 domain controllers.

The network design plan dictates that you must deploy Windows Server 2012 R2 in the head office to replace the existing Windows Server 2008 servers and domain controllers. In addition, the servers in the regional hub offices must be replaced with servers running Windows Server 2012 R2.

The smaller branch offices and regional distribution centers have a line-of-business (LOB) application that has been ported across from UNIX to Windows Server 2012 R2. Each location is to be equipped with its own server to support this application. In addition, network infrastructure services must be moved to Windows Server 2012 R2.

The network design team has stipulated that a number of server roles, including those of domain controllers, must be virtualized where possible. Your deployment design should encompass the likelihood of A. Datum acquiring two new companies in the near future, and the possibility that additional servers running Windows Server 2012 R2 will be deployed. In addition, you must create a standard server image for deployment, and prepare deployment tools to implement your deployment plan.

Objectives

After completing this lab, you will be able to:

- Plan an automated server installation and deployment strategy.
- Prepare the Windows Server 2012 R2 image.
- Deploy Windows Server 2012 R2.

Lab Setup

Estimated Time: 75 minutes

Virtual machines	20413C-LON-DC1 20413C-LON-SVR1 20413C-LON-SVR3
User name	Adatum\Administrator
Password	Pa$$w0rd

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

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Microsoft Hyper-V® Manager, click **20413C-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.

4. Sign in using the following credentials:
   - User name: **Administrator**
   - Password: **Pa$$w0rd**
   - Domain: **Adatum**
5. Repeat steps 2 through 4 for **20413C-LON-SVR1**.
6. Important: Do not start 20413C-LON-SVR3 until instructed to do so in the lab steps.

Exercise 1: Planning an Automated Server Installation and Deployment Strategy

Scenario

You have been assigned the task of planning a server installation strategy, selecting suitable installation and deployment tools, and planning an appropriate server deployment method.

A. Datum Automated Server Installation and Deployment Strategy Document

Document Reference Number: BS00929/1

Document Author Date	Brad Sutton 5th July

Requirements Overview:

To design an automated server installation and deployment strategy document, consider the following design factors:

- Familiarity with operating system image management by the IT staff
- Number of servers that must be deployed
- Variations in server configurations
- Use of retail or volume license media
- Network configuration, both in terms of the distribution of the servers for deployment, and in terms of the services installed currently to support the deployment process

The IT staff is planning to deploy Windows Server 2012 R2 to various offices, and has some experience with imaging and deployment. The configuration of the various servers is expected to be fairly consistent.

There is no requirement to upgrade settings from existing servers because there are plans to introduce new server hardware and to virtualize workloads where possible.

Proposals

1. What kind of image will you use: thin or thick?
2. Would lite-touch or zero-touch deployment be applicable for this scenario?
3. Which deployment technologies would you consider to implement the server upgrade plan?
4. What are the requirements for implementing this deployment technology?
5. Create a list of the deployment components necessary to support your server deployment plan.

The main tasks for this exercise are as follows:

1. Read the supporting documentation.

2. Update the proposal document with your planned course of action.

3. Examine the suggested proposals in the Lab Answer Key.

4. Discuss your proposed solution with the class, as guided by your instructor.

▶ Task 1: Read the supporting documentation

- Read the documentation in the lab Exercise Scenario.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the A. Datum Automated Server Installation and Deployment Strategy document.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones in the Lab Answer Key.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you should have planned an automated server installation and deployment strategy for A. Datum Corporation.

Exercise 2: Preparing the Windows Server 2012 R2 Image

Scenario

You must now implement the server deployment plan. You will install Windows DS on LON-SVR1, and then modify the default installation image to include the Web Server (IIS) role.

The main tasks for this exercise are as follows:

1. Create the image store, and map a network drive to the image store.

2. Mount the relevant image.

3. Add the Web Server (IIS) role to the image.

▶ Task 1: Create the image store, and map a network drive to the image store

1. Switch to LON-SVR1.
2. On the **Allfiles (E:)** drive, create a new folder named **Images**.
3. In the **E:\Images** folder, create a new folder named **Custom Images**.
4. On your host, in the 20413C-LON-SVR1 window, on the toolbar, click **Media**, point to **DVD Drive**, and then click **Insert Disk**.
5. In the **Open** dialog box, in the **File name** text box, type **D:\Program Files\Microsoft Learning\20413\Drives\Windows2012R2.iso**, and then click **Open**.
6. Copy **D:\Sources\install.wim** (which is the default installation .wim image) from the installation media to the newly created **E:\Images\Custom Images** folder.
7. Share the **E:\Images** folder.

8. Grant the **Administrator** user **Full Control** on the shared folder.

9. Map the network drive **Z** to **\\lon-svr1\Images** so that drive Z is the default share for all images.

▶ Task 2: Mount the relevant image

1. Open an elevated command prompt.

2. At the command prompt, type the following command, and then press Enter:

```
Mkdir c:\mounted
```

3. At the command prompt, type the following command, and then press Enter:

```
Dism /get-imageinfo /imagefile:”z:\Custom Images\install.wim”
```

Note: This command lists all images that are contained in install.wim. Notice the index for the Windows Server 2012 R2 Datacenter edition.

4. At the command prompt, type the following command, and then press Enter:

```
Dism /mount-wim /wimfile:”z:\Custom Images\install.wim” /index:4 /mountdir:c:\mounted
```

Note: This command will mount the install.wim image for offline servicing. After you mount the image, you can add drivers, add packages, or enable features. This step will take approximately five minutes for the mounting of the image finish.

▶ Task 3: Add the Web Server (IIS) role to the image

1. At the command prompt, type the following command, and then press Enter:

```
Dism /image:c:\mounted /get-features
```

Note: This command lists all available features and their state in the image file. If you want to see more detail, redirect the output in the text file by using the following command:

```
Dism /image:c:\mounted /get-features > c:\All Features.txt
```

2. At the command prompt, type the following command, and then press Enter:

```
Dism /image:c:\mounted /get-featureinfo /featurename:IIS-WebServerRole
```

Note: Notice that the state of Web Server (IIS) role is disabled.

3. At the command prompt, type the following command, and then press Enter:

```
Dism /image:c:\mounted /enable-feature /featurename:IIS-WebServerRole –all
```

Note: This command will install the Web Server (IIS) role with all the depended features. Note that the name of the role is case sensitive.

Ensure that **The operation completed successfully** message displays.

4. At the command prompt, type the following command, and then press Enter:

```
Dism /unmount-wim /mountdir:c:\mounted /commit
```

Ensure that **The operation completed successfully** message displays.

Note: This command commits the changes in the install.wim image, which you will use later for deploying to a machine that has no operating system installed. This step to commit the changes in install.wim can take approximately five minutes.

5. Leave the Command Prompt window open.

Results: After completing this exercise, you should have prepared the image, and added the Web Server (IIS) role to the image.

Exercise 3: Deploying Windows Server 2012 R2

Scenario

You will now deploy the Windows Server 2012 R2 image to LON-SVR3.

The main tasks for this exercise are as follows:

1. Install the Windows DS role.
2. Configure Windows DS.
3. Use WDSUtil to add a boot image.
4. Add an install image.
5. Configure automatic naming.
6. Launch the deployment process.

▶ Task 1: Install the Windows DS role

- On LON-SVR1, use Server Manager to install Windows DS with both the **Deployment Server** and **Transport Server** role services.

▶ Task 2: Configure Windows DS

1. In the Windows Deployment Services console, configure Windows DS integration with AD DS by accepting the default PXE options.
2. Select the **E:\RemoteInstall** folder as the default share to where all images will be stored.

▶ Task 3: Use WDSUtil to add a boot image

1. Switch back to the command prompt, type the following command, and then press Enter:

```
Wdsutil /add-image /ImageFile:”D:\sources\boot.wim” /ImageType:boot
```

Ensure that **The operation completed successfully** message displays.

2. Switch to the Windows Deployment Services console, and verify that one boot image for the 64-bit architecture is listed.

▶ Task 4: Add an install image

1. In the Windows Deployment Services console, create an image group with the name **ImageGroup1**.
2. Add the **Windows Server 2012 SERVERDATACENTER** image from the custom image file that you previously edited (**z:\custom images\install.wim**).

Note: This process of adding the install image will take 5–10 minutes.

▶ Task 5: Configure automatic naming

1. Use the Windows Deployment Services console to enable Windows DS to respond to all client computers (known and unknown).
2. In the **Windows DS Server Properties** dialog box, on the **AD DS** tab, set the naming format type as **LON-SVR%0#**.

▶ Task 6: Launch the deployment process

1. Start the 20413C-LON-SVR3 virtual machine, and press the F12 key for network service boot.
2. When the Windows Deployment Services Wizard starts, sign in to the Windows DS server with the following credentials:
   - Username: **Adatum\Administrator**
   - Password: **Pa$$w0rd**
3. Select the **Windows Server 2012 R2 SERVERDATACENTER** install image.

   Note: Step 4 is optional. You can choose to not complete the deployment process, or you can follow step 4 to complete the deployment.
4. When the installation finishes, complete the deployment process:

   a. Accept the license terms.

   b. Accept the regional settings.

   c. Specify the built-in administrator password as **Pa$$w0rd**.

   d. Sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.

   e. In Server Manager, verify that Web Server (IIS) role displays.

Results: After completing this exercise, you should have deployed Windows Server 2012 R2 by using Windows DS.

▶ Task: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20413C-LON-SVR1** and **20413C-LON-SVR3**.

Module Review and Takeaways

Review Question(s)

Question: Your organization has different server builds, none of which are identical. You have chosen to use customized images to aid in deployment. Should you think about using thick or thin images?

Question: What tools do you need to automate High Touch with Retail Media deployments?

Question: Your organization wants to implement a lite-touch deployment strategy. Aside from using MDT 2013, what tools would be useful for performing lite-touch deployments?

Real-world Issues and Scenarios

Although Windows DS provides opportunities for deploying Windows operating systems, mid-size and enterprise companies should consider implementing MDT to customize more complex migration scenarios. For zero-touch implementation, Configuration Manager provides robust, scalable, and a controlled deployment environment. Configuration Manager also enables Windows operating system deployments, and provides ongoing management on already-installed computers.

Tools

Tool	Use to	Where to find it
ACT 6.0	Check application compatibility for Windows 8	http://go.microsoft.com/fwlink/?LinkID=391890 ACT 6.0 is available for download as a component of the Windows Assessment and Deployment Kit (Windows ADK) for Windows 8.1.
Windows Assessment and Deployment Kit (Windows ADK) for Windows 8.1	Assess and deploy Windows 8.1	http://go.microsoft.com/fwlink/?LinkID=391890
Windows SIM	Create and edit answer files	Windows ADK
USMT	Migrate user settings	Windows ADK
DISM	Service .wim-based image files	Windows ADK

Best Practices

Best practice	Description
Always install the most recent security updates on the reference computer.	Starting with an up-to-date reference computer helps lessen the window of vulnerability for new computers coming online.
Implement access controls to protect bootable media.	When you create bootable media, you should always assign a password and control physical access to the media.

Best practice	Description
Use PXE service points only on secure network segments.	PXE service point require User Datagram Protocol (UDP) ports to be open on switches and servers
If you must deploy operating systems to an unknown computer, implement access controls to prevent unauthorized computers from connecting to the network.	Although provisioning unknown computers can be a convenient way to bring up multiple computers on demand, it can also allow a malicious user to become a trusted client on your network.
Reduce the size of the boot image to speed up TFTP downloads	Ensure that you prepare the boot image by using the **PEIMG.exe /prep** command.

Module 3

Planning and Deploying Servers Using Virtual Machine Manager

Contents:

Module Overview

When you plan for and deploy servers on virtual machines, you should adhere to the following high-level process:

1. Analyze existing workloads.
2. Identify application resources and requirements.
3. Configure suitable virtual machines for deployment to the best available hosts.

By using Microsoft® System Center 2012 R2 Virtual Machine Manager (VMM), you can configure, deploy, and manage a server virtualization environment.

In this module, you will learn how VMM manages Windows Server® 2012 Hyper-V® host physical computers, host groups, and the storage, networking and fabric management features of VMM. You will learn about Virtual Machine Manager libraries, and how to create reusable profiles and templates to aid in virtual machine deployment. You also will review the planning steps and considerations for a VMM deployment.

Objectives

After completing this module, you will be able to:

- Describe the features and capabilities of VMM.
- Implement a Virtual Machine Manager library and profiles.
- Plan and deploy service templates and virtual machine templates.

Lesson 1
System Center 2012 R2 Virtual Machine Manager Overview

VMM provides administrators with a single administrative tool for deploying and managing a virtualization infrastructure, including components such as hosts, storage, networks, libraries, and update servers. This infrastructure provides the foundation for managing the configuration and deployment of virtual machines.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe key features of VMM.
- Describe the VMM architecture.
- Manage hosts and host groups.
- Add hosts to VMM.
- Manage non-domain joined hosts.
- Plan VMM storage options.
- Describe VMM networking options.
- Explain VMM fabric management.
- Manage the VMM fabric.
- Explain the planning considerations for VMM.

Virtual Machine Manager

VMM provides administrators with a single administrative tool for deploying and managing a virtualization infrastructure. Using VMM, administrators can deploy and manage components such as hosts, storage, networks, and libraries. Administrators can also use VMM to update servers.

VMM provides the foundation for managing virtual machine configuration and deployment. Using VMM, administrators can manage a single virtual machine host computer, or manage as many as 1,000 hosts and 25,000 guests.

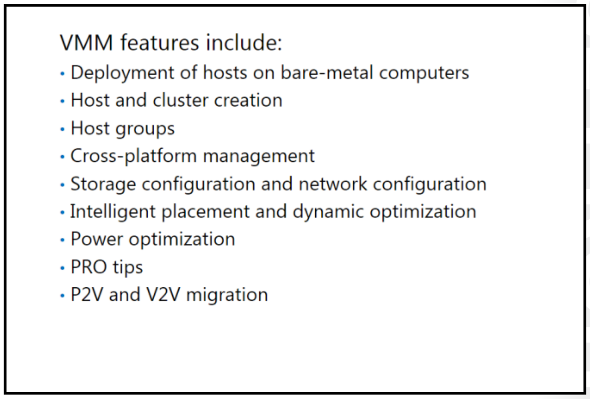
VMM features include:
- Deployment of hosts on bare-metal computers
- Host and cluster creation
- Host groups
- Cross-platform management
- Storage configuration and network configuration
- Intelligent placement and dynamic optimization
- Power optimization
- PRO tips
- P2V and V2V migration

VMM consists of a VMM server, a Virtual Machine Manager database, and a VMM console. These are the core components that a deployment requires, and you can deploy these components to a single server or to multiple servers. A later section of this module details VMM deployment prerequisites.

The following list describes some of the key VMM features:

- Deployment of hosts on bare-metal computers. You can automate deployment of Windows Server host machines on physical servers with an installed baseboard management controller (BMC), and which meet discovery and deployment prerequisites. System Center 2012 R2 VMM enables administrators to discover more information about a target host's resources, and configure more networking settings, such as logical switches.
- Host and cluster creation. You can create Hyper-V hosts and clusters easily by using the VMM console, which simplifies manual deployment and reduces the possibility of configuration errors.
- Host groups. You can group hosts for logical separation, such as business use, performance, and geographical location, and you also can apply changes to multiple hosts.
- Cross-platform management. VMM supports Citrix XenServer host and pool management, and supports VMware ESX hosts through integration with VMware vSphere.

Note: System Center 2012 VMM only works with Windows Server 2008 R2 or Windows Server 2012 hosts. You cannot use System Center 2012 VMM to manage Windows Server 2008 host machines. In addition, System Center 2012 VMM supports only VMware ESX or VMware ESXi versions 4.1 or newer, and Citrix XenServer versions 6.0 and 6.1. However, to manage Citrix XenServers, you must first install the System Center Integration Pack on the XenServer hosts, and then add the hosts in VMM.

- Storage configuration. VMM supports discovery, classification, and provisioning of storage for Hyper-V hosts, including thin provisioning capabilities. VMM storage discovery works with Storage Management Initiative Specification Common Information Model (CIM) XML, and with symmetric multiprocessing (SMP) storage providers. Additionally, System Center 2012 Service Pack 1 (SP1) VMM also supports the new Windows standards-based Storage Management Service and the Server Message Block (SMB) 3.0 protocol.
- Network configuration. Network configuration enables you to create logical networks, media access control (MAC) address pools, and supported load balancers. Additionally, VMM supports the new Windows Server 2012 network virtualization features, including the ability to run overlapping addresses on the same physical network.
- Intelligent placement. Intelligent placement helps you select an appropriate host based on the virtual machine that you are deploying, and includes ratings of hosts against expected utilization thresholds, such as percentage of CPU, I/O, and network throughput.
- Dynamic optimization. VMM can balance workloads automatically, according to configurable thresholds for core resources such as CPU, memory, disk, and network utilization.
- Power optimization. You can configure VMM to use power thresholds that you specify. This enables VMM to evaluate the performance requirements of a Hyper-V host cluster and shut down hosts if the hosts are not needed to provide adequate performance. Before shutting down the host, VMM will migrate all virtual machines to other hosts in the cluster. As performance requirements increase, VMM can restart the Hyper-V host.
- Performance and Resource Optimization (PRO) tips. PRO tips can offer preset remediation based on alerts. For example, you can use PRO tips to initiate the live migration of a virtual machine from a heavily utilized host machine to a host machine with more capacity.

- Microsoft Application Virtualization Server (Server App-V). Server App-V allows the virtualization of server-based applications. VMM has a built-in service and template designer that can help you construct single and multitier applications. You then can deploy these applications as services, which you can scale through automation.
- Physical-to-virtual migration (P2V). VMM has built-in functionality that enables you to perform P2V and virtual-to-virtual (V2V) migrations.

Additional Reading: For more information on the new features in System Center 2012 VMM, refer to the article What's New in System Center 2012 - Virtual Machine Manager, at http://go.microsoft.com/fwlink/?LinkId=253224.

VMM Architecture

VMM is a System Center 2012 component that provides a management solution for a virtualized data center. You can use VMM to create and deploy virtual machines and services to private clouds by configuring and managing your virtualization host, networking, and storage resources. By using VMM, you can discover, capture, and aggregate information about the virtualization infrastructure and enable automatic management of policies and processes. In the private cloud infrastructure, VMM helps transition enterprise IT from an infrastructure-focused deployment model into a service-oriented, user-centric environment.

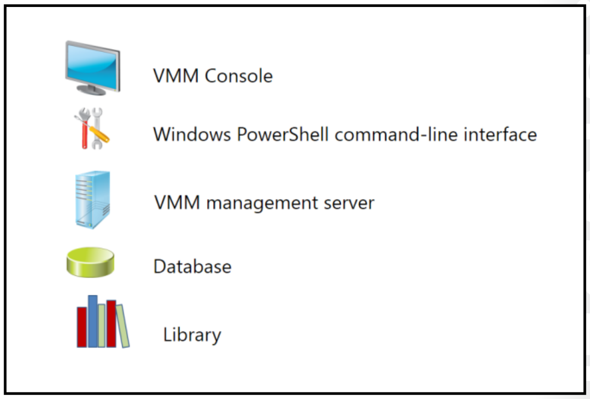

VMM architecture consists of several different, interrelated components, including:

- VMM console. The VMM console is a program that you use to connect to a VMM management server. Through the VMM console, you can view and manage physical and virtual resources, including virtual machine hosts, virtual machines, services, and library resources.
- Command shell. Windows PowerShell® is the command-line interface in which you use cmdlets to perform all available VMM functions. The VMM console is built using Windows PowerShell. You can use VMM-specific cmdlets to manage all the actions in a VMM environment.
- VMM management server. The VMM management server is the computer on which the VMM service runs. The VMM management server processes commands and controls communications with the database, the library server, and the virtual machine hosts.
- Virtual Machine Manager database. VMM uses a Microsoft SQL Server® database to store the information that you view in the VMM console, such as managed virtual machines, virtual machine hosts, virtual machine libraries, jobs, and other virtual machine-related data.
- Virtual Machine Manager library. The Virtual Machine Manager library is a catalog of resources such as virtual hard disks, templates, and profiles that are used to deploy virtual machines and services. A library server 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.

Managing Hosts and Host Groups

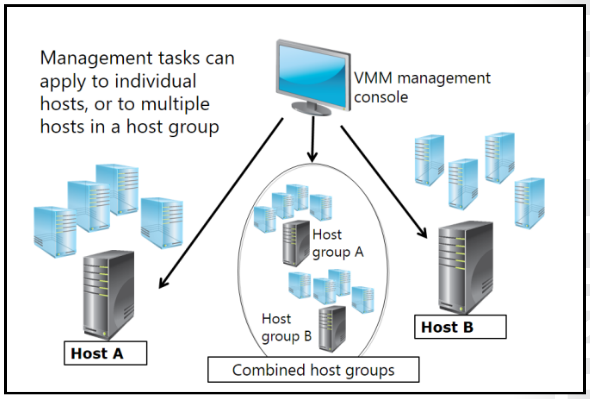
Management tasks can apply to individual hosts, or to multiple hosts in a host group
VMM management console
Host group A
Host group B
Host A
Host B
Combined host groups

Using a Hyper-V server to manage multiple virtual machines offers several advantages. The Hyper-V Manager console becomes the single, central location to conduct all virtual machine configuration and management. You can then add the Hyper-V host to VMM along with other hosts to further centralize your administrative and management oversight by creating host groups. You can then add a selected host to these groups. When you need to manage several hosts (but not all) in a particular manner, you can set distinct properties to host groups, which simultaneously configures all the hosts belonging to that host group.

Managing Hosts

For VMM to manage a Hyper-V virtualization host, you must deploy the VMM software to the host by using the Add hosts function in the VMM console. In the case of a host in a perimeter network, you deploy the agent software manually, and then add the host in the VMM console.

To deploy a Hyper-V host in a trusted domain, perform the following procedure:

1. Open the Virtual Machine Manager console, click the **VMs and Services** workspace, from the ribbon, click **Add Resources**, and then click **Hyper-V Hosts and Clusters**.
2. On the **Resource location** page, click **Windows Server computers in a trusted Active Directory domain**, and then click **Next**.
3. On the **Credentials** page, choose to either use a RunAs account (an account already configured with domain privileges) or manually enter credentials of an account with privileges to install the agent on the host server, and then click **Next**.
4. On the **Discovery Scope** page, specify computer names by entering them on separate lines in the **Computer name** field, or you can click **Specify an Active Directory query to search for Windows Server computers**, type a query, and then click **Next**.
5. On the **Target resources** page, either click each host or click **Select all**, and then click **Next**. A dialog box prompts you that you are enabling the Hyper-V role on any servers as part of the process. If you choose to enable the role, the servers reboot during the process. To close the dialog box, click **OK**.
6. On the **Host settings** page, assign the host or hosts to a Host group. (A later section of this module details host groups.) Additionally, if you have multiple VMM servers, and if another VMM environment currently is managing your host, you can reassociate the host with this environment by clicking **Reassociate**. You also can assign default placement paths, which is the location in which Windows will store new or migrated Hyper-V virtual machine files. Additionally, you can assign these paths after you add the host, and then click **Next**.
7. On the **Summary page**, confirm the settings, and then click **OK**. When the Jobs window launches, you can review the progress of the agent deployments.

Managing Host Groups

You can use host groups to organize and manage your servers. This makes it easier to apply management settings at a group level. All servers reside in the default host group named All Hosts, unless you specify another location.

Host groups may be nested. Therefore, unless you clear the inherited parent host group settings, the parent group's settings will apply to the hosted group. You can make this change in the Properties page of the selected child object.

You can create host groups by clicking the VMs and Services workspace, then right-clicking in the Navigation pane, and then clicking Create Host Group. The default host group is called All Hosts. To edit host group properties, right-click a host group, and then click Properties. From the Host group Properties dialog box, on the General page, you can edit the following settings:

- Name the group
- Move the group
- Provide a group description
- Allow unencrypted files transfers to the group

By default, a host group uses the placement setting from the parent host group. If you opt to configure custom placement rules at the individual group level, you can block inheritance by modifying the parent host-group setting.

On the Placement Rules page of the host group properties, you can assign custom placement rules. For example, you can assign custom values to hosts and virtual machines that will determine placement based upon criteria, including one of the following criteria:

- The virtual machine *must* match the host.
- The virtual machine *should* match the host.
- The virtual machine *must not* match the host.
- The virtual machine *should not* match host.

Demonstration: Adding Hosts to VMM

In this demonstration, you will see how to:

- Set the default domain Group Policy to allow domain members to become hosts.
- Add hosts to the VMM console.

Demonstration Steps

Set the default domain Group Policy to allow domain members to become hosts

1. On LON-DC1, in Server Manager, open the Group Policy Management Editor, and then edit the **Default Domain Policy**.
2. Navigate to **Computer Configuration**, **Administrative Templates\Network\Network Connections\Windows Firewall\Domain Profile**, and then apply the following settings:

   a. In Windows Firewall:

   - Allow inbound file and printer sharing exception: **Enabled**
   - Options: Type an asterisk (*****) (which indicates all IP addresses)

   b. In Windows Firewall:

   - Allow ICMP exceptions: **Enabled**
   - Options: **Allow inbound echo request**

c. In Windows Firewall:

- o Define inbound port exception: **Enabled**
- o Options: Define port exceptions, click **Show**, and under **Value**, type **5985**

3. In the Group Policy Management Editor, navigate to **Administrative Templates\Windows Components\Windows Remote Management (WinRM)\WinRM Service**.
4. In the **Allow remote server management through WinRM** dialog box, click **Enabled**, in **Options**, for both **IPv4** and **IPv6**, type an asterisk (*****).
5. Close the Group Policy Management Editor.
6. On the LON-HOST1 physical machine, update the group policy with the following Windows PowerShell cmdlet:

```
gpupdate.exe /force
```

Add LON-HOST1 to VMM

1. On LON-VMM1, open the VMM console, and add **LON-HOST1** as a Hyper-V server to the **All Hosts** node in **VMs and Services**, using the following settings:

   a. **Resource Location** page: **Windows Server computers in a trusted Active Directory**.

   b. **Credentials** page:

   - o User name: **ADATUM\Administrator**
   - o Password: **Pa$$w0rd**

   c. **Discovery Scope** page, Computer names: **lon-host1.adatum.com**

   d. **Target resources** page, Discovered computers: **lon-host1.adatum.com**

   e. **Host Settings** page: **All Hosts**

   f. **Summary** page: **Finish**

2. Observe that LON-HOST1 now displays in the VM's and Services console tree.

Managing Non-Domain Joined Hosts

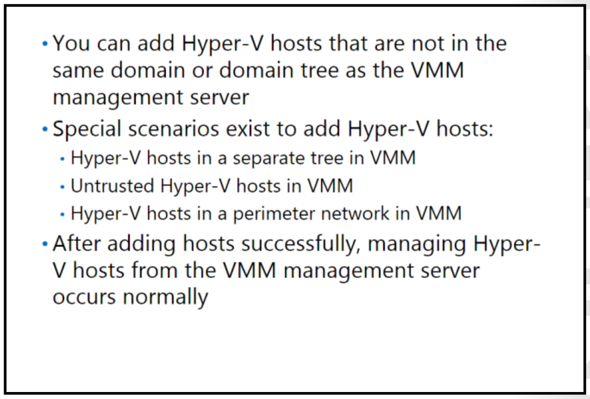

- You can add Hyper-V hosts that are not in the same domain or domain tree as the VMM management server
- Special scenarios exist to add Hyper-V hosts:
  - Hyper-V hosts in a separate tree in VMM
  - Untrusted Hyper-V hosts in VMM
  - Hyper-V hosts in a perimeter network in VMM
- After adding hosts successfully, managing Hyper-V hosts from the VMM management server occurs normally

In most cases you would want to deploy Hyper-V hosts that are part of the same overall forest structure. This enables you to utilize forest trust relationships between all domains in that forest. However, there may be times where this is not possible, such as when Hyper-V hosts are stand-alone or workgroup servers, or if they are deployed in a perimeter network or in a business partner's physical location. You might also need to deploy the Hyper-V host in a domain that is in a different tree, but still in a particular Active Directory® forest. For example, the Adatum.com domain and forest might have a separate tree named Contoso.com, and you need to deploy a Hyper-V host in Contoso.com that is managed by a VMM management server in Adatum.com.

Hyper-V Hosts in Different Trees in VMM

To add Hyper-V hosts to a VMM management server in a domain in another tree, you must do the following:

- Meet the following prerequisites for VMM:
  - The VMM service must use an account that has permission to register its service principal name (SPN) in Active Directory Domain Services (AD DS).
  - You must add the DNS suffix of the domain's tree in the TCP/IP settings of the VMM management server.
  - Only certain Group Policy settings for Windows Remote Management can be used. You can set only the following settings:
    - Allow automatic configuration of listeners
    - Turn on Compatibility HTTP Listener
    - Turn on Compatibility HTTPS Listener

Any other settings configured for Windows Remote Management in Group Policy might cause the Virtual Machine Manager agent on the Hyper-V host that you are trying to install to fail. In addition, the Allow automatic configuration of listeners setting must allow messages from any IP addresses. This could potentially risk your overall security.

- When using the Add Resource Wizard to install Hyper-V hosts, you must enter the Hyper-V host's fully qualified domain name (FQDN), and do not check the AD verification check box. VMM attempts to write an SPN into AD DS for the FQDN of the Hyper-V host. If that fails, you can attempt to create the SPN manually by typing the following commands into a Command Prompt window:

```
Setspn –A HOST/<FQDN> <NETBIOS Name>
```

For example, using LON-HOST3 in the Contoso domain, you would type the following command:

```
Setspn –A HOST/Lon-Host3.contoso.com Lon-host3
```

Then try the Add Resource Wizard again.

Untrusted Hyper-V Hosts in VMM

To add Hyper-V hosts that are in an untrusted Active Directory domain that is managed by VMM in another domain, you must do the following:

- Meet the following prerequisites:
  - Group Policy settings for Windows Remote Management cannot be enabled in an untrusted domain for setting access for a VMM management server in the trusted domain.
  - While not required, we recommend using a specific Run As account for the untrusted domain's Hyper-V hosts. This account must have local administrative rights on the Hyper-V hosts. The advantage to using a specific Run As account is that it allows you to differentiate between the permissions of trusted domain account and the permissions of the untrusted domain.
- In the Add Resource Wizard, in an untrusted Active Directory domain, select the Windows Server computer radio button, and if you have already created a Run As account, use the Run As account discussed previously. Then complete the rest of the wizard using your standard replies.

Hyper-V Hosts in a Perimeter Network in VMM

A perimeter network (also known as a *demilitarized zone* or *DMZ*), is a network subnet that normally exists between two firewalls; an outside firewall between an organization's resources and the Internet service provider (ISP), and an inside firewall that resides between an organizations internal resources and the perimeter network itself. Both firewalls have certain rules that either allow or block certain internetworking packets, depending on a table of settings that the firewall administrator has created. Certain servers that must be enabled to allow access from outside entities reside in the perimeter network, and the outside firewall has rules in the table to allow this access. The inside firewall would block the same packets trying to get into the internal resources. In most cases, the servers in the perimeter network are not domain-joined, as this can potentially provide a hacker with an entry point into the domain.

In this case, you can install the Virtual Machine Manager agent on the Hyper-V host in the perimeter network:

1. Access the VMM installation media from the Hyper-V host.
2. On the Installation Flash page, run the Optional Installations, and select the Local Agent hyperlink.
3. In the installation wizard, select the **This host is on a perimeter network** check box.
4. Supply a Security file encryption key, which is any value that you chose, similar to a password value. This encryption key will be stored locally as “SecuirtyFile.txt” in a folder of your choosing. Remember where you stored the security file.

Note: Later, when you run the Add Resource Wizard on the VMM management server, in the “To ensure that the Security.txt file is available to VMM” section, you must transfer the security file to a location that is accessible to the VMM management server.

5. Supply the IP address of the VMM management server to the installation wizard.
6. Return to the VMM management server and run the Add Resource Wizard.
7. On the Resource location page, click the radio button for Windows Server computers in a perimeter network.
8. The wizard then requests the Hyper-V host name or IP address, and the SecurityFile.txt key that you created earlier.

Once you add the Hyper-V hosts by using any of the scenarios listed above, you can now manage the hosts just as you would a local domain-joined host.

Virtual Machine Storage

A key factor when provisioning virtual machines is ensuring that the underlying storage infrastructure is reliable and can provide sufficient performance. This includes successfully managing peak utilization times such as backups, antivirus sweeps, and multiple, concurrent virtual machine startups. Storage is one of the more complicated and costly resources to manage in virtualization projects, and it benefits an organization to design storage solutions that have the flexibility to scale up and meet future growth, but not overprovision capacity.

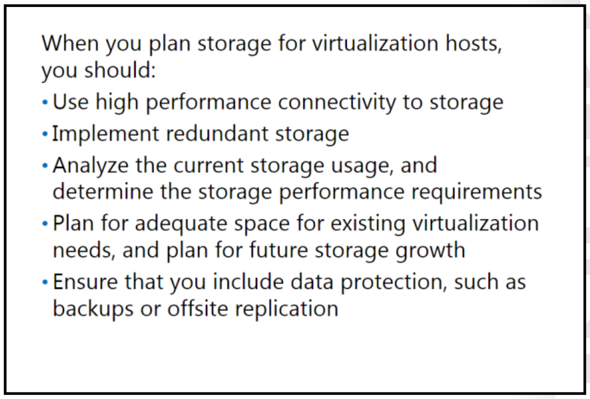

When you plan storage for virtualization hosts, you should:

- Use high performance connectivity to storage
- Implement redundant storage
- Analyze the current storage usage, and determine the storage performance requirements
- Plan for adequate space for existing virtualization needs, and plan for future storage growth
- Ensure that you include data protection, such as backups or offsite replication

Windows Server 2012 R2 builds upon and introduces new storage options for virtualization, which support small-to-midsize corporations' highly available storage solutions. These solutions were previously available only by investing in storage area network (SAN) technologies or in non-Microsoft software.

The .vhdx file performance can affect a virtual machine's performance. Servers that you otherwise provision well with random access memory (RAM) and processor capacity can still experience unsatisfactory performance if you misconfigure the storage system or it becomes overwhelmed with traffic. You should ensure that the storage design provides adequate performance, and that your design includes a plan for monitoring storage for availability and performance.

Consider the following factors when planning for storage options.

- Storage connectivity. You can locate .vhd or .vhdx files on local or remote storage. When you locate these files on remote storage, you must ensure that there is adequate bandwidth and minimal latency between the host and the remote storage. Slow network connections to storage, or connections where there is high latency, result in poor virtual machine performance.
- Storage redundancy. The volume that the .vhdx files are stored on should be fault tolerant. This should apply regardless of whether the .vhdx file is stored on a local disk or a remote SAN device. Hard disks often fail. Therefore, the virtual machine and the Windows Server 2012 Hyper-V host should remain in operation after a disk failure. Replacement of failed disks should not affect the operation of the Hyper-V host or virtual machines.
- Storage performance. The storage device where .vhdx files are stored should have excellent I/O characteristics. Many enterprises use solid-state drive (SSD) hybrid drives in a redundant array of independent disk (RAID) 1+0 arrays. This achieves maximum performance and redundancy, particularly when multiple virtual machines are running simultaneously on the same storage. This can place a tremendous I/O burden on a disk subsystem, so you need to ensure that you choose high performance storage, or your virtual machine performance may suffer. The Assessment and Planning Tool measures I/O, and its output can assist in storage planning.
- Storage capacity. If you configure .vhdx files to grow automatically, it is important that there is adequate space in which these files indeed can grow. Additionally, you need to monitor growth carefully so that you experience no service disruptions if a .vhdx file consumes all available space.
- Data protection. Consider the performance of your backup solution, its impact on your storage design, and the amount of data that your virtual machines will host. Review existing data and ensure that you can back up required virtual machines and their storage within an acceptable timeframe.
- Use of Hyper-V. Hyper-V offers flexible storage options including most of the options that Windows Server supports, for example: locally-attached storage such as Serial Advanced Technology Attachment (SATA), small computer system interface (SCSI), and SSD. Hyper-V supports remotely-connected Fibre Channel, Internet SCSI (iSCSI), and Serial Attached SCSI storage. Hyper-V also supports running virtual machines in file shares using the SMB 3.0 protocol. Shared .vhdx allows guest virtual machines to be clustered without needing iSCSI or Fibre Channel SANs. Hyper-V and VMM support live migration outside of a clustered environment, sometimes referred to as *shared-nothing live migration*.

Note: When choosing your virtualization storage options, you should examine closely all the features and components that you plan to use. You also should review carefully the prerequisites for each technology to ensure compatibility. For example, if you plan to use the Windows Server 2012 R2 Offloaded Data Transfer (ODX) feature for virtual machine SAN transfers, you cannot have Windows Server Data Deduplication or BitLocker® Drive Encryption enabled.

- Storage Availability. Using VMM, you can manage block storage and file storage for deploying and storing your virtual machines. You can do this through the use of Windows Storage Management application programming interface (API) deployment.

VMM Networking

In VMM, the networking infrastructure is comprised of a group of configurable network resources that you can use to create, model, organize, and manage your virtualized server network connectivity. The following sections describe the configurable components and their subcomponents.

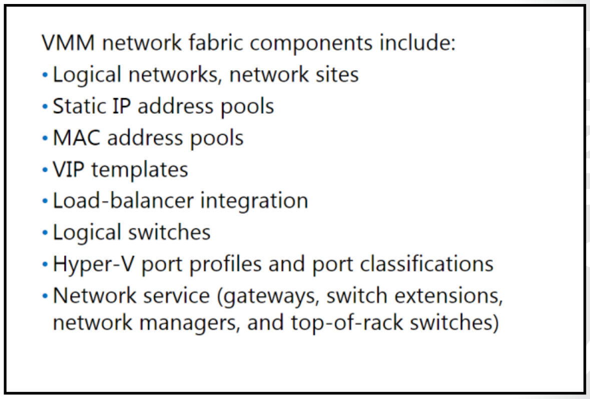

VMM network fabric components include:

- Logical networks, network sites
- Static IP address pools
- MAC address pools
- VIP templates
- Load-balancer integration
- Logical switches
- Hyper-V port profiles and port classifications
- Network service (gateways, switch extensions, network managers, and top-of-rack switches)

Logical Networks

A *logical network* is a set of logical network objects that you can use to model your network environment. You can create multiple logical networks and associate them with one or more host groups. For example, you can create a perimeter logical network, a development logical network, and a production logical network. When administrators or application administrators deploy virtual machines and services, they can select a logical network without the need to be familiar with the underlying networking infrastructure.

Network Sites

You can create network sites to associate subnets and virtual local area networks (VLANs) with a location or department. You associate sites with the logical network, and then assign the host group that can use the network site.

MAC Address Pools

VMM can assign static MAC addresses automatically to new virtual network devices on Windows-based virtual machines that are running on any managed Hyper-V, VMware ESX, or Citrix XenServer host. VMM has two default static MAC address pools: the default MAC address pool for Hyper-V and Citrix XenServer, and the default VMware MAC address pool for VMware ESX hosts. You should use the default static MAC address pools only if you set the MAC address type for a virtual machine to Static. If you set the virtual machine setting to Dynamic, the hypervisor assigns the MAC address. You can use the default MAC address pools, or you can configure custom MAC address pools that you scope to specific host groups.

Virtual IP Templates

A virtual IP template contains a load balancer and related configuration settings for a specific type of network traffic. For example, you could create a template that specifies the load-balancing behavior for HTTPS traffic on a specific load balancer manufacturer and model. These templates represent the best practices from a load balancer configuration standpoint. After you create a virtual IP template, users (including self-service users), can specify the virtual IP template to use when they create a service. When users model a service, they can select an available template that best matches the needs of their load balancers and type of application.

Load Balancer Integration

By adding a load balancer to VMM, you can load-balance requests to the service tier’s virtual machines. You can use Network Load Balancing (NLB), or you can add supported hardware load balancers through the VMM console. VMM includes NLB as an available load balancer, and it uses a round robin method for load balancing. To add supported hardware load balancers, you must install a configuration provider that is available from the load balancer manufacturer. The configuration provider is a plug-in to VMM that translates Windows PowerShell commands to API calls, which are specific to a load balancer manufacturer and model. Supported hardware load balancer devices are F5 BIG-IP, Brocade ServerIron, and Citrix Netscaler. You must obtain the load balancer provider from the load balancer vendor, and install the provider on the VMM management server.

Logical Switches

You can use logical switches to apply a single configuration to multiple hosts. You configure logical-to-Hyper-V port profiles and uplink profiles, port classification, and virtual-switch extensions. By using logical switches, you can enforce compliance among the host servers, and reduce the time required to deploy and administer hosts.

Port Profiles

You can create and use two Hyper-V port profiles in VMM:

- Virtual network adapter port profiles. You create this type of profile for virtual machines and hosts. These profiles have configurable offload, security, and bandwidth settings.
- Uplink port profiles. You configure this type of profile to use with uplink ports. You can configure the load-balancing algorithm and teaming mode.

Port Classifications

You can create port classifications, and then use them across multiple logical switches to help identify and group sets of features.

Network Service

A network service in VMM includes components such as gateways, virtual switch extensions, top-of-rack switches, and network managers. To add a network service, you must first install the associated provider, and then restart the System Center Virtual Machine Manager service.

You can configure each of the following components by using the Add Network Service Wizard:

- Gateway. In VMM, you can configure a gateway to allow network traffic in and out of a virtual machine network that is using network virtualization. You can configure this for local network routing, which routes traffic between the virtual machine network and the physical network. Alternatively, you can configure it for remote network routing, which first creates a virtual private network (VPN) connection with another endpoint of a site-to-site VPN, and then routes packets in and out of the virtual machine network through the VPN tunnel.
- Virtual switch extensions. Virtual switch extensions enable non-Microsoft vendors to add monitoring, filtering, and forwarding extensions. For example, Cisco has created the Cisco Nexus 1000V for Hyper-V. This forwarding extension enables Cisco administrators to configure networking in VMM by using familiar Cisco commands. An example of a monitoring extension is Host sFlow, which exports performance metrics using the sFlow protocol.
- Network managers. Network managers enable you to use a non-Microsoft network management console to configure forwarding extensions. With network managers, you can manage settings, such as logical networks, sites, and virtual machine networks.
- TOR switches. You can manage TOR switches by using VMM, which enable you to control physical switch ports. For example, you can create the corresponding VLAN and apply it to the physical port, thus keeping both physical and virtual switch settings synchronized.

VMM Fabric Management

- Includes network and storage infrastructure, host computers and groups, and WDS and WSUS servers
- Aggregates and abstracts everything into resources that can be consumed and deployed
- Is accessed by administrator and designated user roles in private cloud resource allocation

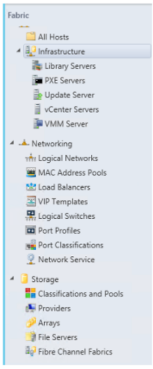

In VMM, the *fabric* is the infrastructure and services that you use to manage and deploy hosts, and that you use to create and deploy virtual machines and services to both the data center and the private cloud. The fabric includes:

- Host groups
- Networking
- Storage elements
- Pre-Boot Execution Environment (PXE)
- Windows Deployment Service (WDS)
- Windows Server Update Services (WSUS) servers
- Virtual Machine Manager libraries
- VMware ESX and Citrix XenServers

The main benefits of using the fabric are:

- Aggregate private cloud resources. The goal for the fabric is to aggregate private cloud resources in meaningful ways that enable you to deploy fabric resources more easily and comprehensively. The fabric is a logical manifestation of the networks, storage, and services that are available as resources in your cloud environment.
- Abstract your networking resources. The fabric combines logical networks with Hyper-V virtual networks to define IP address assignments and route traffic, and set up static addresses for host servers. The VMM fabric can supply IP addresses by using combinations of IP ranges, MAC address pools, and virtual IP templates. The VMM fabric also provides IP load balancer support.
- Storage. VMM uses the Windows Standards-Based Storage Management service extensively to create this storage aspect of the fabric. You can automate storage assignments across your public or private cloud, provided the storage device is supported through the Storage Management Initiative Specification (SMI-S). Additionally, if you are using Windows Server 2012 R2 with the File Server role and the iSCSI Target Server role enabled, you can attach storage, create storage pools, create discs and volumes, and create iSCSI disks and targets, which you can then add into your fabric storage.
- Management. The VMM console has a workspace devoted to the fabric, which enables you to manage the overall fabric that makes up all of these resources mentioned in this list. In VMM, the fabric workspace has an additional element named *infrastructure*. Your VMM management servers, PXE servers, VMware servers, and library servers are now located in this infrastructure.

Demonstration: Managing the VMM Fabric

In this demonstration, you will see how to manage Fabric resources by creating a:

- Logical network
- Logical network IP pool

Demonstration Steps

Create a logical network

1. On LON-VMM1, launch the Virtual Machine Manager console. Click the **Fabric** workspace, and on the ribbon, click **Create Logical Network**.
2. On the **Name** page, in the **Name** text box, type **Adatum UK**, click **Allow new VM networks created on this logical network to use network virtualization**, and then click **Next**.
3. On the **Network Site** page, add two **Network Sites** with the **Host groups that can use this network site** section set to **All Hosts**.
4. In the **Associated VLANs and IP subnets** area, use the following settings for the Central network site:
   - Network site name: **Central**
   - VLAN: **0**
   - IP Subnet: **192.168.1.0/24**
5. Repeat step 4 using the following settings
   - Network site name: **West Side**
   - VLAN: **0**
   - IP Subnet: 192.168.2.0/24
6. Click **Next**, click **Finish**, and close the Jobs window.

Create a logical network IP Pool

1. From the workspace, on the ribbon, click **Create IP Pool**.
2. On the **Name** page, in the **Name** text box, type **Central IP Pool**, verify that the logical network is **Adatum UK**, and then click **Next**.
3. On the **Network site** page, click **Use an existing network site**, and ensure that **Central** is selected. Verify that the **IP subnet** selected is **192.168.1.0/24**, and click **Next**.
4. On the **IP address range, Gateway, DNS and WINS** pages, review the options, and then click **Next**.
5. On the **Summary** page, click **Finish**, and then close the Jobs window.
6. Using the same steps as above, create another IP Pool with the name West Side instead of Central.

VMM Planning Considerations

When you plan a VMM deployment, you should consider the following factors:

- Number of hosts
- Number of branch sites with hosts
- Security, administrative groups, and self-service options that you require
- Availability and recovery time that each of the components require

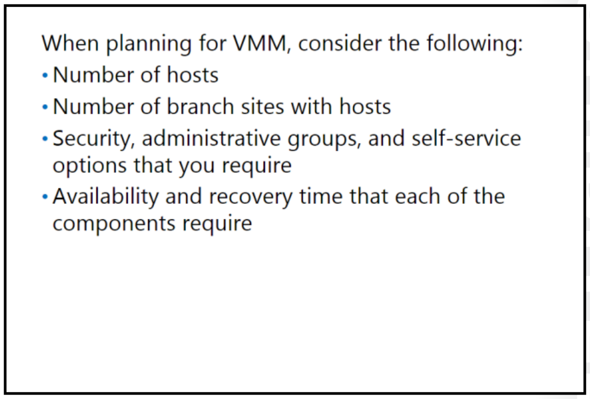

When planning for VMM, consider the following:

- Number of hosts
- Number of branch sites with hosts
- Security, administrative groups, and self-service options that you require
- Availability and recovery time that each of the components require

The number of hosts determines the physical or virtual resources that each component server in the VMM deployment requires. In VMM, a VMM management server has the capacity to manage 1,000 hosts and 25,000 virtual machines. However, the demand on a single management server would suggest that you should use multiple VMM management servers. You can use System Center 2012 App Controller with up to five VMM VMM management servers. Therefore, in theory, you could manage resources of over 500,000 virtual machines. If your deployment has thousands of hosts, you should consider contacting your regional Microsoft office for guidance on a personalized deployment to fit your environment.

The number of branch sites with hosts, and the wide area network (WAN) link capabilities between the branches and the VMM management server determines if you should have a single VMM deployment with multiple Virtual Machine Manager library servers, or if you should have individual VMM deployments at each branch.

VMM offers delegated administration and self-service. You can use App Controller, Service Manager, or your own customized portals to provide self-service to your users. When you determine what type of VMM deployment is appropriate for your environment, you can then plan a self-service deployment that is appropriate for the design. For example, App Controller can span five VMM deployments. However, your security requirements may require that you have an App Controller deployment for each VMM deployment.

The availability and recovery time for VMM components is also important when determining the topology for your VMM deployment.

VMM is a cluster-aware application that you can configure to be highly available. SQL Server is also cluster-aware, and you can install the Virtual Machine Manager library server on a Windows file server cluster, but not on the same cluster that hosts a clustered VMM instance. System Center 2012 Data Protection Manager (DPM) can back up your VMM components and, if required, you can locate DPM at a remote site and use it to restore one or more offsite components.

If you are deploying VMM, you should consider the following factors:

- The Virtual Machine Manager database no longer supports SQL Server Express. Therefore, you must move your database to a supported version of SQL Server.
- A Windows Deployment Services (WDS) server is required for deploying Hyper-V hosts to bare-metal computers, which means a computer that does not have an operating system.
- At least one library server is necessary, but you should consider at least one Virtual Machine Manager library for each site that you separate with a low-speed WAN link.
- You should use WSUS or System Center 2012 R2 Configuration Manager for update management.
- App Controller has replaced the self-service portal. There is no longer an upgrade path from existing self-service portals to App Controller.
- System Center 2012 R2 Operations Manager is required to use VMM reporting, and to leverage PRO tips.
- Managing VMware ESX and VMware ESXi hosts requires that you integrate VMware vSphere. If you need more than the maximum number of hosts for business or network reasons, you must have multiple VMM servers. You can use App Controller to view resources for up to five VMM servers.
- Consider which VMM services you use in your topology, and review the associated ports that VMM uses to communicate between its components. Ensure that firewalls are not blocking any ports, and determine whether the component coexists with another application that these ports review. If you need to amend a default port, ensure that you update the associated firewall rules.

Lesson 2
Implementing a Virtual Machine Manager Library and Profiles

The Virtual Machine Manager library is a catalog that provides access to file-based resources that are required for building virtual machines. These file-based resources can be Sysprep scripts, .iso image files, and virtual hard disks that your library servers store. From the Virtual Machine Manager library, you also can manage virtual machine templates, guest operating system profiles, and hardware profiles that reside in the Virtual Machine Manager database. You also can store virtual machines in the Virtual Machine Manager library when you are not using them.

VMM profiles are very important components of rapid virtual machine deployment. Instead of configuring various virtual machine and operating system settings each time you deploy a new virtual machine, you can use preconfigured values from profiles. The various profiles are Virtual Machine Manager library resources. System Center 2012 SP1 VMM and System Center 2012 R2 VMM have several new profiles.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe Virtual Machine Manager libraries.
- Plan for Virtual Machine Manager library servers.
- Describe considerations for configuring hardware and compatibility profiles.
- Describe guest operating system profiles.
- Configure profiles.
- Describe application and SQL Server profiles.
- Deploy virtual machines using profile objects.
- Explain planning considerations for implementing profile objects.

Virtual Machine Manager Libraries

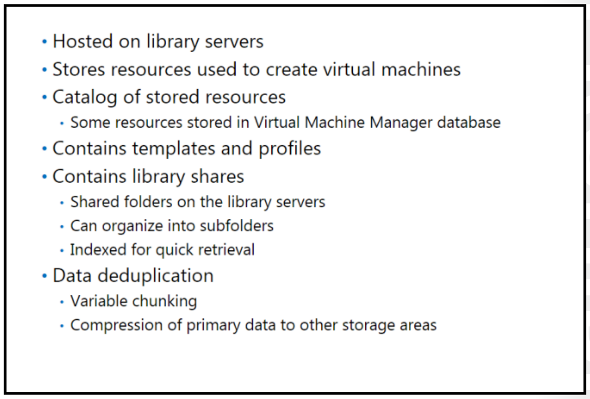

- Hosted on library servers
- Stores resources used to create virtual machines
- Catalog of stored resources
  - Some resources stored in Virtual Machine Manager database
- Contains templates and profiles
- Contains library shares
  - Shared folders on the library servers
  - Can organize into subfolders
  - Indexed for quick retrieval
- Data deduplication
  - Variable chunking
  - Compression of primary data to other storage areas

The Virtual Machine Manager library is a catalog of resources that you can use to store objects that are not running or associated with a host. You also can repeatedly use those resources for building new virtual machines. The Virtual Machine Manager library contains the following files:

- Files that are stored on library shares
- Templates for services and virtual machines
- Application profiles
- Capabilities profiles
- Guest operating systems profiles

- Hardware profiles
- Physical computer profiles
- SQL Server profiles that are stored on the Virtual Machine Manager database

There are only two places in VMM where an object can reside and be managed: the object can be registered to a host, or it can be stored in the Virtual Machine Manager library.

The Virtual Machine Manager library server hosts the Virtual Machine Manager library. When you install VMM, the VMM server is configured as a default library server. The VMM server indexes files that are stored on library shares. You cannot remove or modify the default library server that is created during the installation process. However, you can add additional library servers, if necessary.

Each library server can have one or more library shares. A *library share* is a file share that contains the resources that you use to build virtual machines. When you add a new library share, the New Library Share Wizard does not create the share for you. Instead you must create and configure a file share before adding it as a new library share.

You can organize content in a library share by creating subfolders. This is similar to creating folders in a file share. However, if the folders do not have any content, they do not display in the VMM console.

You can copy resources such as virtual hard disks and .iso files to a file share by using File Explorer. However, when you add new files to a library share, they are not immediately available. The VMM server must refresh the content before it displays in the VMM console. Content refreshes (indexed) once per hour, by default. One hour is the minimum setting possible, but you also can trigger a refresh manually.

During a library refresh, VMM indexes files that are stored on library shares, and then updates the Library view and resource listings. Not all files are indexed, and not all indexed files display in the Library view.

If any of the library resources are attached to a virtual machine and VMM indexes the configuration file for that virtual machine, the resources display as part of the virtual machine rather than as individual components.

You can create other resources, such as templates and profiles, from the VMM console. These resources are metadata that exist only in the Virtual Machine Manager database and not in the library share file system. However, they are visible in the Library view.

Enabling Data Deduplication

Data deduplication is a new Windows Server 2012 feature. It is designed for use on industry-standard hardware, and does not require extensive server resources. Data deduplication uses variable chunking that ranges from 32 kilobytes (KB) to 128 KB. It also uses compression of primary data to other storage areas from one disk to another. You can run data deduplication on a small server with a single CPU, 4 gigabytes (GB) of RAM and a Serial ATA (SATA) drive. By placing the library share on a separate hard drive, you can grow the drive to accommodate various library components. By turning on per-volume data deduplication, you can achieve significant space savings on moderately used libraries.

Planning Virtual Machine Manager Library Servers

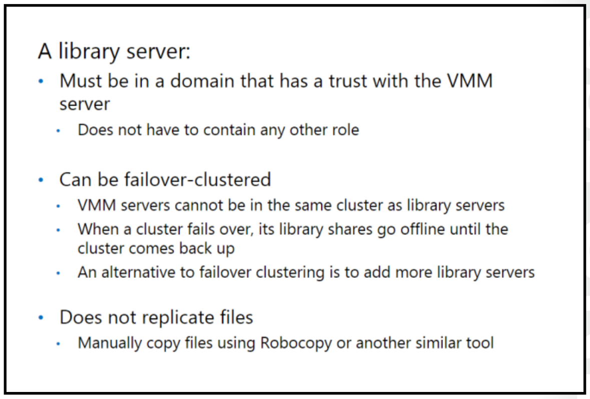

A library server:

- Must be in a domain that has a trust with the VMM server
  - Does not have to contain any other role
- Can be failover-clustered
  - VMM servers cannot be in the same cluster as library servers
  - When a cluster fails over, its library shares go offline until the cluster comes back up
  - An alternative to failover clustering is to add more library servers
- Does not replicate files
  - Manually copy files using Robocopy or another similar tool

A *library server* is a central repository (or *storage area*), of the resources that you use to create virtual machines. By storing these resources centrally, you can simplify the process of creating virtual machines. Additionally, you can provide security for the resources.

When you install VMM, the VMM server is configured automatically as a default library server. Additionally, VMM creates a default library share during the installation process. You cannot remove or modify this library share, and the default library server might be the only library server you ever need. This typically is the case for small and medium size environments, which typically do not contain multiple sites that require bandwidth utilization considerations. However, you can add more library servers and library shares, depending on your current business needs and objectives, and to scale out as your virtual environment grows.

Each library server can have multiple library shares. To enhance performance and reduce network traffic during virtual machine creation, you should store the files that you use to create virtual machines, near the hosts that you use to stage virtual machine creation.

You can associate library servers with specific host groups. For example, you may have a library server that you dedicate to resources in a test lab environment. In this case, you will associate the library server with the host groups that contain the hosts for the lab environment. A library server should have fast network connectivity to the host group with which you associate it.

A library server must meet the following requirements:

- The library server must have Windows Server 2008 R2 with SP1 or newer Windows Server operating system. For highly available file servers, the failover cluster must have been created in Windows Server 2008 R2 or newer.
- The library server must be in an Active Directory domain that has a two-way trust relationship with the VMM server's domain.

VMM does not support file servers that you configure with the case-sensitive option for Windows Services for UNIX (NFS Case Control is set to Ignore).

The Virtual Machine Manager library server role does not have to run any other VMM role. It just needs to be a file server.

Highly Available Library Servers

To make your Virtual Machine Manager library highly available, you can use two approaches. The simpler solution is to deploy multiple Virtual Machine Manager library servers with redundant content. By using that approach, you still have library resources available on another server if one server fails. However, the biggest drawback to this solution is the synchronization between two (or more) library servers. When you add a resource to one library server, you also must add it manually to other library servers. Alternatively, you can use a script with scheduled tasks to automate this process.

A second drawback to deploying multiple Virtual Machine Manager library servers is that you use more disk storage for duplicated content. However, this approach still can be appropriate if you do not have significant library resources, and if these resources do not change frequently.

Another approach to making your Virtual Machine Manager library highly available is to use failover clustering technology from Windows Server 2008 and newer Windows Server operating systems. By using the failover clustering technology, you can make a file server failover cluster that can provide high availability to Virtual Machine Manager library resources.

Configuring Hardware and Capability Profiles

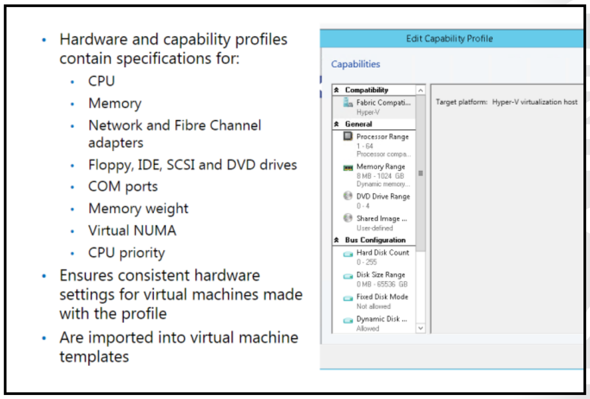

- Hardware and capability profiles contain specifications for:
  - CPU
  - Memory
  - Network and Fibre Channel adapters
  - Floppy, IDE, SCSI and DVD drives
  - COM ports
  - Memory weight
  - Virtual NUMA
  - CPU priority
- Ensures consistent hardware settings for virtual machines made with the profile
- Are imported into virtual machine templates

Edit Capability Profile

Capabilities

Compatibility

Fabric Compati...
Hyper-V

General

Processor Range
1 - 64
Processor compa...

Memory Range
8 MB - 1024 GB
Dynamic memory...

DVD Drive Range
0 - 4

Shared Image ...
User-defined

Bus Configuration

Hard Disk Count
0 - 255

Disk Size Range
0 MB - 65536 GB

Fixed Disk Mode
Not allowed

Dynamic Disk ...
Allowed

Target platform: Hyper-V virtualization host

Hardware Profile

In VMM, a *hardware profile* is a library resource containing hardware specifications that you can apply to a new virtual machine or a virtual machine template. A hardware profile can contain the following specifications:

- Cloud compatibility
- CPU, memory
- Network and Fibre Channel adapters
- Floppy drives
- IDE drives
- SCSI drives
- DVD drives
- COM ports
- Memory weight
- Virtual non-uniform memory access (NUMA)

The hardware profile also can contain the priority given to the virtual machine when allocating resources on a virtual machine host. By using hardware profiles, you can ensure consistent hardware settings in virtual machines.

You can update any existing hardware profile to modify settings for one or more virtual machine hardware components. After you change the profile, any new virtual machines that you create by using that hardware profile use the updated hardware configuration settings. Changes do not affect existing virtual machines that were created earlier by using this profile, nor do they affect settings on a template or virtual machine into which you imported this profile earlier. VMM maintains no association with the hardware profile after you create a virtual machine or template.

You can create a hardware profile by using the new hardware profile action in library view, or you can save a new hardware profile based on the hardware configuration of a virtual machine or a template. You also can create the hardware profile while creating a new virtual machine or virtual machine template.

You can create hardware profiles that import a standard hardware configuration into a template or a virtual machine. The options are the same whether you update the hardware configuration of a virtual machine, a hardware profile, or a template. You manage hardware profiles in the Library workspace.

You can also create a new hardware profile by using the Hardware Profile Wizard. To access the Hardware Profile Wizard, right-click the Hardware Profiles element in the Profiles node of the Library workspace console tree. The wizard has two pages. On the General page, you enter the name and description of the new hardware profile, and on the Hardware Profile page, you can select numerous elements to preconfigure the hardware aspects of a deployable virtual machine.

Capability Profile

In the Library workspace, under the Profiles node of the console tree, there is a Capability Profiles sub node. When selected, this node displays three existing capability profiles in the Profiles details pane: ESX Server, used for VMware hosts, Hyper-V, used for Microsoft Hyper-V hosts, and XenServer, used for Citrix XenServer hosts. You can use these profiles to provide the built-in fabric capability profile for their respective virtual machine host platforms. These profiles are set at read-only access level, and you should not modify them unless you want to globally change the host platform capability profile permanently.

Instead of altering these default profiles, you can create your own capability profiles. For example you might have a Hyper-V cluster that is used in a private cloud. To ensure that everything that is put in this cloud is configured as highly available, you can create a capability profile.

To create a capability profile, navigate to the Library workspace, click Create on the home tab of the ribbon, and from the selection menu, click Capability Profile. This brings up the Create Capability Profile Wizard. On the General page, you provide a name and optional description of the profile. On the Capabilities page, you will find the following options:

- Fabric capability. VMM ensures that the settings provided are compatible with the destination locations. The three options in this setting are Hyper-V, ESX Server, and XenServer virtualization host.
- Processor Range. You use this option to select the number of processors that the host will use. You can use a default range or select a minimum and maximum number of processors. You can also provide compatibility with different processor versions, and limit the processors that a virtual machine can use. You have the option of using the default range, selecting a user-defined or required processor compatibility, or disabling the capability altogether.
- Memory Range. Use this setting to specify how much memory should be allocated to the virtual machine, or let the virtualization host manage the amount dynamically within a range.
- DVD Drive Range. You use this option to set the number of DVD drives that you can use.
- Shared Image Mode. Use this option to enable sharing between virtual machines of .iso image files that are stored in the Virtual Machine Manager library.
- Hard Disk Count. Use this option to set the number of virtual hard disks in use. The maximum number allowed is 255.
- Disk Size Range. This option sets the size of virtual hard disks. The maximum size is 64 terabytes (TB).
- Fixed Disk Mode. You use this option to select the fixed, dynamic, or differencing disks option.
- Dynamic Disk Mode. This option is identical to the Fixed Disk Mode.
- Differencing Disk Mode. This option is identical to the Fixed Disk Mode.
- Network Adapter Range. This option allows you to select up to 12 Network Adapters.
- Network Optimizations. By selecting this option, you enable network optimization
- Availability. This option enables you to select a highly available virtual machine mode. When you configure VMM to be highly available, VMM attempts to place the virtual machine on a clustered server.

Configuring Guest Operating System Profiles

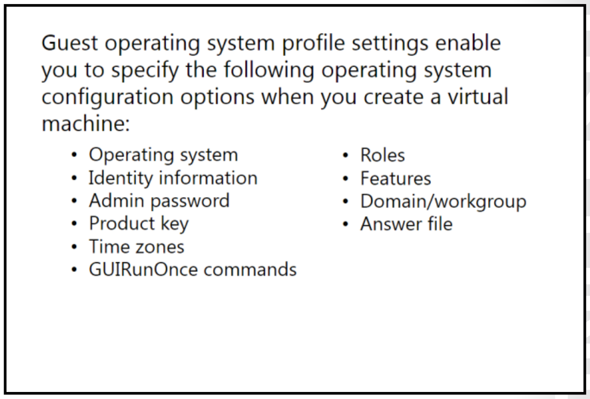

Guest operating system profile settings enable you to specify the following operating system configuration options when you create a virtual machine:

- Operating system
- Identity information
- Admin password
- Product key
- Time zones
- GUIRunOnce commands
- Roles
- Features
- Domain/workgroup
- Answer file

In a virtual environment, a guest operating system runs on a virtual machine, and the host operating system runs on the physical host computer on which you deploy one or more virtual machines. In VMM, a guest operating system profile is a collection of operating system settings that you can import into a virtual machine template to provide a consistent operating system configuration for virtual machines that you create from that template.

To ensure standard settings for the operating systems on virtual machines, you can create guest operating system profiles. Guest operating system profiles are Virtual Machine Manager database objects. They are not associated with any physical files. You configure the profiles in the Library workspace, where they display in the Profiles node. You also can view templates by selecting the Templates node in the Library workspace console tree.

You can create a guest operating system profile by using the New Guest OS Profile Wizard in the Library workspace, or you can specify guest operating system settings while you are creating a template. After you create a template, VMM does not maintain an association between the template and the guest operating system profile that was used with it. Any changes that you make to the guest operating system profile affect only new templates that you create after you make changes.

The following settings are available to a guest operating system profile:

- Operating System. You use this setting to specify the operating system of the virtual machine. VMM provides you with a drop-down list of 37 separate operating system editions, from the Microsoft Windows® 2000 Server to Windows Server 2012 R2 Datacenter.
- Identity Information. Use this field to add the computer name. You can provide a pattern for the computer name here. For example, consider a scenario where you are creating a profile to deploy different virtual machines that are running Windows Server 2012 in a Server Core installation of Windows Server 2012. You could type W2012-Core##, and then use this as a template where the first server would be named W2012-Core01, the second server would be named W2012-Core02, and so on. You can also let VMM assign a random name of numbers and letters by typing an asterisk in the computer name text box.
- Admin Password. This setting offers you three choices. You can select no local administrator credential required, provide the specific password of the local administrator account, or you can use the run as account for the local administrator account.
- Product Key. You can use this setting to specify the product key for use with the virtual machine. If you use a multiple activation key (MAK) key, record the number of virtual machines that you create with it to avoid running out. If you use a Key Management Server (KMS) key or Active Directory key server, ensure that you set up the infrastructure to support it.
- Time Zone. This setting provides a drop-down list to select the specific time zone.
- Roles. Use this setting to select various server roles available to Windows Server 2008 R2 and newer Windows Server operating systems. The roles are listed alphabetically, and the various role services that are available to that role are included directly underneath and tabbed to the right.
- Features. Similar to the Roles option above, you can select from features that are available on Windows Server 2008 R2 and newer Windows Server operating systems. If a feature has sub-elements, these are included underneath the feature name and tabbed to the right.

- Domain/Workgroup. By default, the Guest OS Profile Wizard chooses a workgroup named Workgroup. You can supply a FQDN. If so, you must supply credentials for a user who is allowed to join a computer to a domain, or you can select the Run As account credentials.
- Answer File. You can set additional settings typically found in installation answer files. You can attach a Sysprep.inf file for Windows Server 2003 and older Windows Server operating systems, or an Unattend.xml file for Windows Vista® and newer Windows client operating systems. The answer file must be stored on the library share.
- GUIRunOnce commands. You can use these commands to run a command automatically the first time a user logs on. Normally these would be command line commands, executables, and scripts. You can add as many commands as deemed necessary.

The wizard also includes a Dependencies page that is empty by default, and an Access page, where you can specify the users with whom the profile can be shared.

Demonstration: Configuring Profiles

In this demonstration, you will see how to create a guest operating system profile and create a hardware profile.

Demonstration Steps

- Navigate to the Library workspace, and create a Guest OS Profile and Hardware Profile with the following settings:
- Guest OS Profile
  - Name: **DemoGuestOS**
  - Description: **Demonstration creating a Guest OS profile**
  - Operating System: **Windows Server R2 2012 Standard**
  - Identity Information, Computer name: **WS2012-Core###**
  - Admin Password: **Specify the password of the local administrator account**
  - Password: **Pa$$word**
- Hardware Profile
  - Name: **DemoHWProfile**
  - Description: **Demonstration creating a hardware profile**
  - Compatibility: **Hyper-V**
  - Memory: **Dynamic**
  - Maximum memory: **1024**
  - Network Adapter 1: **External Network**

Configuring Application Profiles and SQL Server Profiles

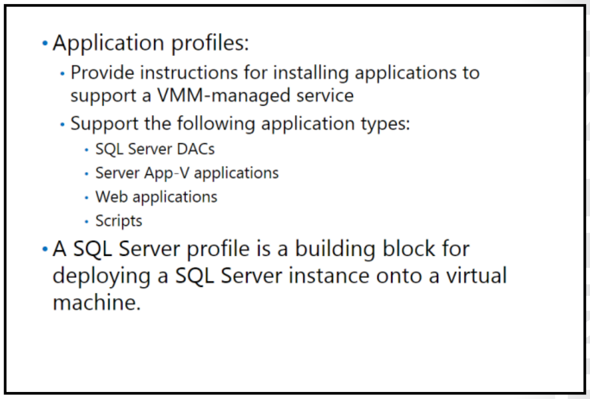

- Application profiles:
  - Provide instructions for installing applications to support a VMM-managed service
  - Support the following application types:
    - SQL Server DACs
    - Server App-V applications
    - Web applications
    - Scripts
- A SQL Server profile is a building block for deploying a SQL Server instance onto a virtual machine.

Application profiles provide configuration instructions for installing specific application types. Application profiles support the following application types:

- SQL Server data-tier applications (DACs)
- Server App-V applications
- Web applications
- Scripts

SQL Server DACs

SQL Server 2008 R2 supports a new package type called a DAC. A *DAC* contains all of the database and instance objects that the application uses, and is typically targeted towards departmental-based applications.

SQL database developers create DACs by using one of the following methods:

- Author and build a DAC using the SQL Server data-tier application project type that is available in Microsoft Visual Studio®.
- Extract a DAC from an existing database by using the Extract Data-Tier Application Wizard in the SQL Server Management Studio.

After developers create DACs, they can import the DACs into the Virtual Machine Manager library, which is then accessible from the application profile.

Server App-V

Server App-V is a technology that creates virtual application packages that you can then deploy to servers that run the Server App-V agent. A virtual application package does not require a local installation; however the package runs as if it is a locally installed application.

Web Applications

A web application is a package that is stored within the Virtual Machine Manager library. It contains the content, websites, certificates, and registry settings of a web-based application. You can package and deploy web applications with the Microsoft Web Deployment Tool. VMM also uses this tool to deploy web applications as a service when deploying a web application as specified in an application profile.

Scripts

When deploying a virtual machine as part of a service, you also can use the application profile to run scripts. You use scripts during the preinstallation and the postinstallation phases of a specific application. For example, you might need to copy updated configuration files to a deployed web application, or you may have to run specific virtual application commands to finalize a virtual application deployment. You can also use scripts to help you with preconfiguration or postconfiguration tasks when you uninstall applications. Scripts must be available in the Virtual Machine Manager library as a resource package.

To create an application profile, complete the following steps:

1. Open the VMM console, and then click the **Library** workspace.
2. In the navigation pane, expand **Profiles**, and then click **Application Profiles**.
3. On the ribbon, click **Create**, and then click **Application Profile**. The **New Application Profile** dialog box opens.

4. On the **General** page, provide a name and description for the application profile.
5. In the **Compatibility** drop-down list box, click **General** to allow for all types of supported applications in the profile. Alternatively, you can use the **SQL Server Application Host** selection if you are using this application profile to deploy a SQL Server DAC to an existing SQL Server computer. Selecting the second option enables you to add only SQL Server DAC packages and SQL Server scripts.
6. On the **Application Configuration** page, click **OS Compatibility**, and then select the guest operating systems that are compatible with the application.
7. Click **Add**, and then select the appropriate application type. Note that you can add an application script only after you have added an application.
8. Click **OK** to accept the application configuration settings.

You can add one or more applications as required by the service that you are configuring.

VMM allows you to configure a SQL Server instance when you deploy a virtual machine as part of a service. The process for installing and configuring a SQL Server instance includes a number of components as described in the following high-level steps:

1. Prepare a SQL Server image. The virtual machine that you deploy must contain a version of SQL Server 2008 R2 or SQL Server 2012 that you prepared previously using Sysprep. SQL Server 2008 R2 provides a built-in Sysprep functionality that you can use to deploy and configure SQL Server.
   - You can use SQL Server 2012 using System Center 2012 SP1 VMM or System Center 2012 R2 VMM only. You cannot use the original System Center 2012 – VMM version.
2. Create a SQL Server profile. The SQL Server profile contains configuration settings such as the instance name and ID, product key, media source, SQL Server administrators, and service account designations.
3. Create a virtual machine template. The virtual machine template specifies the hardware, operating system, and SQL Server profile that you use to deploy to a new virtual machine.
4. Create a service template. A service template provides the foundation for deploying a virtual machine and for using the SQL Server profile to configure the instances that are defined within the profile settings. A *service* is a set of virtual machines that you configure and deploy together to support specific infrastructure requirements. For example, you may have a multitier web-based application that requires a SQL Server database. A service template gathers all of the configuration settings into a single managed entity for the multiple servers. You can only configure and deploy a virtual machine with SQL Server when you deploy the application as a service.
5. Deploy the service. Deploying the service essentially deploys and configures all servers and applications associated with the service.

Before you can deploy a SQL Server virtual hard drive image, you must prepare the image by using SQL Server Sysprep. You run SQL Server Sysprep prior to running Windows Sysprep to create an image that includes a prepared operating system and an unconfigured SQL Server installation.

SQL Server Sysprep is a two-step installation process. It begins with image preparation. During image preparation, SQL Server Setup installs the product binaries without configuring any SQL Server settings for the instance that is being prepared. After this first step completes, Windows Sysprep begins, and the image is then captured.

You perform the second step of the installation process during image deployment. After you deploy an image to a virtual machine, you can proceed with the final installation and completion of a SQL Server–prepared instance. VMM uses the SQL Server profile that you prepared to provide the configuration settings for each Sysprepped instance in the image.

The SQL Server profile provides most of the common settings for use during deployment. However, you can also use a SQL Server configuration file to provide the additional configurations for settings that are not available in the SQL Server profile. A SQL Server configuration file is an .ini file, which is similar to a Windows operating system answer file (Unattend.xml). If you use a SQL Server configuration file, you must save it to a Virtual Machine Manager library share so that it is available to the template.

Demonstration: Deploying Virtual Machines Using Profiles

In this demonstration, you will see how to deploy virtual machines using profiles.

Demonstration Steps

1. In the VMs and Settings workspace, create a new virtual machine with the following settings:
   - Source: **SmallCore.vhd**
   - Name: **DemoProfileVM**
   - Description: **Demonstration using profiles to create a virtual machine**
   - Hardware profile: **DemoHWProfile**
2. After the virtual machine creation job starts, return in about 10 minutes. Connect to the **DemoProfileVM** in Hyper-V Manager, and skip the **Enter Product ID** page.

Planning Considerations for Implementing Profile Objects

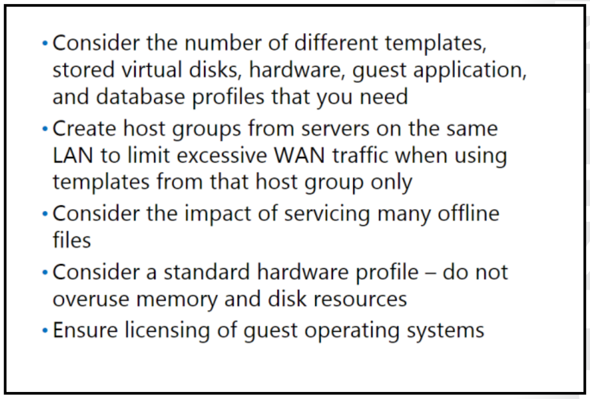

- Consider the number of different templates, stored virtual disks, hardware, guest application, and database profiles that you need
- Create host groups from servers on the same LAN to limit excessive WAN traffic when using templates from that host group only
- Consider the impact of servicing many offline files
- Consider a standard hardware profile – do not overuse memory and disk resources
- Ensure licensing of guest operating systems

As part of your virtualization strategy and infrastructure design, consider the number of different hardware, guest application, and database profiles you will need. Plan how many different operating systems you will deploy, and where you will need to store your files. The following list summarizes some considerations for working with profiles:

- Working with library items. Consider creating a plan for the number of templates you think you will need, and configure some or all of these before starting deployment. Consider keeping only the number of .vhdx files that you require. In a mixed host environment, consider having both .vhd and .vhdx formats. Remove legacy and unused profiles and templates. Back up a library occasionally, so if you must recover an older image, you can retrieve it from the backup copy.
- Try to keep the Virtual Machine Manager library tidy and prevent virtual sprawl. Remove unused virtual machines and virtual hard drives. Virtual sprawl includes offline files, which can end up being stored across file and infrastructure servers other than hosts.
- If you need to have the same templates and files across multiple Virtual Machine Manager libraries, you can send large files offline, and then import them where required. If you need to avoid using slow WAN links, set up equivalent objects at multiple locations for virtualization deployment resources that you need and do not want deployed over a WAN.

- Performance. Consider the impact of servicing many offline files. In larger organizations, collaborate to ensure that someone is not servicing images while someone else is trying to deploy images. You can designate maintenance windows during which you can perform maintenance functions. This helps prevent someone attempting to deploy images that you are working on.
- Consider an appropriately sized hardware profile. If you set the base configuration for all of your virtual servers with more memory, processors, and disk space than is necessary, then resources will be wasted, and the full value of virtualization will not be achieved.
- Licensing. You can use the guest operating system profiles to help enforce licensing requirements. For example, you can preconfigure an image for MSDN®, the Microsoft Developer Network, and then assign this to the developers who have the MSDN agreement. Consider licensing when using a template that is based on another machine, and ensure that only the people that should use a template are using it.
- Systems integration, automation, and self-service. Other applications deploy from VMM and its libraries. If necessary, create multiple libraries with appropriate security, and ensure that you keep the deployed files and images up to date.
- Service Templates. When building services for scale-out applications, consider versions and revisions, and try to keep them consistent. For example, if you have a template for a three-tier application, when you are updating tiers, you need to remember to increment the revisions appropriately.

Lesson 3
Planning and Deploying VMM Services

Services are essentially a set of virtual machines, the applications to install on the virtual machines, and the networking configuration needed for the service. You configure and deploy them together and then manage as a single entity. You can create and modify service templates that allow this deployment and management. Additionally, System Center 2012 VMM introduces the Service Template Designer, which offers a graphical management tool to create and configure virtual machine templates. In addition, you might want to change the scale of a service, update a service, or even import a service template into another VMM infrastructure.

In this lesson, you will focus on planning services and their virtual machines, and then deploying them using service templates.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe virtual machine templates.
- Describe service templates.
- Configure a virtual machine and service templates.
- Explain how to design service templates.
- Explain how to service and update service templates.
- Explain how to deploy virtual machines by using App Controller.

Virtual Machine Templates

The Virtual Machine Manager Library Templates node contains three sub nodes: service deployment configuration (this node is a storage location), service templates, and virtual machine templates. You can create service templates and virtual machine templates by using wizards.

You can also modify the settings in an existing template. VMM then incorporates the updated settings into new virtual machines that you create from the updated template. However, the updated template does not affect existing virtual machines that you have created previously using the template. When you modify an existing template, there are new configurable properties in addition to those that are available when you create the template in the New Template Wizard.

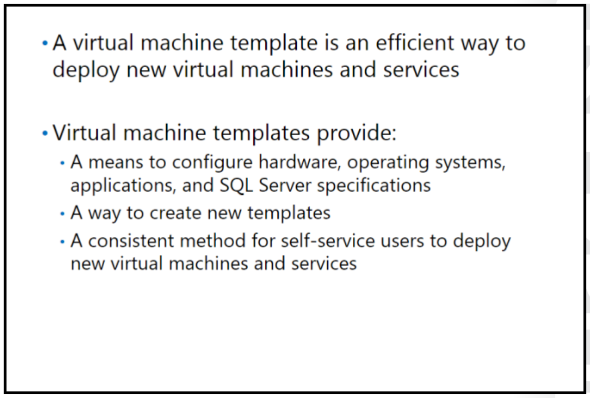

- A virtual machine template is an efficient way to deploy new virtual machines and services
- Virtual machine templates provide:
  - A means to configure hardware, operating systems, applications, and SQL Server specifications
  - A way to create new templates
  - A consistent method for self-service users to deploy new virtual machines and services

To modify the settings of virtual machine template, open the Library workspace, expand the Templates node in the console tree, click VM Templates, and then double-click the template in the details pane.

- On the General page, you can modify following fields:
  - Name (required). Identifies the template to VMM.
  - Description (optional). Describes the characteristics and purpose of a template.

  - Cost Center (optional). Specifies the cost center for a virtual machine that you create from a template. Identifying a cost center enables you to collect data about the allocation of virtual machines (or resources allocated to virtual machines) by cost center.
  - Tag (optional). Specifies a word or phrase that you can use to group specific virtual machines as a set. You can use the tag as a filter to search for a specific set of virtual machines.

- On the Hardware Configuration page, you can modify settings as described in the topic about hardware configuration settings.
- On the Operating System Configuration page, modify the guest operating system settings as described previously in the topic about guest operating system settings.
- On the Application Configuration page, were you to make changes to any of these, you can modify settings to add any of the following: a compatible operating system, SQL Server data-tier applications, virtual applications, web applications, and scripts.
- On the SQL Server Configuration page, you can modify a SQL Server instance, or the configuration of an instance in the service account and the agent that SQL Server uses to run its various services.
- On the Custom Properties page, add or modify the custom fields (optional). You can add up to ten custom properties to each virtual machine that you create by using this template.
- On the Settings page, add or modify quota points (optional). Only self-service users who share a self-service policy use quota points. You can specify a value for the Quota Points setting if the virtual machines that you create by using this template are for self-service users. You can use quota points to limit the number of virtual machines that self-service users can deploy at one time. The quota applies to all virtual machines that you deploy on a host, including virtual machines that are not running.
- On the Dependencies page, you can select and modify dependency properties by clicking on the hyperlinks that make up the dependency’s name, and then modifying the various properties.
- The Validation Errors page displays any validation errors that the template may encounter.
- On the Access page, you can modify the self-service owner, and add or remove self-service user roles that are allowed to use the template.

Service Templates

Deploying a new service requires a high level of automation, predefined components, and management software support. VMM provides service templates for this purpose. 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, a user can use service templates to install and start new applications on demand.

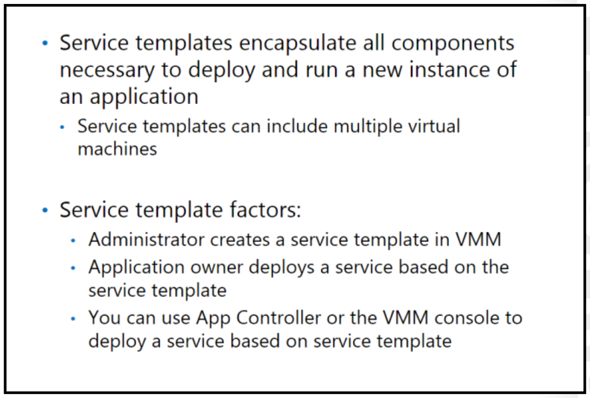

- Service templates encapsulate all components necessary to deploy and run a new instance of an application
  - Service templates can include multiple virtual machines
- Service template factors:
  - Administrator creates a service template in VMM
  - Application owner deploys a service based on the service template
  - You can use App Controller or the VMM console to deploy a service based on service template

Process for Deploying a New Service

To deploy a new service or application by using VMM service templates, use the following procedure:

1. The system administrator creates and configures service templates in VMM by using the Service Template Designer.

2. The application owner (such as a developer who needs to deploy the application environment), opens the App Controller portal, and requests a new service deployment based on available service templates that he or she can access. The user can then deploy the service to a private cloud to which he or she has access. As an alternative to App Controller, the user can also use the VMM console.
3. The VMM management server evaluates the request. VMM searches for available resources in the private cloud, calculates the user quota, and verifies that the cloud can execute the requested service deployment.
4. VMM deploys the virtual machines and applications (if any) on to the host chosen by VMM when it automatically creates the service.
5. The application owner gains control over service virtual machines through the App Controller portal, or by Remote Desktop Protocol (RDP).

If manual approval for resource creation is required, you can use System Center 2012 Service Manager to create workflows.

Information Included in the Service Template

The service template includes information about the virtual machines that are deployed as part of the service. The service template also includes 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 also make use of existing virtual machine templates. Although you can define the service without using any existing virtual machine templates, it is much easier to build a template if you have already created virtual machine templates. After creating the service template, you configure it for deployment by using the Configure Deployment option in the Designer canvas area.

Demonstration: Configuring Virtual Machine and Service Templates

In this demonstration, you will see how to create a:

- Virtual machine template
- Service template

Demonstration Steps

Create a virtual machine template

1. Navigate to the Library workspace, and create a **VM Template** with the following settings:
   - Name: **DemoVMTemplate**
   - Description: **Demonstration creating a VM template**
   - Configure Hardware: Use the **DemoHWProfile**
   - Configure Operating System: Use the **DemoGuestOS** Profile created earlier.
   - Application Profile: None
   - SQL Server Profile: None
2. After you create the template **DemoVMTemplate**, open its properties and review all pages in the console tree.
3. Close all open windows.

Create a service template

1. In the Virtual Machine Manager console, click the **Create Service Template** icon, and create a service template with the following properties:

   a. **General** page:

      - Name: **DemoServiceVM**

   b. **Hardware Configuration** page. In the **Compatibility** section, click **Hyper-V**. In the Select a virtual hard disk pop-up window, click **SmallCore.vhd**. In the **Network Adapter 1 (Legacy)** details pane, click the **Connected to a VM network** option, and then click **External Network**.

   c. **Operating System Configuration** page. In the **Operating system** drop-down list, click **Windows Server 2012 R2 Standard**.

   d. **Application Configuration** page. Click **None – do not install any applications**

   e. **SQL Server Configuration** page. Accept default settings (**None**).

   f. **Custom Properties** page. Accept default settings.

   g. **Settings** page. Accept default settings.

   h. **Dependencies** page. Accept default settings.

   i. **Validation Errors** page. If there were any validation errors they would display here. Accept default settings.

   j. At the bottom of the Single Tier Properties window, click **OK**.

2. Use your mouse to drag the **External Network** box beside the **NIC 1** box.

3. Click **Save and Validate** the service, and then click the **Configure Deployment** icon beside it. Provide the name of the service as **Demo Service**.

4. In the **Deploy Service – Demo Service** window, if a pink shaded area displays in the Deploy Service – Demo Service console saying it could not find a host, click the **Refresh Preview** button.

5. Click **Deploy Service**, and in the **Deploy service** pop-up window, click **Deploy**. The Jobs window displays.

6. On **LON-HOST1**, in Hyper-V Manager, connect to the new virtual machine and **Skip** the product key input page.

7. Close the Jobs window.

8. Close all windows.

Designing Service Templates

In the VMM console, you use the Service Template Designer to create a service template, which defines the configuration of the service.

When you start the Service Template Designer, a few available preconfigured patterns will display. From here, you can create additional templates either by modifying the Blank pattern, or by selecting either the Single Machine pattern, the Two-Tier Application pattern, or the Three-Tier Application pattern. Deploying tiers actually defines levels of your application. For example, one tier of your application can be a web server (or servers), while a second tier could be database servers.

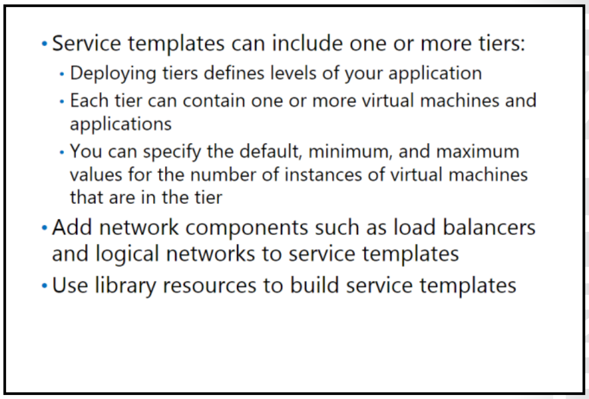

- Service templates can include one or more tiers:
  - Deploying tiers defines levels of your application
  - Each tier can contain one or more virtual machines and applications
  - You can specify the default, minimum, and maximum values for the number of instances of virtual machines that are in the tier
- Add network components such as load balancers and logical networks to service templates
- Use library resources to build service templates

Note: A tier is not equivalent to a virtual machine. A tier—or more specifically a machine tier—contains one or more virtual machines of an identical type.

When you create a tier, you specify the default, minimum, and maximum values for the number of instances of virtual machines that are in the tier. You also can add a virtual IP load balancer to a tier that has virtual machines with services that need load balancing. By creating tiers, you define levels on which your application is working.

The simplest way to add a tier is to use the Service Template Designer. In the Service Template Designer, a list of available virtual machine templates displays in the left pane. You select the virtual machine template that you want to use to create a tier, and then drag the virtual machine template on to the canvas. Service Template Designer then creates the tier using the properties of the selected virtual machine template.

For each tier that you have in your service template, you can configure options such as name, scale-out capabilities, hardware configuration, operating system configuration, and application configuration.

If you created a service template with a pattern that created default tiers, you can drag the virtual machine template on to one of those default tiers. The tier is then configured with the properties of that virtual machine template. You can also add more tiers.

No link or relationship is created between the virtual machine template and the tier that you create. Any subsequent changes that you make to the virtual machine template in VMM are not made to the tier in the service template. Furthermore, any configuration settings that you make to the tier are not made to the virtual machine template. The virtual machine template that you drag to the tier in the Service Template Designer provides you with a configuration template that you can further modify, but establishes no permanent connection between the virtual machine template, tier, or service template.

Updating and Servicing Service Templates

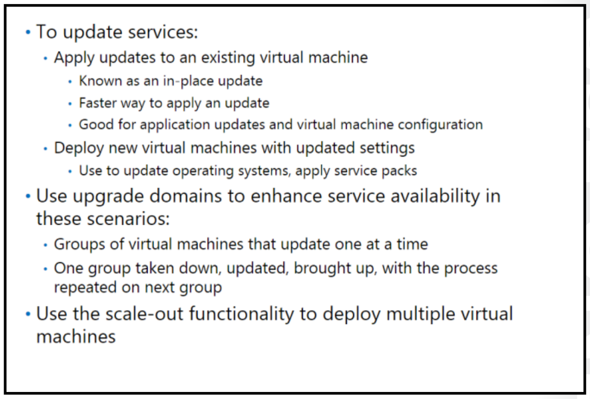

- To update services:
  - Apply updates to an existing virtual machine
    - Known as an in-place update
    - Faster way to apply an update
    - Good for application updates and virtual machine configuration
  - Deploy new virtual machines with updated settings
    - Use to update operating systems, apply service packs
- Use upgrade domains to enhance service availability in these scenarios:
  - Groups of virtual machines that update one at a time
  - One group taken down, updated, brought up, with the process repeated on next group
- Use the scale-out functionality to deploy multiple virtual machines

Business requirements should dictate whether a particular service is current and practical. If requirements change, you can make changes to a deployed service by updating that service. Because you use a service template to deploy a service, you also can update the service template to make changes to the deployed service.

In VMM, you can update a deployed service in one of two ways: by applying updates to the existing virtual machines, or by deploying new virtual machines with the updated settings. Applying updates to the existing virtual machines is the faster of the two options. This type of update is called an *in-place update*. You can apply most application updates and configuration changes to virtual machines by using an in-place update.

You can also create objects called upgrade domains. You can use an upgrade domain to minimize service interruptions when you update a tier in place. Note that upgrade domains have nothing to do with Active Directory domains. When you set the number of upgrade domains, VMM arbitrarily assigns virtual machines to an upgrade domain. When you update a tier in a service, VMM updates the virtual machines in the tier according to the upgrade domain to which they belong. The upgrade domains are updated one at a time, and the virtual machines in that upgrade domain are shut down, updated, and then brought back online. The VMM then repeats the process in the next upgrade domain. This means that updates can take place with the least possible impact to the running service.

Alternatively, you can use VMM to update a deployed service by creating new virtual machines with the updated settings. This takes more time because you are replacing the service's existing virtual machines with new virtual machines. However, this is the preferred way to deploy operating system updates (such as service packs) on the virtual machine. You can use a script to save the state of certain applications before taking down the virtual machines, and then restore the application state to the new virtual machines when you deploy them. You can also use Server App-V, which supports automatic saving and restoring of application states without requiring scripting.

After you deploy a service, you may need to deploy additional virtual machines to a tier in that service. You can use the scale-out functionality of VMM for such scenarios.

There may be times when you need to expand your service quickly to meet rapid growth in demand. For example, a company that sells a particular product only during a holiday season might see a big increase in visits to its website during that time. In this situation, having the capability to scale up more virtual machines to host additional web servers is ideal.

You specify within the properties of a particular tier in a service template whether that tier can be scaled out. You can set the minimum and maximum number of virtual machines that can be deployed in that tier. Note that you will receive a warning if you try to scale out a tier beyond the maximum number of specified virtual machines. However, VMM does not prevent you from scaling out that tier. Instead, the tier and service displays in the VMs and Services workspace a status of needs attention.

Use the following procedure to scale out a tier in a deployed service:

1. Open the VMs and services workspace.
2. Select the private cloud or host group to which you deploy the service.
3. Select the service to be scaled out.
4. On the **Home** tab of the ribbon, click the **Services** icon.

5. On the **Service** tab of the ribbon, click the **Scale Out** icon. This opens up the Scale Out Tier Wizard.
6. The first page of the wizard is the **Select Tier** page. On this page, the **Tier details** section displays the number of virtual machines currently deployed, and the minimum and maximum tier sizes.
7. On the **Select Tier** page, use the Tier drop-down list box to select the tier that you want to scale out, and then click **Next**.
8. Since you will be creating a new virtual machine, on the **Specify Virtual Machine Identity** page, type a name for the new virtual machine, and then click **Next**.
9. If the tier is in a service that is deployed to a private cloud:
   - In the **Configure Settings** page, select the **Identity Information** item in the settings tree, in the **Computer name** text box, type the computer name, and then click **Next**.
10. If the tier is in a service that is deployed to a host group:
    a. On the **Configure Settings** page, select the **Identity Information** item in the settings tree, and then in the **Computer name** text box, type the computer name.
    b. Update any other virtual machine settings as needed, and then click **Next**.
11. On the **Add Properties** page, you can select actions to take when the host server starts or stops, and then click **Next**.
12. On the **Summary** page, click **Scale Out**.
13. The Jobs window displays, and shows the Create virtual machine task. This can take several minutes. When this task completes successfully, go to the VM's and Services workspace, and verify that the new virtual machine was added to the service tier.

Deploying Virtual Machines by Using App Controller

In a private or public cloud solution, end users focus not on virtual machines or servers, but rather on applications and services. Because VMM focuses primarily on virtual machines and service management, you need an additional tool that enables application owners to view services and applications. In previous VMM versions such as System Center Virtual Machine Manager 2008 R2, the Self-Service Portal enabled end users to create and manage virtual machines from their permission scope. However, the Self-Service Portal is orientated to virtual machines, and not to services or applications.

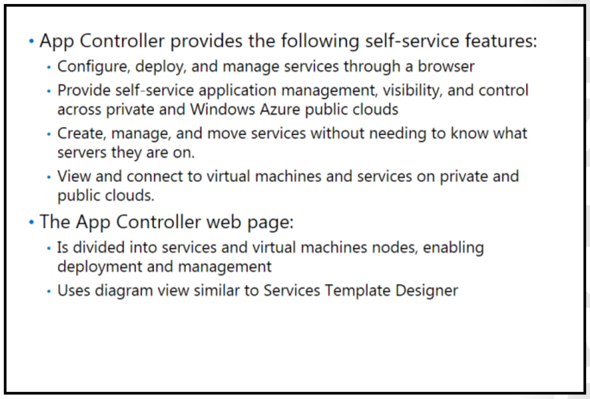

- App Controller provides the following self-service features:
  - Configure, deploy, and manage services through a browser
  - Provide self-service application management, visibility, and control across private and Windows Azure public clouds
  - Create, manage, and move services without needing to know what servers they are on.
  - View and connect to virtual machines and services on private and public clouds.
- The App Controller web page:
  - Is divided into services and virtual machines nodes, enabling deployment and management
  - Uses diagram view similar to Services Template Designer

Benefits of App Controller

App Controller enables users to self-manage application components directly from within their browser. It also provides them with a unified view that enables them to manage applications and services across private clouds and Windows Azure™.

App Controller provides the self-service component of this solution by enabling application owners to:

- Configure, deploy, and manage services through a browser window, using a library of standard templates.

- Provide self-service application management, visibility, and control across both the Microsoft private cloud services, and the Microsoft public cloud services such as Windows Azure.
- Create, manage, and move services without needing to know what servers they are on, or needing to use server-level tools.
- View and connect to virtual machines and services on private and public clouds. Job tracking and history views of jobs and actions taken are also available.

App Controller also enables data center administrators to delegate authority to application owners. Predefined templates ensure compliance with company IT standards and policies. By using App Controller, data center administrators can create for application owners a customized, role-based view of private and public cloud services, and a view of consumed and available resources. In addition, application owners can customize all service components, including virtual machines, network resources, and load balancing.

You can also use App Controller to move applications and components within public and private cloud environments. You can copy Windows Azure configuration, package files, and .vhd/.vhdx files among Windows Azure subscriptions. You can also copy service templates and resources from one VMM server to another.

App Controller has been updated to work with VMM. App Controller can connect to the version of the Service Provider Foundation that shipped with System Center 2012 SP1 and System Center 2012 R2.

You use the Services page in the App Controller web-based console to deploy new services to public and private clouds, and even change the properties of the services that are already deployed. You also can deploy virtual machines to either VMM, Windows Azure, or to another hosting provider. If a virtual machine is part of a VMM service, it deploys when the service deploys.

The App Controller web-based console can also provide management control of the services in virtual machines that are deployed already on VMM private clouds, and of services that are deployed on Windows Azure.

On the services page, you can list the deployed services and display a diagram, much like the Services Template Designer. The diagram enables you to view or change deployed service properties, and view other tasks that you can perform on deployed services.

To deploy a service to a private cloud, go to the Clouds node in the App Controller console tree. You can then right-click a named cloud in the Clouds details pane, and then click Deploy. Alternatively, on the control bar at the top of the Clouds details pane, you can click the Deploy button. App Controller uses the New Deployment diagram view to configure the settings for the service deployment. The New Deployment view has a Deploy button that is not available until all required settings have been supplied.

You can also manage deployed services by selecting the Services node in the App Controller console tree. The All Deployed Services details pane has a list of the various deployed services. A VMM administrator in the VMM console creates service templates. After you create these templates and delegate them to a user role, they display as deployed services in the All Deployed Services details pane. You can then right-click a deployed service name, or select the name and then use the various buttons on the control bar above. The Open Diagram button brings up a diagram view that enables you to change settings for a particular deployed service. You can also start, stop, suspend, and resume a deployed service. There is also a Servicing icon that enables you to upgrade and delete deployed services and resolve any issues found.

If your organization has a subscription to Windows Azure, you can add the items in that subscription to the App Controller web-based console. You can find Windows Azure hyperlinks and icons in in the App Controller Library.

To deploy a service to Windows Azure, you must first create the Windows Azure configuration and package files. You must first upload these files to the Windows Azure storage account. After you complete this step and after you select the particular configuration file, the diagram view loads with the proper information, which enables you to click hyperlinks to configure settings. After this configuration is done, the Deploy button in the diagram view becomes available.

The Virtual Machines node of the App Controller web-based console also has a Deploy button. When you select it, the console displays a New Deployment diagram view that you can use to create a virtual machine. Similar to the deployed services diagram view, you can click hyperlinks in the view to configure the various settings. The Deploy button is available only after you complete all of the required configurations.

You can also use the Virtual Machines node to select and then right-click a listed virtual machine, which opens a context menu. Alternatively, you can highlight the virtual machine in the Virtual Machines node, and then choose an item from the control bar above to manage it. The functionality in the Virtual Machines node enables you to:

- Open a diagram of an existing virtual machine
- View its properties
- Start the virtual machine
- Store it in a virtual machine library
- Mount an .iso image file to it
- Open a Console to the virtual machine
- Delete the virtual machine

Lab: Planning and Deploying Virtual Machines by using Virtual Machine Manager

Scenario

A. Datum Corporation had decided to virtualize workloads across their network infrastructure to optimize their data center, and to facilitate the potential deployment of a private cloud. To enable the centralized management and automated deployment of the required multiple virtual servers, you are required to plan and deploy the VMM components.

You have already deployed VMM, now you must plan for the configuration objects on the server, and perform a virtual machine pilot deployment.

Objectives

After completing this lab, you will be able to:

- Plan Virtual Machine Manager library components.
- Plan VMM virtual machine templates and service templates.
- Deploy virtual machines using VMM templates.

Lab Setup

Estimated Time: 75 minutes

Virtual machines	20413C-LON-HOST1 20413C-LON-DC1-B 20413C-LON-VMM1
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. Restart your host machine and boot into 20413C-LON-HOST1.
2. On LON-HOST1, launch the Microsoft Hyper-V® Manager console.
3. On LON-HOST1, in Hyper-V Manager, In the Virtual Machines details pane, right-click **20413C-LON-DC1-B**, and then click **Settings**.
4. In **Settings for 20413C-LON-DC1-B**, in the Hardware console, click **Network Adapter**.
5. In the Network Adapter details pane, in the **Virtual Switch** drop-down list box, click **External Network**, and then click **OK**.
6. Repeat steps 3-5 for 20413C-LON-VMM1.
7. In the Virtual Machines details pane, right-click **20413C-LON-DC1-B**, click **Connect**, and then click **Start**.
8. Ensure that the 20413C-LON-DC1-B virtual machine is running.

9. Sign in to LON-DC1 using the following credentials:

- User name: **Administrator**
- Password: **Pa$$w0rd**
- Domain: **Adatum**

10. Repeat steps 7 through 9 for 20413C-LON-VMM1.

Exercise 1: Planning Microsoft® System Center 2012 R2 Virtual Machine Manager Components

Scenario

A. Datum plans to virtualize a wide variety of servers using Hyper-V and VMM. The team that is planning the implementation of virtual machines has identified the following server roles for virtualization:

- Domain controllers
- File servers
- Web servers
- Database servers
- Email servers
- Application servers

Supporting Documentation

----- Original Message -----
From: Charlotte Weiss [Charlotte@adatum.com]
Sent: 03 Feb 08:45
To: Brad@adatum.com
Subject: Adatum Virtual Machine Management (VMM) Project Plan

Brad,

See the list of Servers we intend to virtualize for the VMM Project. I need you to consider what virtual machine options we should include for each server role. Also, we need you to identify two to three hardware profiles and at least two guest operating system profiles for these servers.

Any other information gratefully received!

Charlotte

Charlotte Weiss

From: Brad Sutton [Brad@adatum.com]
Sent: 04 Feb 9:05
To: Charlotte@adatum.com
Subject: Re: Adatum Virtual Machine Management (VMM) Project Plan
Attachments: Adatum Server Virtualization Proposal.docx

Charlotte,

We have eight domain controllers: four in London, two in Toronto, and two in Sydney. All of them are running Windows Server 2012 R2, which means we can virtualize all our domain controllers without having to worry about USN rollback errors. We'll need to create a separate hardware profile for the London domain controllers as they need more performance abilities. Toronto and Sydney's requirements are basically the same. My team will draw up similar requirements for the other servers.

Regards,

Brad

The list of A. Datum servers are shown in the following table.

Server type	Description
Domain controllers	Eight domain controllers, all running Windows Server 2012 R2, four in London, two each in Toronto and Sydney
File servers	Five file servers, three running Windows Server 2008 R2, and two running Windows Server 2012 R2
Web servers	Three web servers, all running Windows Server 2012 R2
Database servers	Two servers running Windows Server 2012 R2 and SQL Server 2012
Email servers	Two servers running Windows Server 2012 R2 and Microsoft Exchange Server 2013
Application servers	Two servers running Windows Server 2008 R2 and Microsoft SharePoint® Server 2010

Proposals

The following is from the proposal document for the Adatum Virtualization Test Project. The profile requirements for the domain controllers are supplied. Complete the profile requirements for all of the other servers.

Domain controllers

Profile requirements: Eight domain controllers all running Windows Server 2012 R2, four in London, and two each in Toronto and Sydney. The hardware/performance requirements are shown in the following table.

Domain controllers	Hardware/performance requirements
London	4 Processor cores, 8 GB of RAM and a separate disk drive to house the NT Directory Service database and the Sysvol directory
Toronto	2 Processor cores, 4 GB of RAM
Sydney	2 Processor cores, 4 GB of RAM

Question: What profiles do you recommend to meet these requirements?

Answer:

File servers

Profile requirements: Five file servers, three running Windows Server 2008 R2, and two running Windows Server 2012 R2

Hardware/performance requirements: two processor cores, 4 GB of RAM, two 2-TB drives for data

Question: What profiles do you recommend to meet these file server requirements?

Answer:

Web servers

Profile requirements: Three web servers, all running Windows Server 2012 R2

Hardware/performance requirements: Two processor cores, 4 GB of RAM, two 1-TB drives for data

Question: What profiles do you recommend to meet these web server requirements?

Answer:

Database servers

Profile requirements: Two SQL servers running Windows Server 2012 R2 and Microsoft® SQL Server® 2012

Hardware/performance requirements: Four processor cores, 8 GB of RAM, two 2-TB drives for data

Question: What profiles do you recommend to meet these database server requirements?

Answer:

Email servers

Profile requirements: Two Exchange servers running Windows Server 2012 R2 and Exchange Server 2013

Hardware/performance requirements: Four processor cores, 8 GB of RAM, two 2-TB drives for data

Question: What profiles do you recommend to meet these email server requirements?

Answer:

Application servers

Profile requirements: Two servers running Windows Server 2008 R2 and SharePoint Server 2010

Hardware/performance requirements: Two processor cores, 4 GB of RAM (data stored on SQL Server databases)

Question: What profiles do you recommend to meet these application server requirements?

Answer:

The main tasks for this exercise are as follows:

1. Read the supporting documentation

2. Update the proposal document with your planned course of action

3. Examine the suggested proposals in the Lab Answer Key

4. Discuss your proposed solution with the class, as guided by your instructor

▶ Task 1: Read the supporting documentation

Read the documentation provided.

▶ Task 2: Update the proposal document with your planned course of action

Answer the questions in the proposal section of the Adatum Virtualization Test Project.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

Compare your proposals with the ones in the Lab Answer Key.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

1. Go over your answers for the proposal with the class.
2. Explain any choices you made that differ from the Lab Answer Key.

Results: At the end of this exercise, you should have created hardware profiles, guest operating system profiles, and SQL Server profiles to meet requirements.

Exercise 2: Planning Virtual Machine and Service Templates

Scenario

Based on the virtualization requirements, you will now plan the virtual machine templates and service templates that are required to meet your deployment goals. The virtual machine templates will build on the profile objects that you planned for in the previous exercise.

Supporting Documentation

----- Original Message -----
From: Charlotte Weiss [Charlotte@adatum.com]
Sent: 15 Feb 08:45
To: Brad@adatum.com
Subject: Adatum VMM Project Deployment Plan

Brad,

Thanks for that proposal document. We need to go ahead and set up templates for the deployments. I've talked to the Development group, and they would like to be able to quickly roll out small database and web servers. Do you think you could create a Service Template to do so?

Charlotte

Charlotte Weiss

From: Brad Sutton [Brad@adatum.com]
Sent: 15 Feb 9:05
To: Charlotte@adatum.com
Subject: Re: Adatum VMM Project Deployment Plan

Attachments: Adatum Server Service Template Requirement.docx

Charlotte,

We can do it! We can easily create virtual machine templates from the various profiles we just wrote. Also, we can create a Service Template to deploy the small database and web servers the Development group wants. I'll have my team get right on it!

Regards,

Brad

Adatum Server Service Template Requirement.docx

Developer Group Virtual Machine Requirement: Service Template

Document Reference Number: GW00777/1

Document Author: Brad Sutton
Date: 24th February

Requirements Overview

The developer group would like to have the following virtual machines created on an as-needed basis, and each will be used no longer than one week. When no longer needed, the virtual machines can be deleted and their resources returned to the host computers.

- Windows Server 2012 R2: 2-GB RAM, single processor, 60-GB partition, Role: Microsoft Internet Information Services (IIS) 8.0
- Windows Server 2012 R2: 2-GB RAM, single processor, 100-GB partition, install SQL Server 2012 Developer edition

Additional Information

Need profiles and service template that can create the virtual machines as needed.

Proposals

Specify the types of profiles and templates needed and write them into the proposals section of the Adatum Server Service Template Requirement document.

The main tasks for this exercise are as follows:

1. Read the supporting documentation
2. Update the proposal document with your planned course of action
3. Examine the suggested proposals in the Lab Answer Key
4. Discuss your proposed solution with the class, as guided by your instructor

▶ Task 1: Read the supporting documentation

Read the documentation provided.

▶ Task 2: Update the proposal document with your planned course of action

Specify the types of profiles and templates needed and write them into the proposals section of the Adatum Server Service Template Requirement document.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

Compare your proposals with the ones in the Lab Answer Key.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

1. Go over your answers for the proposal with the class.
2. Explain any choices you made that differ from the lab answers.

Results: At the end of this exercise, you should have planned virtual machine and service templates.

Exercise 3: Implementing Virtual Machine Manager Components

Scenario

Now that you have created the VMM component design, the next step is to configure the components and validate the configuration by deploying virtual machines using a service template.

The main tasks for this exercise are as follows:

1. Adding LON-HOST1 as a Host Server to VMM
2. Create File Server guest operating system and hardware profiles
3. Create developers' group web server guest operating system profile and hardware profile
4. Create developer's group database guest operating system profile and hardware profile
5. Create a developers' group database SQL Server profile
6. Configure a virtual machine template
7. Create a virtual machine template from the developers' group web server profiles
8. Create the developers' group web service template
9. Deploy virtual machines using the service template
10. To prepare for the next module

▶ Task 1: Adding LON-HOST1 as a Host Server to VMM

Set the default domain Group Policy to allow domain members to become hosts

1. On LON-DC1, in Server Manager, open the Group Policy Management Editor, and then edit the Default Domain Policy.
2. Navigate to **Computer Configuration**, **Administrative Templates\Network\Network Connections\Windows Firewall\Domain Profile**, and then apply the following settings:
   a. In Windows Firewall:
      - Allow inbound file and printer sharing exception: **Enabled**
      - Options: type an asterisk (*****) (which indicates all IP addresses).

   b. In Windows Firewall:
      - Allow ICMP exceptions: **Enabled**
      - Options: **Allow inbound echo request**.

   c. In Windows Firewall:
      - Define inbound port exception: Enabled
      - Options: Define port exceptions, click **Show**, and in the **Value** text box, type **5985**.
3. In the Group Policy Management Editor, navigate to **Administrative Templates\Windows Components\Windows Remote Management (WinRM)\WinRM Service**.
4. In the **Allow remote server management through WinRM** dialog box, select **Enabled**, in **Options**, for both **IPv4** and **IPv6**, type an asterisk (*****).
5. Close the Group Policy Management Editor.

6. On the LON-HOST1 physical machine, update the group policy using the following Windows PowerShell cmdlet:

```
gpupdate.exe /force
```

Add LON-HOST1 to VMM

1. On LON-VMM1, open the VMM console, and add **LON-HOST1** as a Hyper-V server to the **All Hosts** node in **VMs and Services**, using the following settings:

   a. **Resource Location** page: **Windows Server computers in a trusted Active Directory**.

   b. **Credentials** page:

      - User name: ADATUM\Administrator
      - Password: Pa$$w0rd

   c. **Discovery Scope** page: Computer names: **lon-host1.adatum.com**

   d. **Target resources** page: Discovered computers: **lon-host1.adatum.com**

   e. **Host Settings** page: **All Hosts**

   f. **Summary** page: **Finish**

2. Observe that LON-HOST1 now displays in the VM's and Services console tree.

▶ Task 2: Create File Server guest operating system and hardware profiles

1. Open the VMM console (if not already open), and in the Library workspace, create a File Server Guest profile, with the following parameters:

- Name: **FileServerGuestOS**
- Description: **Adatum File Server Guest OS profile**
- Operating System: **edition of Windows Server 2012 R2 Standard**
- Identity Information: **LON-FS1##**
- Admin Password: **Pa$$w0rd**
- Roles: **File and Storage Services**

2. Create a File Server Hardware profile with the following parameters:

- Name: **FileServerHWProfile**
- Description: **Create File Server hardware profile**
- Compatibility: **Hyper-V**
- Memory: **Dynamic, Maximum memory 2048**
- Network Adapter: **Network Adapter 1, Connected to a VM network, External Network**

▶ Task 3: Create developers' group web server guest operating system profile and hardware profile

1. Create a Developers' Group Web Guest OS profile with the following parameters:

- Name: **DevGroupWebGuestOS**
- Description: **Developers' Group Web Server Guest OS profile**
- Operating System: **Windows Server 2012 R2 Standard**

- Identity Information: **DevWeb##**
- Admin Password: **Pa$$w0rd**
- Roles: **Web Server (IIS)**

2. Create a Developers' Group Hardware profile with the following parameters:

- Name: **DevGroupWebHWProfile**
- Description: **Create Developers' Group Web Server hardware profile**
- Compatibility: **Hyper-V**
- Memory: **Dynamic, Maximum memory 2048**
- Network Adapter: **Network Adapter 1, Connected to a VM network, External Network**

▶ Task 4: Create developer's group database guest operating system profile and hardware profile

1. Create a Developers' Group Database guest operating system profile with the following parameters:

- Name: **DevGroupDBGuestOS**
- Description: **Developers' Group Database Server Guest OS profile**
- Operating System: **Windows Server 2012 R2 Standard**
- Identity Information: **DevDB##**
- Admin Password: **Pa$$w0rd**

2. Create a Developers' Group Database Hardware profile with the following parameters:

- Name: **DevGroupDBHWProfile**
- Description: **Create Developers' Group Database Server hardware profile**
- Compatibility: **Hyper-V**
- Memory: **Dynamic, Maximum memory 2048**
- Network Adapter: **Network Adapter 1, Connected to a VM network, External Network**

▶ Task 5: Create a developers' group database SQL Server profile

- Create a Developers' Group Database SQL Server profile with the following parameters:
  - Name: **DevGroupDBSQL**
  - Description: **Developers' Group Database Server SQL Server profile**
    - SQL Server Deployment Name: **DevDeploy**
    - Instance Name: **MSSQLSERVER**
    - Instance ID: **1**
- SQL Server administrators: **adatum\DevDBO**
- Media Source: **D**
- System Administrator (SA) Password Run As Account:
  - Create Run As Account Name: **DevDBO**
  - Description: **Developer's Group Database Owner**
  - Username: **Adatum\DevDBO**
  - Password and Confirm Password: **Pa$$w0rd**

- **Clear the Validate domain credentials**
- Service Accounts (all): **NT AUTHORITY\System**

▶ Task 6: Configure a virtual machine template

1. Navigate to the **Library** workspace and create a **VM Template** with the following settings:

- Name: **FSVMTemplate**
- Description: **Create the File Server VM template**
- Select VM Template Source: **SmallCore.vhd**
- Configure Hardware: Use the **FileServerHWProfile**
- Configure Operating System: Use the **FileServerGuestOS** Profile created earlier.
- Application Profile: **None – do not install any applications**
- SQL Server Profile: **None – no SQL Server configuration settings**

2. After you create the **FSVMTemplate**, open its properties and review all pages in the console tree.

▶ Task 7: Create a virtual machine template from the developers' group web server profiles

1. Navigate to the Library workspace, and create a **VM Template** with the following settings:

- Select VM Template Source: **SmallCore.vhd**
- Name: **DevGrpWebVMTemplate**
- Description: **Create the Developers' Group Web Server VM template**
- Configure Hardware: Use the **DevGroupWebHWProfile**
- Configure Operating System: Use the **DevGroupWebGuestOS** Profile that you created earlier.
- Application Profile: **None – do not install any applications**
- SQL Server Profile: **None – no SQL Server configuration settings**

2. After you create the template for DevGrpWebVMTemplate, open its properties and review all pages in the console tree.

▶ Task 8: Create the developers' group web service template

1. In the Virtual Machine Manager console, click the **Create Service Template** icon, and create a service template with the following properties:

- Name: **Dev Group Web Service Template**
- Release: **1**

2. Choose the **Blank** pattern.
3. Note the text that says, **Drag VM Templates onto the canvas to create a new Tier and copy the VM Template settings into that tier**. Drag the **DevGroupWeb** virtual machine template onto the Designer canvas.
4. Use your mouse to drag the **External Network** box next to the **NIC 1** box.

▶ Task 9: Deploy virtual machines using the service template

1. On the **Home** tab, click **Save and Validate** the service, and then next to it click the **Configure Deployment** icon. Provide the name of the service as follows:

- Name: **DevGroup Web Service**.

2. When the **Deploy Service – DevGroup Web Service** console opens, if you see a message in a pink shaded area in the **Deploy Service – DevGroup Web Service** console saying it could not find a host, click the **Refresh Preview** button.
3. Click **Deploy Service**.
4. In the Deploy service pop-up window, click **Deploy**.
5. Verify that the Jobs window displays.
6. Close the Jobs window.
7. Close all open windows.

▶ Task 10: 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 LON-HOST1, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-DC1-B**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for 20413C-LON-VMM1.
5. Restart the host computer, and at the boot menu, select the Windows Server 2012 installation.

Results: At the end of this exercise, you should have implemented Virtual Machine Manager components.

Question: After you created the virtual machine template, when you reviewed its properties, where did the values in the Hardware and Operating Systems tabs come from?

Question: Why did you decide to use a service template to deploy the virtual machines required by the Developers' Group?

Module Review and Takeaways

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
You cannot add a library server.	

Review Question

Question: How can you deploy a virtual machine template in VMM?

Tools

Tool	Use for	Where to find it
VMM console	Creating and deploying virtual machines, hardware profiles, guest profiles, and SQL Server profiles, virtual machine templates, Service Templates, and other library objects.	Installed on the VMM management server
App Controller	Web-based management of virtual machines, Hyper-V hosts, VMM services, private clouds, and other VMM objects.	System Center 2012 VMM installation media
Group Policy Management Editor	Establishing settings for large numbers of computer and users settings. Can also use to set Windows Remote Management settings when deploying the Virtual Machine Manager agent to Hyper-V hosts.	Domain controllers, or add appropriate Remote Server Administrator Tools to the computer

Module 4

Designing and Maintaining an IP Configuration and Address Management Solution

Contents:

Module Overview

When designing network services, you must first design an appropriate IP addressing scheme, and select and configure a method for allocating your chosen scheme. You must also consider how best to manage the introduction of new computer devices, and how to allocate an IP configuration to these devices. You must complete these steps before you can design higher-level network services, such as Domain Name System (DNS), Internet Information Services (IIS), and applications, including email and databases.

Objectives

After completing this module, you will be able to:

- Design Dynamic Host Configuration Protocol (DHCP) servers.
- Plan DHCP scopes.
- Design an IP Address Management (IPAM) provisioning strategy.
- Manage servers and address spaces by using IPAM.

Lesson 1
Designing DHCP Servers

DHCP is an essential component of most IPv4 and IPv6 networks. A majority of network clients rely on DHCP to obtain an IPv6 address for network connectivity. Therefore, when designing a DHCP infrastructure, you must ensure that it is reliable, secure, and highly available.

Lesson Objectives

After completing this lesson, you will be able to:

- Determine the appropriate DHCP server placement.
- Describe the process for configuring DHCP failover.
- Implement DHCP failover.
- Manage IPv6 hosts with DHCP.
- Design your DHCP infrastructure.
- Migrate the DHCP Server role.

DHCP Server Placement

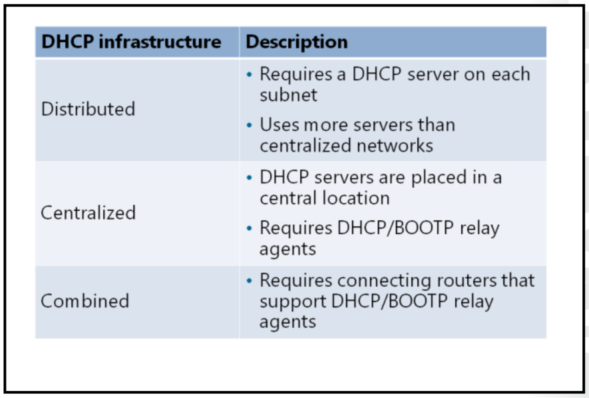

DHCP infrastructure	Description
Distributed	• Requires a DHCP server on each subnet • Uses more servers than centralized networks
Centralized	• DHCP servers are placed in a central location • Requires DHCP/BOOTP relay agents
Combined	• Requires connecting routers that support DHCP/BOOTP relay agents

To ensure correct functioning of your network, you must deploy DHCP servers appropriately within the network infrastructure.

There are three possible strategies for DHCP server placement:

- Distributed. When you select a distributed DHCP infrastructure, you are assuming that no DHCP/Bootstrap Protocol (BOOTP) relays are in place. Consequently, each subnet requires a DHCP server. This is not practical for larger organizations, because it would require a large number of DHCP servers, which would be difficult to manage.
- Centralized. A centralized DHCP infrastructure places DHCP servers in a central location and uses DHCP/BOOTP relays to forward requests from remote subnets. This works well within a single location, but may not be suitable for large organizations with many wide area network (WAN) links. The failure of a WAN link may prevent some clients from obtaining an IP address and accessing network resources.
- Combined. Many organizations use a combined infrastructure. Each WAN location has its own centralized infrastructure with one or more DHCP servers. This allows for centralized management of the DHCP servers, and avoids concerns about WAN link reliability.

What Is DHCP Failover?

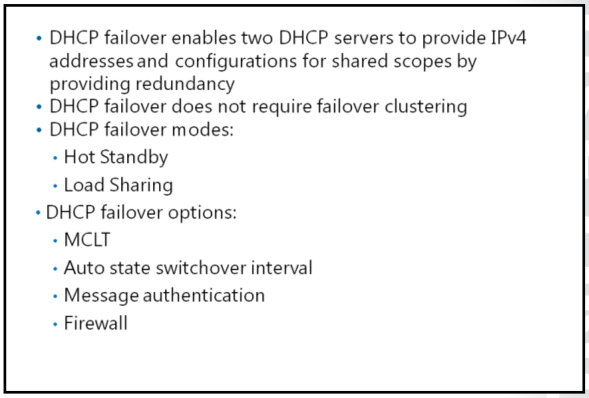

- DHCP failover enables two DHCP servers to provide IPv4 addresses and configurations for shared scopes by providing redundancy
- DHCP failover does not require failover clustering
- DHCP failover modes:
  - Hot Standby
  - Load Sharing
- DHCP failover options:
  - MCLT
  - Auto state switchover interval
  - Message authentication
  - Firewall

In addition to determining DHCP server placement, you also need to consider how to service DHCP clients in case of a server failure. Prior to Windows Server® 2012, techniques that you could use to ensure a DHCP server was available to manage client requests included split scope, and failover clustering. With Windows Server 2012, Microsoft introduced DHCP failover as a highly availability solution for DHCP servers. This solution does not require the extra hardware required for failover clustering, and does not implement split scopes to allow multiple servers to serve the same IP subnets.

DHCP Failover

DHCP clients renew their leases on their IP addresses at regular, configurable intervals. If the DHCP service fails, the leases time out and clients no longer have IP addresses. In the past, failover for DHCP 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, by providing redundant DHCP services, you had to 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. In addition, you can only use failover for IPv4 (and not IPv6) scopes and subnets.

To configure DHCP failover, you need to establish a failover relationship between the two DHCP server 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 so long as they all have unique names. To configure failover, use the Configuration Failover Wizard. You launch this wizard by right-clicking either the IP node or the scope node, and then clicking Configure Failover.

You can configure DHCP failover in one of the following two modes:

- 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 takes over this role if the primary server becomes unavailable.

  A DHCP server can simultaneously act as the primary for one scope or subnet, and 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 one hour. The secondary server assumes control of the entire IP range after the MCLT interval has passed.

Hot Standby mode is best suited for deployments in which a disaster recovery site is physically 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 simultaneously supply IP configuration to clients. The server that responds to IP configuration requests depends on how the administrator configures the load distribution ratio. The default ratio is 50:50.

Additionally, you can configure the following DHCP failover options:

- MCLT
- Auto state switchover interval
- Message authentication
- Firewall

MCLT

The administrator configures the MCLT parameter to determine the amount of time a DHCP server should wait when a partner becomes 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 the communication interrupted state until the administrator manually changes the state to a partner down state. The administrator also can 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 originates from the failover partner.

Firewall Considerations

DHCP uses Transmission Control Protocol (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: Implementing DHCP Failover

This demonstration shows you how to configure DHCP failover.

Demonstration Steps

1. On LON-SVR1, install the DHCP Server role from Server Manager. Run the DHCP post configuration steps.
2. Switch to LON-DC1.
3. In Server Manager, click **Tools**, and then in the drop-down list box, click **DHCP**.
4. In the DHCP console, launch the Configuration Failover Wizard.

5. Configure failover replication with the following settings:

    a. Partner server: **172.16.0.11**

    b. Relationship Name: **Adatum Failover**

    c. Maximum Client Lead Time: **10 minutes**

    d. Mode: **Load balance**

    e. Load Balance Percentage: **50**

    f. State Switchover Interval: **60 minutes**

    g. Message authentication shared secret: **Pa$$w0rd**

    h. Complete the Configuration Failover Wizard.

    i. Switch back to LON-SVR1, and note that the IPv4 node is active and that the Adatum scope is configured.

Managing IPv6 with DHCP

Throughout the autoconfiguration process, the IPv6 host can proceed through several states. There are several ways that you can assign an IPv6 address and other configuration settings during these states. Based on how you set up the router, a client might use stateless configuration (no DHCPv6 service), or stateful configuration with a DHCPv6 server involved. The client uses one of these configurations either to assign an IP address and other configuration settings, or just assign other configuration settings. The other configuration settings can include DNS servers and domain names.

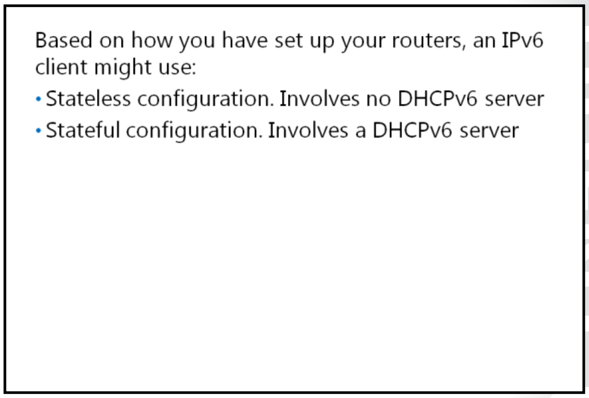

Based on how you have set up your routers, an IPv6 client might use:

- Stateless configuration. Involves no DHCPv6 server
- Stateful configuration. Involves a DHCPv6 server

Types of Autoconfiguration

Types of autoconfiguration include:

- Stateless. With stateless autoconfiguration, address configuration is based only on the receipt of Router Advertisement messages.
- Stateful. Stateful configuration is based on the use of a stateful address configuration protocol (such as DHCPv6) to obtain addresses and other configuration options:
    - A host uses stateful address configuration when it receives instructions to do so in Router Advertisement messages.
    - A host also will use a stateful address configuration protocol when there are no routers present on the local link.
- Both. Configuration is based on the receipt of Router Advertisement messages and DHCPv6.

Autoconfigured Address States

Autoconfigured addresses are in one or more of the following states:

- Tentative. The address is in the process of being verified as unique. Duplicate address detection performs verification. A node cannot receive unicast traffic on a tentative address.

- Valid. The address has been verified as unique, and the computer can send and receive unicast traffic on that address.
- Preferred. The address enables a node to send and receive unicast traffic to and from it.
- Deprecated. The address is valid, but its use is discouraged for new communication.
- Invalid. The address no longer allows a node to send or receive unicast traffic.

Why Use Stateful Configuration?

Organizations use the stateful configuration to control how IPv6 addresses are assigned by using DHCPv6.

If there are any specific scope options that you need to configure—such as the IPv6 addresses of DNS servers—then a DHCPv6 server is necessary.

Communication with a DHCP Server

When IPv6 attempts to communicate with a DHCP server, it will use multicast IPv6 addresses to communicate with the DHCP server. This is different from IPv4, which uses broadcast IPv4 addresses.

Guidelines for Designing a DHCP Solution

Consider the following guidelines when designing a DHCP infrastructure:

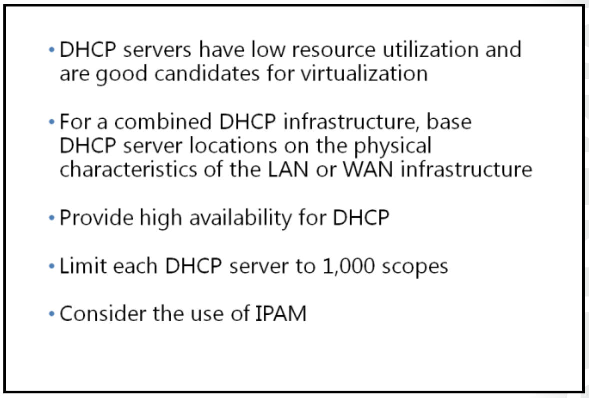

- DHCP servers have low resource utilization and are good candidates for virtualization
- For a combined DHCP infrastructure, base DHCP server locations on the physical characteristics of the LAN or WAN infrastructure
- Provide high availability for DHCP
- Limit each DHCP server to 1,000 scopes
- Consider the use of IPAM

- Virtualize DHCP servers as part of a server consolidation effort. Because they have low resource utilization, DHCP servers are good candidates for virtualization. You can co-host multiple services on the same virtual server, such as DHCP, DNS, and Active Directory® Domain Services (AD DS).
- Plan a combined DHCP infrastructure based on network characteristics, such as WAN links and their reliability. If WAN links are not reliable, ensure there is a DHCP server available in each physical site to avoid DHCP traffic from passing through the WAN link.
- Make DHCP a highly available service in your organization to ensure that clients are able to access network resources. You can achieve DHCP high availability in the following ways:
    - Windows failover clustering. You can configure the DHCP service as a failover resource. However, a failover cluster requires expensive hardware and it is not worth the investment just to make DHCP highly available. In situations where a failover cluster already exists for other services co-hosted on a set of servers alongside DHCP, you can use failover clustering to maintain a highly available DHCP solution.
    - Split scope. You also can provide high availability by having two or more DHCP servers which each contains a portion of a larger scope. However, managing such an environment becomes more complex because you must ensure IP addresses are not reused in multiple scopes on different servers.
    - DHCP failover. Use the new DHCP failover feature in Windows Server 2012 and Windows Server 2012 R2 to maintain a highly available DHCP solution without incurring extra hardware cost.

- Limit each DHCP server to up to 1,000 scopes in centralized and combined DHCP infrastructures. Although you can have more than 1,000 scopes per DHCP server, Microsoft only supports up to 1,000 scopes.
- Consider the use of IPAM to manage DHCP. IPAM helps manage scopes on multiple servers and ensures that IP conflicts are not occurring. (You will learn more about IPAM later in this module.)

Migrating the DHCP Server Role

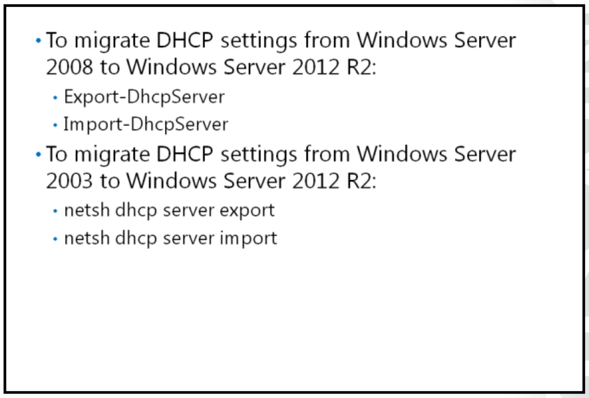

- To migrate DHCP settings from Windows Server 2008 to Windows Server 2012 R2:
  - Export-DhcpServer
  - Import-DhcpServer
- To migrate DHCP settings from Windows Server 2003 to Windows Server 2012 R2:
  - netsh dhcp server export
  - netsh dhcp server import

In certain scenarios, you might already have a DHCP server running Windows Server 2008 or Windows Server 2008 R2 in your environment. If that is the case, you can export the DHCP settings from the existing server, and then import them into a new server running Windows Server 2012 R2 or Windows Server 2012. To migrate a DHCP server to a Windows Server 2012 R2 or Windows Server 2012 server, perform the following steps:

1. Install the DHCP Server role on a computer running Windows Server 2012 R2 or Windows Server 2012.
2. Run the following Windows PowerShell® cmdlet on the new server:

```
Export-DhcpServer –ComputerName oldServer –Leases –File filePath -verbose
```

Note: You need to ensure that the remote server has remote management enabled for the Windows PowerShell cmdlet to run as expected.

3. Run the following command on the new server:

```
Import-DhcpServer –ComputerName newServer –Leases –File filePath -Verbose
```

Although migrating a DHCP server from Windows Server 2008 or newer to Windows Server 2012 R2 requires you to run only two Windows PowerShell cmdlets, you might need additional cmdlets if you have DHCP servers running on Windows Server 2003. In this case, perform the following steps to migrate the DHCP Server role from Windows Server 2003 to a computer running Windows Server 2012 R2:

1. On the Windows Server 2003 server, open an elevated command prompt, and run the following command:

```
netsh
```

2. At the netsh prompt, run the following command:

```
DHCP
```

3. At the netsh dhcp prompt, run the following command:

   ```
   server \\servername_or_IP_address
   ```

4. At the netsh dhcp server prompt, run the following command:

   ```
   export filePath all
   ```

5. The file created in step 4 contains the DHCP settings for the old server. Copy that file to a location that is accessible to the new DHCP server.
6. Install the DHCP Server role on a computer running Windows Server 2012 R2.
7. On the Windows Server 2012 R2 computer, open an elevated command prompt, and run the following command:

   ```
   netsh
   ```

8. At the netsh prompt, run the following command:

   ```
   DHCP
   ```

9. At the netsh dhcp prompt, run the following command:

   ```
   server \\servername_or_IP_address
   ```

10. At the netsh dhcp server prompt, run the following command:

    ```
    import filePath
    ```

11. Restart the new server.

Lesson 2
Planning DHCP Scopes

When you configure a DHCP scope, you have a number of configuration options from which to choose. You should plan on how to determine lease lengths, implement superscopes, use reservations, and implement DHCP class-level options. These considerations will simplify the management of your DHCP infrastructure.

Lesson Objectives

After completing this lesson, you will be able to:

- Determine the appropriate DHCP lease duration.
- Implement superscopes, where appropriate.
- Use DHCP reservations to support specific clients.
- Implement the DHCP option classes.

Determining DHCP Lease Length

DHCP allocates IP configurations to requesting computers on a dynamic basis; this is known as a *DHCP lease*. When the DHCP lease has reached 50 percent of the lease duration, the client computer attempts to renew the lease. The default lease duration for wired clients is eight days, but you can configure this differently.

Note: Client computers also attempt lease renewal during the startup process.

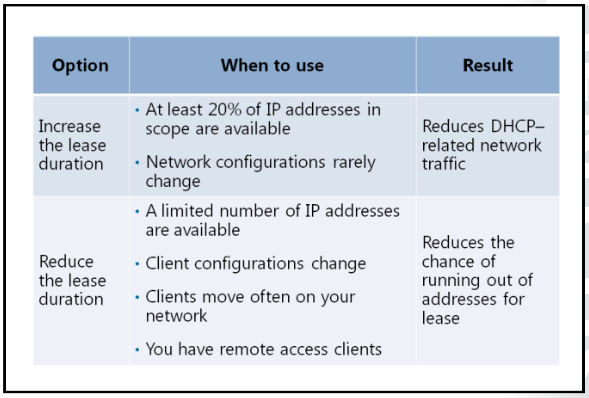

Option	When to use	Result
Increase the lease duration	• At least 20% of IP addresses in scope are available • Network configurations rarely change	Reduces DHCP–related network traffic
Reduce the lease duration	• A limited number of IP addresses are available • Client configurations change • Clients move often on your network • You have remote access clients	Reduces the chance of running out of addresses for lease

When determining the lease length, consider the following:

- Network traffic. Short lease lengths generate additional network traffic because clients renew their leases more often. In most networks, the traffic generated by DHCP requests is minimal and unlikely to affect performance. However, this may be a concern when using a centralized infrastructure over WAN links.
- Address reuse. When roaming clients (such as laptops) obtain an IP address, that address becomes unavailable until the roaming client lease expires, even if the roaming clients are no longer actively using the addresses. In a situation with a high number of roaming clients, this can result in rapid depletion of a large address space.
- Change. If clients have long leases, it may be more difficult to change IP address configuration and DHCP options. For example, a computer with a 60-day lease may keep that information for 60 days without updating. However, new addressing information typically is obtained at half the lease length when the client attempts a renewal.

Implementing Superscopes

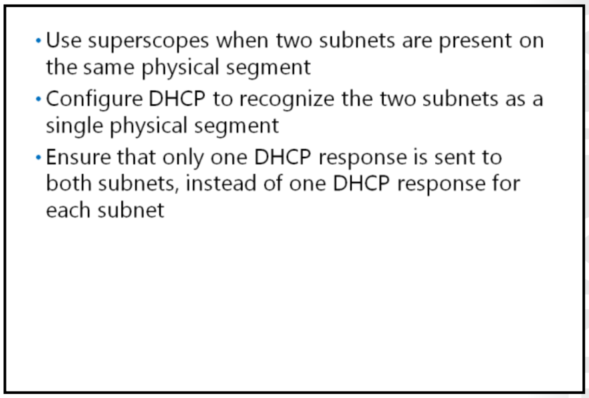

- Use superscopes when two subnets are present on the same physical segment
- Configure DHCP to recognize the two subnets as a single physical segment
- Ensure that only one DHCP response is sent to both subnets, instead of one DHCP response for each subnet

Superscopes are relevant only when a single physical network segment has multiple subnets. For example, consider a physical location that has added additional clients and as a result has expanded beyond the number of addresses that are available in a single subnet. If IP addressing has not been assigned appropriately, it may not be possible to supernet the two addresses, and they must then be managed as two separate networks.

In this scenario, you can use superscopes to combine scopes on a single physical segment. You configure a DHCP server with a scope for each subnet, and then combine the scopes into a single superscope. This configures the DHCP server to recognize that both scopes are on a single physical segment.

Without a superscope, the DHCP server would send a lease offer for each scope. With a superscope in place, the DHCP server only sends a single lease offer.

Using Reservations

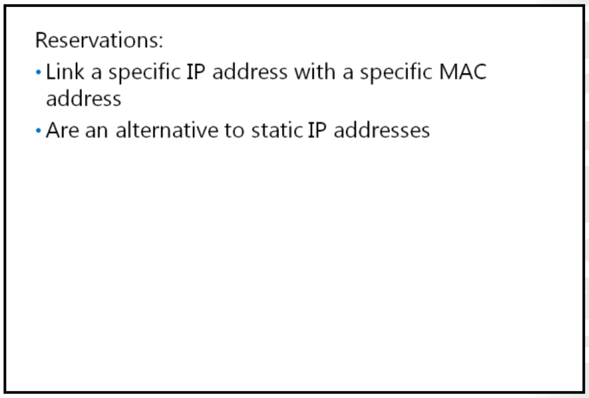

Reservations:

- Link a specific IP address with a specific MAC address
- Are an alternative to static IP addresses

You use DHCP reservations to assign the same IP address to a device each time it leases an address from DHCP. In the reservation, the media access control (MAC) address of the device identifies the device. DHCP reservations enable you to assign a specific IP configuration to a particular host, without having to configure that host manually.

Reservations are an alternative to using static IP addresses for devices. Static IP addresses might be simpler to configure because you need only to configure the individual computer manually. However, static IP addresses do not integrate with the DHCP database, and this makes it more difficult for an administrator to keep track of which IP addresses have been manually assigned to each computer.

Reservations are more complex to create, because they require you to obtain the MAC address of the devices for which you will be making the reservation. However, once you configure the reservation, its information is part of the DHCP database, which acts as a central location for all IP addressing data.

Note: A DHCP reservation does not guarantee that a device will always receive the IP address configured in the reservation. If the device receives a DHCP offer from a different DHCP server on which the reservation was not created, it may use the IP settings offered by the other DHCP server. Therefore, ensure that all DHCP servers have reservations for devices that require static IPs in the subnets for which the DHCP server is active.

Implementing DHCP Options Classes

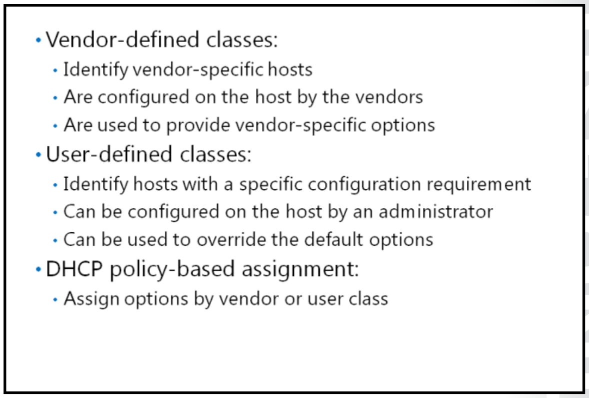

- Vendor-defined classes:
  - Identify vendor-specific hosts
  - Are configured on the host by the vendors
  - Are used to provide vendor-specific options
- User-defined classes:
  - Identify hosts with a specific configuration requirement
  - Can be configured on the host by an administrator
  - Can be used to override the default options
- DHCP policy-based assignment:
  - Assign options by vendor or user class

When designing your DHCP implementation, you need to determine if you will use option classes in your enterprise. You use option classes to provide unique configurations to specific types of client computers. The Windows Server 2012 implementation of DHCP supports two types of option classes: vendor-defined classes, and user-defined classes.

Vendor-Defined Classes

A DHCP client can use vendor-defined classes to identify its vendor type and to configure it to the DHCP server when obtaining a lease. The client must include the vendor class ID option (option code 60) when it requests or selects a lease from a DHCP server.

User-Defined Classes

User-defined classes identify a DHCP client by its type. A client type refers to characteristics such as a dial-up connection or a portable computer. You configure user-defined classes to manage DHCP options that you want to assign to clients requiring a common configuration. You can use user-defined classes to override default options such as a default gateway.

DHCP Policy-Based Assignment

Windows Server 2012 introduces a new policy-based feature that you can use to group client computers based on attributes that are stored in the DHCP client request packets. You can use this feature to provide for more granular, targeted administration of DHCP clients.

You can use the following fields in the DHCP client request packet to identify and categorize clients:

- Vendor Class
- User Class
- MAC address
- Client Identifier
- Relay Agent Information

After determining the characteristics of your DHCP client from the data contained in the fields, you can then use your policy to:

- Assign addresses from a particular range.
- Assign standard DHCP options.
- Assign vendor-specific DHCP options.

Policy-based assignment supports the following common scenarios:

- Multiple device types. You can define IP addressing characteristics based on the device type. You can do this by creating device-specific policies that assign configurations to devices that meet the criteria of your device policies. For example, you can allocate specific IP address ranges to printer devices.

- Multiple roles. You can use policy-based assignment to allocate IP configurations based on the role that a DHCP client performs. For example, you could allocate longer leases to server computers than you would allocate to client computers.
- Virtualization. If you want to differentiate between physical and virtual computers for IP addressing purposes, you can use policy-based assignment for this purpose.

The DHCP policy-based assignment feature provides greater flexibility and control than DHCP option classes.

Lesson 3
Designing an IPAM Provisioning Strategy

Small networks are comparatively easy to administer, at least from an IP configuration management standpoint. You may not even need to implement network services to allocate IP addresses. For example, when a user connects a new laptop to your network, you can configure the wireless access point to allocate the necessary IP configuration.

However, larger networks—particularly enterprise-level networks—are different. You must manage the introduction of new computers and devices, especially the allocation of their IP configurations. In the past, administrators have struggled with a mix of tools and manual procedures to manage this process. With Windows Server 2012 R2, you can use IPAM, which provides a complete framework for all management tasks.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe IPAM.
- Describe the IPAM architecture.
- Describe the requirements for IPAM deployments.
- Deploy IPAM.
- Manage DHCP and DNS servers by using IPAM.
- Integrate DHCP and DNS servers with IPAM.
- Manage and monitor IPAM.
- Design an IPAM deployment topology.
- Execute capacity planning for IPAM.

What Is IPAM?

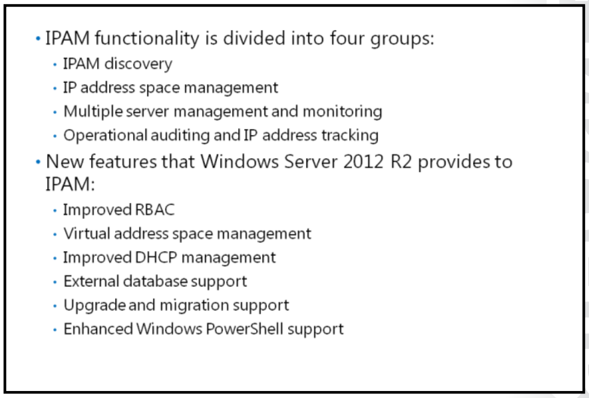

- IPAM functionality is divided into four groups:
  - IPAM discovery
  - IP address space management
  - Multiple server management and monitoring
  - Operational auditing and IP address tracking
- New features that Windows Server 2012 R2 provides to IPAM:
  - Improved RBAC
  - Virtual address space management
  - Improved DHCP management
  - External database support
  - Upgrade and migration support
  - Enhanced Windows PowerShell support

Managing IP address allocation in large networks can be a complex task. IPAM provides a framework for discovering, auditing, and managing the IP address space of your network. It enables you to monitor and administer both DHCP and DNS, and provides a comprehensive view of where specific IP addresses are allocated.

You can configure IPAM to collect statistics from both domain controllers and network policy servers. The resultant data is recorded in a Windows® Internal Database, or an external Microsoft® SQL Server® database.

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 Remote Server Administration Tools (RSAT).

IPAM consists of four modules that provide the following functionality:

- IPAM discovery. You configure IPAM to use AD DS to discover servers running Windows Server 2012, Windows Server 2008 R2, and Windows Server 2008, and that are domain controllers or have either DNS or DHCP installed. 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.
- Multiple server management and monitoring. You can manage and monitor multiple DHCP servers. This enables tasks to run across multiple servers. For example, you can configure and edit DHCP properties and scopes, and track the status of DHCP and scope utilization. You also can 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 collect, manage, and view details of configuration changes from managed DHCP servers. You can also collect address lease tracking from DHCP lease logs, and collect logon event information from Network Policy Server (NPS) and domain controllers.

In addition to these functionalities, Windows Server 2012 R2 provides the following capabilities to IPAM:

- Improved RBAC. You can now create your own roles, and there are more roles readily available.
- Virtual address space management. IPAM integrates with Microsoft System Center 2012 R2 Virtual Machine Manager (VMM) for managing physical and virtual IP addresses used by a VMM infrastructure.
- Improved DHCP management. You can use new operations from IPAM to monitor and manage DHCP servers.
- External database support. You can use a SQL Server database to store IPAM data. Windows Internal Database is still supported as well.
- Upgrade and migration support. You can migrate your IPAM settings from Windows Server 2012 to Windows Server 2012 R2.
- Enhanced Windows PowerShell support. There are 55 new Windows PowerShell cmdlets you can use to manage IPAM.

Note: For the complete list of IPAM Server Cmdlets in Windows PowerShell, visit http://go.microsoft.com/fwlink/?LinkID=392009.

IPAM Components

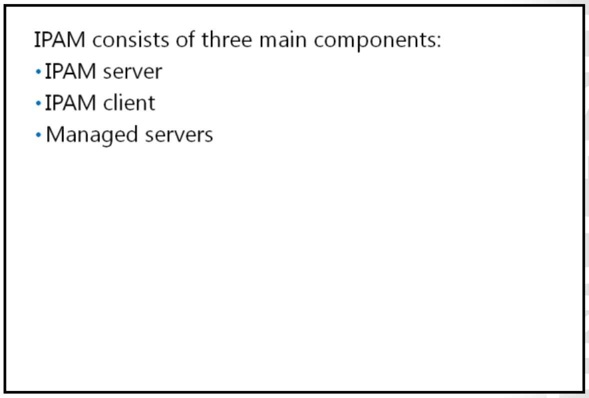

IPAM consists of three main components:

- IPAM server
- IPAM client
- Managed servers

IPAM consists of the following three main components:

- IPAM server. The IPAM server performs data collection from the managed servers. In addition, the IPAM server manages the IPAM database and provides RBAC. You can install the IPAM server on a computer running Windows Server 2012, or Windows Server 2012 R2.
- IPAM client. The IPAM client provides the client computer interface and interacts with the IPAM server. The IPAM client invokes Windows PowerShell cmdlets to perform remote management, DHCP configuration, and DNS monitoring. You can install the IPAM client on a computer running Windows 8.1, Windows 8, Windows Server 2012 R2, or Windows Server 2012.
- Managed servers. You can use IPAM to manage domain controllers, NPS, DNS, and DNS servers running Windows Server 2008 or newer Windows Server operating systems.

IPAM Deployment Requirements

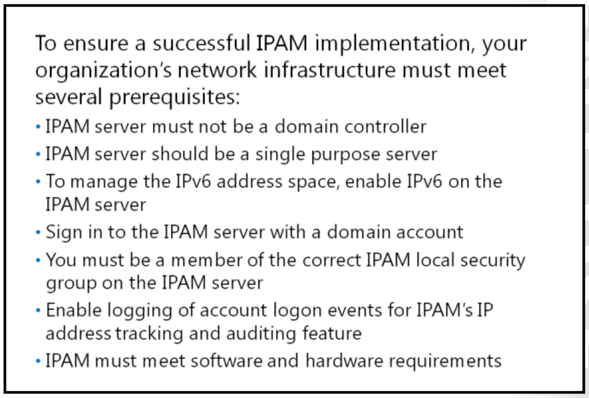

To ensure a successful IPAM implementation, your organization's network infrastructure must meet several prerequisites:

- IPAM server must not be a domain controller
- IPAM server should be a single purpose server
- To manage the IPv6 address space, enable IPv6 on the IPAM server
- Sign in to the IPAM server with a domain account
- You must be a member of the correct IPAM local security group on the IPAM server
- Enable logging of account logon events for IPAM's IP address tracking and auditing feature
- IPAM must meet software and hardware requirements

To ensure a successful IPAM implementation, your organization's network infrastructure must meet several prerequisites:

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

The server on which you intend to deploy IPAM must meet the following hardware and software requirements:

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

When designing an IPAM deployment, consider the following factors:

- IPAM can manage only a single Active Directory forest.
- You should not install IPAM on a domain controller, on a DHCP server, or on a DNS server.
- 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.
- You can define the scope of discovery to a subset of domains in the Active Directory forest.
- A single IPAM server can support up to:
    - 150 DHCP servers and 500 DNS servers.
    - 6,000 DHCP scopes and 150 DNS zones.
- IPAM stores three years of forensic data (IP address leases, host MAC addresses, and user sign-in and sign-out information) for 100,000 users, in a single database. There is no database purge policy, and the administrator must purge the data manually as needed.
- IPAM provides IP address utilization trends only for IPv4.
- IPAM provides IP address reclamation support only for IPv4.
- IPAM does not check for IP address consistency with routers and switches.

Demonstration: Deploying IPAM

This demonstration shows you how to:

- Install IPAM.
- Provision IPAM.

Demonstration Steps

Install IPAM

- On LON-SVR2, in Server Manager, use the Add Roles and Features Wizard to add the IPAM feature and all required supporting features.

Provision IPAM

1. In the IPAM Overview pane, provision the IPAM server by using Group Policy.
2. Enter **IPAM** as the Group Policy Object (GPO) name prefix, and provision IPAM.

Managing DHCP and DNS Servers

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

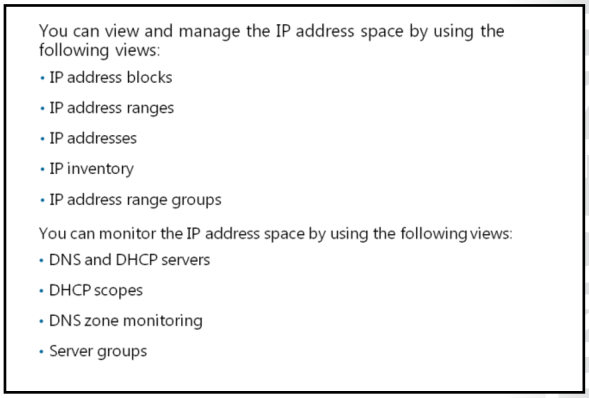

You can view and manage the IP address space by using the following views:

- IP address blocks
- IP address ranges
- IP addresses
- IP inventory
- IP address range groups

You can monitor the IP address space by using the following views:

- DNS and DHCP servers
- DHCP scopes
- DNS zone monitoring
- Server groups

IPAM also enables you to:

- Detect overlapping IP address ranges that are defined on different DHCP servers.
- Find free IP addresses within a range.
- Create DHCP reservations.
- Create DNS records.

IPAM provides a number of ways to filter the view for 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
- 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. An *IP block* is an IP subnet that is marked by a start IP address and an end IP address. You can use IP address blocks to create and allocate IP address ranges to DHCP. You can add, import, edit, and delete IP address blocks. IPAM maps IP address ranges to the appropriate IP address block automatically, based on the boundaries of the range.

IP Address Ranges

IP address ranges are the next hierarchical level of IP address space entities after IP address blocks. An *IP address range* is a portion of a subnet, containing a start IP address and an end IP address. IP address ranges typically correspond to a DHCP scope, or to a static IPv4 or IPv6 address range or address pool that is used to assign addresses to hosts.

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 maps an address to the appropriate range automatically, based on the starting and ending address of the IP address range.

IP Address Inventory

In this 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 within the IP addresses view. You can use this group to customize the way your address space displays for managing and tracking IP usage.

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. You can define logical groups by selecting the grouping criteria from built-in or user-defined custom fields.

Monitoring DHCP and DNS Servers

IPAM enables automated, periodic service monitoring of DHCP and DNS servers across a forest. You can monitor DHCP and DNS servers by using the views listed in the following table.

View	Description
DNS and DHCP servers	By default, IPAM arranges managed DHCP and DNS servers by their network interface in /16 subnets for IPv4, and /48 subnets for IPv6. You can select the view to see only DHCP scope properties, only DNS server properties, or both.
DHCP scopes	You can view utilization statistics that IPAM collects periodically from managed DHCP servers. You can track important scope properties such as **Name**, **ID**, **Prefix Length**, and **Status**.
DNS zone monitoring	By default, IPAM enables zone monitoring for forward and reverse lookup zones. IPAM uses event information collected from DNS servers to define the zone status.
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. You define groups by selecting the grouping criteria from the built-in fields or user-defined fields.

Prior to managing servers by using IPAM, you must discover the servers and then configure them for management in IPAM.

Demonstration: Integrating DHCP and DNS Servers with IPAM

In this demonstration, you will see how to:

- Configure permissions on managed servers.
- Add managed servers to IPAM.

Demonstration Steps

1. On LON-SVR2, in the IPAM Overview pane, start the server discovery process.
2. In the IPAM Overview pane, add the servers to be managed.
3. Verify that IPAM access is currently blocked.
4. 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
```

5. Set the manageability status to **Managed** for both servers.

6. Switch to LON-DC1 and force the update of Group Policy.
7. Switch to LON-SVR1 and force the update of Group Policy.
8. Switch back to LON-SVR2, and refresh the IPv4 view. This process may take up to five minutes for the status to change.
9. In the IPAM Overview pane, retrieve data from the managed server. This action may take five or more minutes to complete.

IPAM Management and Monitoring

You can use the IPAM address space management feature to view, monitor, and manage the IP address space on the network. The address space management feature supports IPv4 public and private addresses, and IPv6 global and unicast addresses.

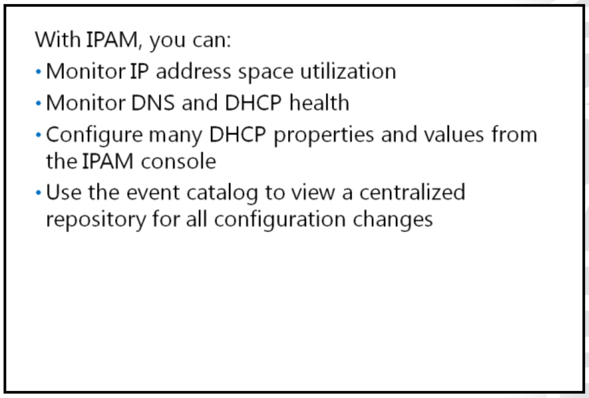

With IPAM, you can:

- Monitor IP address space utilization
- Monitor DNS and DHCP health
- Configure many DHCP properties and values from the IPAM console
- Use the event catalog to view a centralized repository for all configuration changes

Utilization Monitoring

IPAM maintains utilization data for:

- IP address ranges
- IP address blocks
- IP range groups

You can configure thresholds for the percentage of the IP address space that is utilized, and then use those thresholds to determine under-utilization or overutilization. You also can perform utilization trend building and reporting for IPv4 address ranges, blocks, and range groups.

Monitoring DHCP and DNS

Using IPAM, you can monitor DHCP and DNS servers from any physical location in 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.

You can use the IPAM monitoring view to check the status and health of selected sets of Windows Server 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. You can also use the monitoring view to organize the managed servers into logical server groups.

For DHCP servers, you can use the server view 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, you can use this view to track all zones that are configured on the server, along with details of the zone type. You can also use the view 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

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

- Edit DHCP server properties
- Edit DHCP server options
- Create DHCP scopes
- Manage DHCP policies

- Manage DHCP failover
- Manage DHCP filters
- Manage DHCP reservations
- 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 use the central console in the IPAM server to start the DNS management console for any managed DNS server. Using the DNS management console, you can then 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, and 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. You can use these configuration event catalogs to view, query, and generate reports of the consolidated configuration changes, along with details specific for each record.

IPAM Deployment Topologies

When deploying IPAM, you can select from the following three topologies:

- Distributed. You deploy an IPAM server to each site in your forest.
- Centralized. You deploy a single IPAM server for your entire forest.
- Hybrid. In addition to the centralized IPAM server, you also deploy an IPAM server to each site.

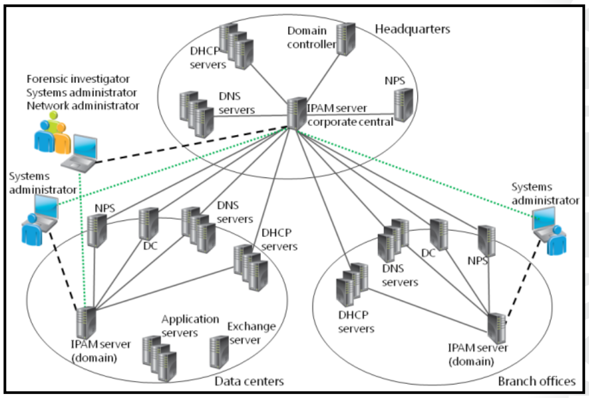

In hybrid and centralized deployments, each IPAM server has its own set of managed servers. However, you can use Windows PowerShell cmdlets to export settings from one server to another.

A managed server in a hybrid deployment can be managed by multiple IPAM servers. For instance, you can have an IPAM server for each site in an organization, and then have one IPAM server that is used to manage all servers in all sites. By doing this, a managed server in any given site will be managed by the IPAM server in the site to which it belongs, and by the central IPAM server for the organization.

IPAM servers do not communicate among themselves, and there is no mechanism for data synchronization between them. Each IPAM server must be configured and managed individually. Hybrid environments contain multiple IPAM servers that manage the same servers. For example, you may have an

IPAM server in the main office managing all DNS and DHCP servers in the organization. You may also have an IPAM server in the regional office that manages only the DNS and DHCP servers in the regional office. In this scenario, the DNS and DHCP servers in the regional office are managed by both IPAM servers.

Capacity Planning for IPAM

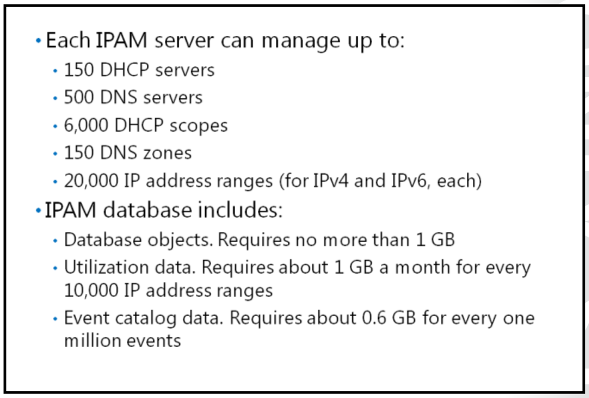

- Each IPAM server can manage up to:
  - 150 DHCP servers
  - 500 DNS servers
  - 6,000 DHCP scopes
  - 150 DNS zones
  - 20,000 IP address ranges (for IPv4 and IPv6, each)
- IPAM database includes:
  - Database objects. Requires no more than 1 GB
  - Utilization data. Requires about 1 GB a month for every 10,000 IP address ranges
  - Event catalog data. Requires about 0.6 GB for every one million events

When designing an IPAM solution and choosing the correct IPAM topology, there are a few considerations that you need to take into account. For instance, each IPAM server has the following limitations:

- Can manage up to 150 DHCP servers
- Can manage up to 500 DNS servers
- Can manage up to 6,000 DHCP scopes
- Can manage up to and 150 DNS zones
- Can manage up to 20,000 IP address ranges for IPv4
- Can manage up to 20,000 IP address ranges for IPv6

If your environment requires more items in any of the categories, you will need to use multiple IPAM servers. You also can use multiple servers to assign different roles per server. For instance, you could have a single IPAM server to manage all DHCP servers in a site, and a separate IPAM server to manage all DNS and domains controllers.

In addition to determining the number of servers to use based on the limitations, topology, and functionality, you also need to plan for the disk capacity that the IPAM servers will use. IPAM servers store the following information in the IPAM database:

- Database objects. These include IP address blocks, IP address ranges, IP address records, custom fields, DHCP settings, and other managed server data. These data requires no more than 1 GB of space in the IPAM database.
- Utilization data. IPAM maintains a sample data set that contains statistical data for IP utilization over time. The amount of data collected for statistics purposes varies according to the number of IP address ranges managed by the IPAM server. Monthly usage data is about 1 GB for every 10,000 IP address ranges. Remember that IPAM does not provide a mechanism for purging historical data. Therefore, you must allocate enough space to store data over time. For instance, if a server manages 20,000 IP address ranges, and if you want to store data for five years (or 60 months), the space required to store utilization data would be 2 x 60 x 1 GB, or 120 GB.
- Event catalog data. IPAM collects DNS, DHCP, NPS, and domain controller event log data. The amount of space necessary to store event data varies according to the number of events generated by each of these services. For each one million events, you need 0.6 GB of space in the database.

Note: IPAM does not provide a mechanism to remove data from its database. You must create your own procedure to ensure that the database does not become too large.

Additional Reading: For more information on IPAM deployment and capacity planning, visit http://go.microsoft.com/fwlink/?LinkID=391892.

Lesson 4
Managing Servers and Address Spaces by Using IPAM

Some organizations have dozens, sometimes even hundreds of DHCP and DNS servers, depending on the number of physical sites and connected devices that they manage. Furthermore, some organizations manage virtual networks in which IP address ranges are duplicated across each virtual network. By using IPAM, you can manage these servers from a centralized location, and cross-reference IP address space information for all your physical and virtual networks.

Lesson Objectives

After completing this lesson, you will be able to:

- Plan DHCP server and scope management.
- Plan DNS server and zone management.
- Plan address space management.
- Manage address spaces.
- Integrate IPAM and VMM.

Planning DHCP Server and Scope Management

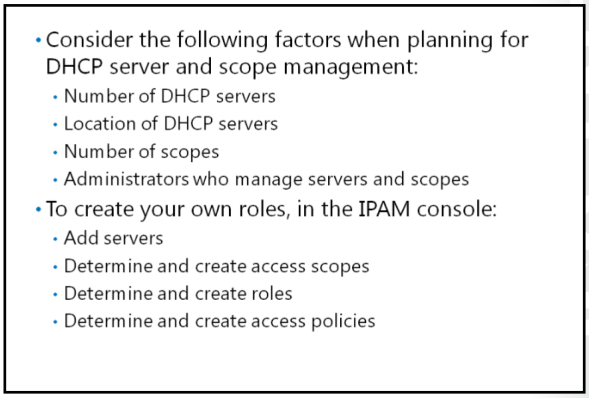

- Consider the following factors when planning for DHCP server and scope management:
  - Number of DHCP servers
  - Location of DHCP servers
  - Number of scopes
  - Administrators who manage servers and scopes
- To create your own roles, in the IPAM console:
  - Add servers
  - Determine and create access scopes
  - Determine and create roles
  - Determine and create access policies

After you install and provision an IPAM server, you can start a discovery task to locate servers that IPAM will manage on your network. After the discovery task locates the servers, you can configure them to allow management through IPAM. You also can change their manageability status to allow IPAM administrators to manage those servers.

However, before you start adding servers to an IPAM server, you need to revisit your topology choice, existing network infrastructure, and operations to determine which DHCP servers to add. You also will need to determine which administrator will be responsible for managing those servers. Consider the following factors to plan for DHCP server and scope management:

- Number and location of DHCP servers to be managed. Remember an IPAM server can manage up to 150 DHCP servers. If you have more than 150 DHCP servers to manage, you need to plan for more IPAM servers, and decide which DHCP servers to add to each IPAM server. You usually divide the DHCP servers by IPAM server based on location.

- Current scopes in each server. Remember an IPAM server can manage up to 6,000 DHCP scopes. Before deciding which DHCP servers an IPAM server will manage, you need to ensure that the number of scopes in each DHCP server is below 6,000.

- Current DHCP administrators. One of the main advantages of using IPAM is the ability to manage roles that you can assign to manage separate objects in IPAM. Based on how you assign administrators to manage different servers and scopes in DHCP, you can create multiple roles and assign administrators to them. The following roles are available by default:

- IPAM DHCP administrator. This administrator manages all DHCP servers and scopes.
- IPAM DHCP reservations administrator. This administrator manages DHCP reservations.
- IPAM DHCP scope administrator. This administrator manages DHCP scopes.

You can create your own roles based on your organization's needs. For instance, imagine that your organization contains two DHCP servers named DHCP1 and DHCP2. Both servers are configured to work by using DHCP failover for four different scopes:

- Scope1 – Intranet
- Scope2 – Perimeter
- Scope3 – Development
- Scope4 – Test

You want to manage both servers by using IPAM; however, you want to allow users from the Development team to manage the Development scope. You can do so by using IPAM RBAC to create the following objects:

- Access scope. Access scopes determine what objects a user can access. In the example, you could create an access scope named DevScope and then link it the Scope3 DHCP scope.
- Role. A role contains a list of operations that a certain user or group associated with the roll can perform. In the example, you want the Development team users to manage the scope named Scope3. You can create a role named DevScope Managers, and then you can:
  - Specify the different operations you want to allow users in that role to perform.
  - Associate the role with a domain group that contains all users in the Development team that will be responsible for managing the DHCP scope for the development environment.
- Access policy. Access policies combine a role with an access scope. In the example, you could create an access role named DevScopePolicy to associate the DevScope access scope to the DevScope Managers role.

Planning DNS Server and Zone Management

You need to follow a similar process for planning your DNS server and zone management settings in IPAM. Consider the following factors when planning for DNS server and zone management:

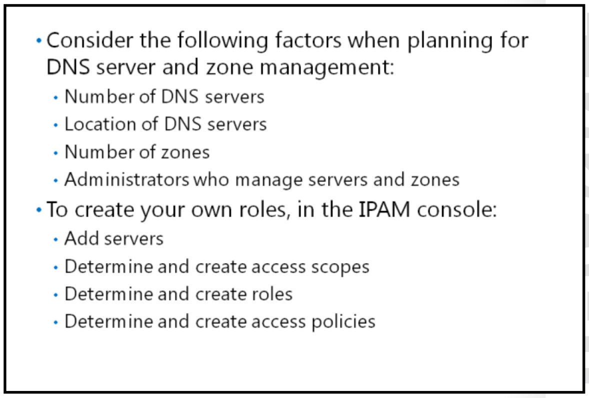

- Consider the following factors when planning for DNS server and zone management:
  - Number of DNS servers
  - Location of DNS servers
  - Number of zones
  - Administrators who manage servers and zones
- To create your own roles, in the IPAM console:
  - Add servers
  - Determine and create access scopes
  - Determine and create roles
  - Determine and create access policies

- Number and location of DNS servers to be managed. Remember an IPAM server can manage up to 500 DNS servers. If you have more than 500 DNS servers to manage, you will need to plan for more IPAM servers, and then decide which DNS servers to add to each IPAM server. You usually divide the DNS servers by IPAM server based on location.
- Current zones in each server. An IPAM server can manage up to 150 DNS zones. Before deciding which DNS servers an IPAM server will manage, you need to ensure that the number of zones in each DNS server is below 150.

- Current DNS administrators. One of the main advantages of using IPAM is the ability to manage roles that you can assign to manage separate objects in IPAM. Based on what administrators can manage different servers and zones in DNS, you can create multiple roles and then assign administrators to them. The DNS record administrator role is available by default for DNS. This role is responsible for managing all DNS resource records.

You can create your own roles based on your organization needs. For example, imagine that your organization contains two DNS servers, named DNS1 and DNS2. DNS1 contains a primary zone for the adatum.com Internet zone. DNS2 contains a primary zone for the adatum.com intranet site.

You want to manage both servers by using IPAM. However, you want to allow users from the Internet team to manage the adatum.com Internet primary zone. You can do so by using IPAM RBAC to create the following objects:

- Access scopes. Access scopes determine what objects a user can access. In the example, you could create an access scope named ADatumInternetZone, and then link it to the adatum.com Internet zone in DNS1.
- Role. A role contains a list of operations a certain user or group associated to it can perform. In the example, you want the Internet team users to manage the adatum.com Internet zone. You can create a role named Internet DNS zone administrators, specify the different operations you want to allow users in that role to perform, and then associate the role with a domain group that contains all users in the Internet team that should be able to manage the DNS zone.
- Access policy. Access policies combine a role with an access scope. In the example, you could create an access role named InternetDNSPolicy to associate the ADatumInternetZone access scope to the Internet DNS zone administrators role.

Planning Address Space Management

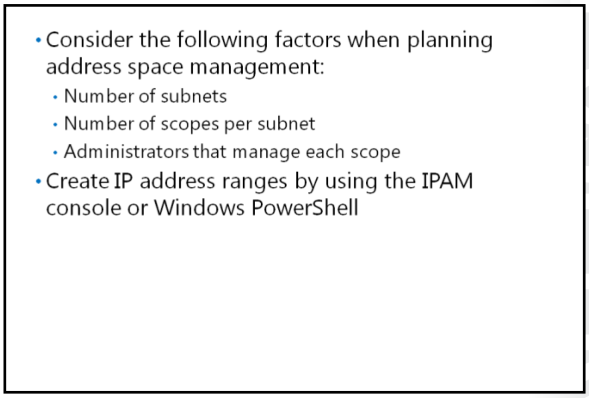

- Consider the following factors when planning address space management:
  - Number of subnets
  - Number of scopes per subnet
  - Administrators that manage each scope
- Create IP address ranges by using the IPAM console or Windows PowerShell

IPAM makes it easier for you to manage your organization's address space by allowing you to create and manage IP address ranges. An *IP address range* is a contiguous space of IPv4 or IPv6 addresses for which you want to track utilization and manage DHCP settings. You can execute the following IP address range operations by using the IPAM client console or Windows PowerShell:

- Create an IP address range
- Add an IP address
- View IP address ranges and addresses
- Track utilization of IP ranges
- Assign IP addresses to devices
- Reclaim IP addresses

Before creating IP address ranges, consider the following:

- Existing physical subnets
- IP scopes for each subnet
- Users that manage each scope and subnet

An IP address range must contain contiguous IP addresses in a single subnet. However, a subnet may contain different DHCP scopes. You need to determine if different users or the same set of users should manage the different scopes. If the same set of users manage the different DHCP scopes, you can create a single IP address range for all the scopes within the same subnet.

Demonstration: Managing Address Spaces

In this demonstration, you will see how to:

- Create an IP address range.
- Add used IP addresses in the range.
- View IP address ranges.
- Track utilization of IP ranges.
- Assign a free IP address to a device.

Demonstration Steps

1. Create an IP address range for all IP addresses in the 172.16.1.0/24 subnet by using the **Add-IpamRange** cmdlet.
2. Add the IP address for LON-DC1, LON-SVR1, and LON-SVR2, and record them as being in use by using the **Add-IpamAddress** cmdlet.
3. View the existing IP address ranges in IPAM by using the **Get-IpamRange** cmdlet.
4. View all IP address ranges that have less than 50 percent of their IP address in use by using the **Get-IpamRange** cmdlet.
5. Find a free IP address in the 172.16.1.0 range by using the **Find-IpamFreeAddress** cmdlet, and store it in a variable named **$freeIP**.
6. View the contents of the **$freeIP** variable.
7. Add the free IP address to a printer device with a MAC address of AA-AA-AA-BB-BB-BB by using the **Add-IpamAddress** cmdlet, and save it to a variable named **$IP**.
8. View the contents of the **$IP** variable.
9. Unprovision the IP address for the printer by using the **Remove-IpamAddress** cmdlet.

IPAM Integration with VMM

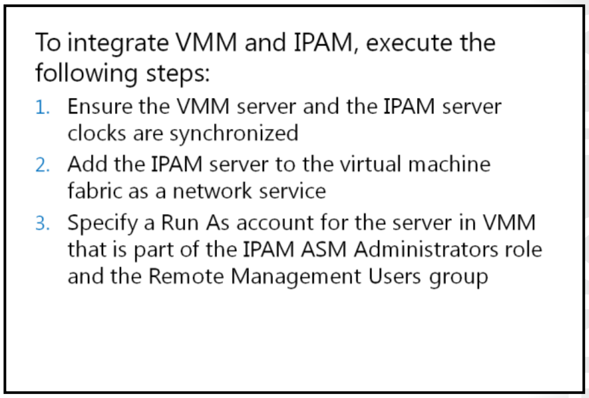

To integrate VMM and IPAM, execute the following steps:

1. Ensure the VMM server and the IPAM server clocks are synchronized
2. Add the IPAM server to the virtual machine fabric as a network service
3. Specify a Run As account for the server in VMM that is part of the IPAM ASM Administrators role and the Remote Management Users group

In Windows Server 2012 R2, IPAM integration with VMM allows you to manage virtualized IP address spaces from the IPAM console itself. You can manage them through the new Virtualized Address Space node in the IPAM console.

After you integrate IPAM and VMM, you can perform the following tasks from VMM:

- Configure address space by creating subnets, pools, virtual networks, and virtual local area networks (VLANs). All settings are pushed to IPAM.
- Detect conflicts. IPAM detects any conflicts and raises notifications. IPAM administrators can make changes to settings, which are synchronized back to VMM.
- Create logical networks. Each logical network is treated as an address space in IPAM. IP addresses can overlap on logical networks, providing VMM is using network isolation.

To integrate VMM and IPAM, perform the following steps:

1. Ensure that the VMM server and the IPAM server clocks are synchronized.
2. Add the IPAM server to the virtual machine fabric as a network service.
3. Specify a Run As account for the server in VMM that is part of the IPAM ASM Administrators role and the Remote Management Users group.

Lab: Designing and Maintaining an IP Configuration and IP Address Management Solution

Scenario

Over the years, IP address configuration and management at A. Datum Corporation has been conducted on mainly a reactive basis. A. Datum recently acquired a new server for a department, and the server was allocated an IP configuration. DHCP was introduced to help automate and centralize the IP address allocations, but introducing this key infrastructure service was often not well planned. This resulted in little logic to the way in which devices were allocated their IP configurations.

During the planned move to Windows Server 2012 R2, the manager of the IT department has seen an opportunity to implement IPAM with Windows Server 2012 R2. After you have selected an appropriate addressing scheme for the Contoso, Ltd network, you must plan how best to use IPAM to manage your IP addressing strategy. You must consider whether to centralize or distribute the topology, what to integrate with IPAM (DNS or DHCP), and how best to secure IPAM.

Having selected a suitable IPAM configuration, you also must consider how to allocate IP addressing to server and client computers by using DHCP. It is important that the DHCP service is highly available.

Objectives

After completing this lab, you will be able to:

- Plan DHCP to support your proposal.
- Plan an IPAM deployment.
- Implement DHCP and IPAM.

Lab Setup

Estimated Time: 75 minutes

Virtual machines	20413C-LON-DC1 20413C-LON-SVR1 20413C-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**, and then click **Hyper-V Manager**.
2. In Hyper-V® Manager, click **20413C-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on using the following credentials:

- User name: **Administrator**
- Password: **Pa$$w0rd**
- Domain: **Adatum**

5. Repeat steps 2 through 4 for 20413C-LON-SVR1 and 20413C-LON-SVR2.

Exercise 1: Planning Dynamic Host Configuration Protocol (DHCP) to Support Your Proposal

Scenario

You must now consider where best to place DHCP servers within the Contoso organization to support the proposed addressing scheme.

Supporting Documentation

Contoso DHCP Deployment Strategy

Document Reference Number: BS01107/2

Document Author	Brad Sutton
Date	11th July

Requirements Overview

Plan a DHCP deployment strategy for Contoso to support the following objectives:

- All client and server computers must be able to obtain an IPv4 configuration automatically.
- Your plan must provide for high availability of the DHCP role. A failure of a DHCP server anywhere in the Contoso organization should not result in clients failing to obtain an IP address.

Additional Information

- Routers are DHCP-capable; they support the propagation of DHCP broadcast traffic.
- There is no requirement to provide support for IPv6 at present.
- You must implement the least number of DHCP servers as possible.
- The IP addressing scheme for Contoso will be as follows:

Europe regional hub locations

Location	Branches	Subnet
Paris	Main office	172.32.32.0/19
Munich	Branch 1	172.32.41.0/25
	Branch 2	172.32.42.0/25
	Branch 3	172.32.43.0/25
Barcelona	Branch 1	172.32.51.0/25
	Branch 2	172.32.52.0/25
	Branch 3	172.32.53.0/25
	Branch 4	172.32.54.0/25
	Branch 5	172.32.55.0/25
Rome	Branch 1	172.32.61.0/25
	Branch 2	172.32.62.0/25
	Branch 3	172.32.63.0/25
	Branch 4	172.32.64.0/25
	Branch 5	172.32.65.0/25
Athens	Branch 1	172.32.71.0/25

Contoso DHCP Deployment Strategy

	Branch 2	172.32.72.0/25

European branch office/regional distribution locations

Location	Branches	Subnet
Germany	Branch 1	172.32.81.0/25
	Branch 2	172.32.82.0/25
	Branch 3	172.32.83.0/25
Spain	Branch 1	172.32.91.0/25
	Branch 2	172.32.92.0/25
	Branch 3	172.32.93.0/25
	Branch 4	172.32.94.0/25
	Branch 5	172.32.95.0/25
Italy	Branch 1	172.32.101.0/25
	Branch 2	172.32.102.0/25
	Branch 3	172.32.103.0/25
	Branch 4	172.32.104.0/25
	Branch 5	172.32.105.0/25
Greece	Branch 1	172.32.111.0/25
	Branch 2	172.32.112.0/25

Proposals

1. How should clients and servers in the head office in Paris obtain an IP configuration?

2. How should clients in regional hub offices obtain an IP configuration?

3. How will you provide high availability for DHCP in the Paris office?

Contoso DHCP Deployment Strategy
4. How will you provide high availability for DHCP in the regional hub sites?
5. How many scopes do you need to configure on the DHCP servers in the regional hub sites?

The main tasks for this exercise are as follows:

1. Read the supporting documentation.
2. Update the proposals document with your planned course of action.
3. Examine the suggested proposals in the Lab Answer Key.
4. Discuss your proposed solution with the class, as guided by your instructor.

▶ Task 1: Read the supporting documentation

- Read the documentation provided.

▶ Task 2: Update the proposals document with your planned course of action

- Answer the questions in the proposals section of the Contoso DHCP Deployment Strategy document.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones in the Lab Answer Key.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you will have planned DHCP to support the Contoso IP addressing scheme.

Exercise 2: Planning an IP Address Management (IPAM) Deployment

Scenario

The increasing size and complexity of the network means that the historically reactive approach to IP configuration and management is no longer acceptable. Having successfully designed an IP configuration for Contoso, you must now consider how best to implement IPAM in A. Datum, and then in Contoso and Trey Research. Specifically, you must consider whether to deploy the IPAM role in a distributed or centralized way.

Supporting Documentation

A. Datum IPAM Deployment Plan
Document Reference Number: BS01207/1

A. Datum IPAM Deployment Plan

Document Author	Brad Sutton
Date	12th July

Requirements Overview

Plan an IPAM deployment for A. Datum to support the following objectives:

- You must initially deploy IPAM to A. Datum.
- You will then deploy IPAM to Contoso and Trey Research in the near future.

Additional Information

- Network policy servers are deployed in A. Datum to support virtual private network (VPN) connections.
- Windows Server 2012 R2 and Windows Server 2008 R2 are deployed in A. Datum. All domain controllers are installed with Windows Server 2012 R2.
- Windows Server 2008 R2 is deployed in Trey Research in a separate AD DS forest.
- Currently, Contoso is only running UNIX. Plans are in place to deploy Windows Server 2012 at Contoso.

Proposals

1. Is a centralized or distributed topology better for the DHCP server infrastructure?

2. Which server roles will you manage?

3. What AD DS considerations should you take into account for the inclusion of the Contoso and Trey Research organizations?

4. What are the organizational and server-level prerequisites for IPAM?

The main tasks for this exercise are as follows:

1. Read the supporting documentation.

2. Update the proposals document with your planned course of action.

3. Examine the suggested proposals in the Lab Answer Key.

4. Discuss your proposed solution with the class, as guided by your instructor.

▶ Task 1: Read the supporting documentation

- Read the documentation provided.

▶ Task 2: Update the proposals document with your planned course of action

- Answer the questions in the proposals section of the A. Datum IPAM Deployment Plan document.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones in the Lab Answer Key.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you will have planned an IPAM deployment strategy for A. Datum.

Exercise 3: Implementing DHCP and IPAM

Scenario

You must now deploy DHCP and IPAM in the A. Datum forest to support your DHCP and IPAM implementation strategies.

The main tasks for this exercise are as follows:

1. Install the DHCP Server role.
2. Configure a DHCP failover relationship.
3. Install IPAM.
4. Configure Group Policy Object (GPO) Settings.
5. Configure IP management server discovery.
6. Configure managed servers.
7. Configure and verify a DHCP scope with IPAM.
8. Configure IP address blocks, record IP addresses, and create DHCP reservations.

▶ Task 1: Install the DHCP Server role

1. If necessary, sign in to LON-SVR1 as **Adatum\Administrator** with a password of **Pa$$w0rd**.
2. Use Server Manager to install the DHCP Server role.
3. Complete the DHCP post install configuration.

▶ Task 2: Configure a DHCP failover relationship

1. Switch to LON-DC1.
2. If necessary, sign in to LON-DC1 as **Adatum\Administrator** with a password of **Pa$$w0rd**.
3. In Server Manager, click **Tools**, and then in the drop-down list box, click **DHCP**.
4. In the DHCP console, launch the Configuration Failover Wizard.
5. Configure failover replication with the following settings:
   - Partner server: **172.16.0.11**
   - Relationship Name: **Adatum DHCP Failover**
   - Maximum Client Lead Time: **15 minutes**
   - Mode: **Load balance**

   - Load Balance Percentage: **50%**
   - State Switchover Interval: **45 minutes**
   - Message authentication shared secret: **Pa$$w0rd**

6. Complete the Configuration Failover Wizard.
7. Switch back to LON-SVR1, open the DHCP console, and note that the IPv4 node is active, and that the Adatum scope is configured.
8. Close the DHCP console on both LON-SVR1 and LON-DC1.

▶ Task 3: Install IPAM

1. If necessary, sign in to LON-SVR2 as **Adatum\Administrator** with a password of **Pa$$w0rd**.
2. In Server Manager, use the Add Roles and Features Wizard to add the **IP Address Management (IPAM) Server** feature and all required supporting features.

▶ Task 4: Configure Group Policy Object (GPO) Settings

1. In the Server Manager navigation pane, click **IPAM**.
2. In the IPAM Overview pane, provision the IPAM server by using Group Policy.
3. Enter **IPAM** as the GPO name prefix, and provision IPAM.

▶ Task 5: Configure IP management server discovery

- In the IPAM Overview pane, configure server discovery for the Adatum domain, and then start the server discovery process.

Note: Discovery may take 5 to 10 minutes to run. The yellow bar indicates when discovery is complete.

▶ Task 6: Configure managed servers

1. In the IPAM Overview pane, add the servers to be managed.
2. Verify that IPAM access is currently blocked.
3. 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
```

4. Set the manageability status to **Managed** for both servers.
5. Switch to LON-DC1.
6. Force the update of Group Policy.
7. Verify the IPAM GPOs were applied.
8. Switch to LON-SVR1.
9. Force the update of Group Policy.
10. Verify that the IPAM_DHCP GPO has been applied.

11. Switch back to LON-SVR2 and refresh the server access status and IPv4 view.
12. In the IPAM Overview pane, retrieve data from the managed server.

Note: This action may take five minutes or more to complete.

▶ Task 7: Configure and verify a DHCP scope with IPAM

1. Switch to LON-SVR2.
2. Use IPAM to create a DHCP IPv4 scope for the Paris Office by using the following settings:
   - Start IP address: **172.32.32.2**
   - End IP address: **172.32.32.200**
   - Subnet mask: **255.255.224.0**
   - Router: **172.32.32.1**
   - DNS Server: **172.32.32.2**
3. Verify the new scope in IPAM.
4. Verify the new scope in LON-DC1.
5. On LON-SVR2, use IPAM to configure failover for the Paris Office scope by using the existing failover relationship between LON-DC1 and LON-SVR1.

▶ Task 8: Configure IP address blocks, record IP addresses, and create DHCP reservations

1. On LON-SVR2, in Server Manager, in the IPAM pane, click **IP address Blocks**.
2. In the **Current view** list, click **IP Address Ranges**. Note that due to DHCP failover, 172.32.32.0 is listed twice.
3. Right-click the **172.32.32.0/19** range for **lon-dc1.adatum.com**, and then click **Edit IP Address Range**.
4. In the **Edit IP Address Range** dialog box, click **Reservations**.
5. In the **Reservation** text box, type **172.32.32.2**, and then click **Add**.
6. In the **Edit IP Address Range** dialog box, click **OK**.

Results: After completing this exercise, you will have deployed DHCP and IPAM to support your proposals.

▶ Task: 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 Microsoft Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 through 3 for 20413C-LON-SVR1 and 20413C-LON-SVR2.

Module Review and Takeaways

Review Question(s)

Question: You have two subnets in your organization and want to use DHCP to allocate addresses to client computers in both subnets. You do not want to deploy two DHCP servers. What factors must you consider?

Question: Your organization has grown, and your IPv4 scope has few addresses remaining. What could you do?

Question: What information do you require to configure a DHCP reservation?

Module 5

Designing and Implementing Name Resolution

Contents:

Module Overview

Name resolution is an essential service in modern computer networks. Consequently, your Domain Name System (DNS) namespace structure and server locations can have significant implications on the availability and performance of your networked applications. Therefore, to optimize availability and performance of networked applications, you must design your DNS infrastructure carefully. This includes determining DNS server locations, DNS zone types, and how and where to store DNS zone data.

After you have completed your high-level DNS design, you must consider both how to optimize DNS queries and how you can ensure high availability and security for DNS. Furthermore, in many environments it may be necessary to implement the DNS GlobalNames zone to support NetBIOS applications.

Objectives

After completing this module, you will be able to:

- Design a DNS server-implementation strategy.
- Design a DNS namespace.
- Design and implement a DNS zone strategy.
- Design and configure DNS zone replication and delegation.
- Optimize the DNS server configuration.
- Design DNS for high availability and security.

Lesson 1
Designing a DNS Server Implementation Strategy

The Windows Server® 2012 DNS server role can provide all DNS services within your network. However, to optimize your DNS configuration, you first must determine how many physical or virtual DNS servers you need, where you want to locate those servers, and the specific functional role of each server. Additionally, you must ensure that you secure these DNS servers, to safeguard both the DNS zone data and the network resources to which the DNS zone data is pointing. Without taking the necessary steps to secure your DNS infrastructure, you run the risk of compromising both DNS and the network applications that rely upon it.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe DNS server requirements.
- Describe the specifications that determine DNS server capacity.
- Describe the factors that determine DNS server placement.
- Select suitable DNS server roles.
- Describe the security considerations for DNS servers.
- Explain how to install the DNS Server role.

DNS Server Requirements

To plan your individual DNS Server requirements properly, you must gather the following information about your organization's physical locations:

- The existence of, or plans to include an Active Directory® Domain Services (AD DS) service infrastructure. This is your most important consideration. If you plan to implement Active Directory–integrated zones, then as part of the overall server requirements you must consider the combination of both DNS and AD DS domain controller requirements. If you intend to implement AD DS, you must factor it in to your DNS namespace design.

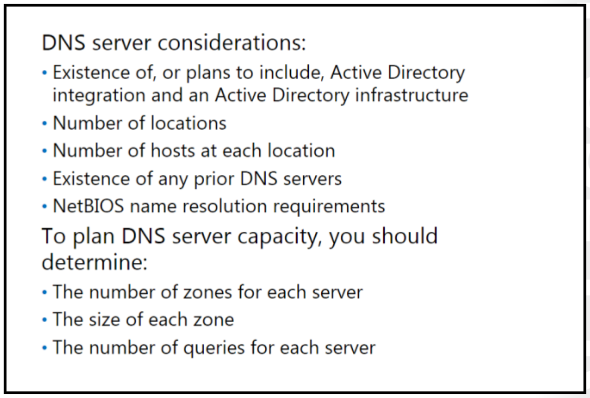

DNS server considerations:

- Existence of, or plans to include, Active Directory integration and an Active Directory infrastructure
- Number of locations
- Number of hosts at each location
- Existence of any prior DNS servers
- NetBIOS name resolution requirements

To plan DNS server capacity, you should determine:

- The number of zones for each server
- The size of each zone
- The number of queries for each server

It is possible to configure the Active Directory domain names to match the DNS domain names. It also is possible to configure both to be different. Additionally, if you implement AD DS, you have the option of including Active Directory–integrated zones in your DNS design, which has a number of benefits that a later topic in this module covers.

- The number of physical locations. Because almost all network applications use host names rather than IP addresses, lack of access to a DNS server can prevent or affect network communications adversely. Although smaller locations may have fewer network hosts, you must provide a mechanism for resolving names and locating services. Consequently, each location usually requires at least one DNS server.

- The number of hosts at each location. The number of hosts and the number of client workstations and server computers determines the number of DNS clients that each location must support. More DNS clients mean more DNS queries, with a corresponding increase in the workload on any DNS server that you configure to support those clients. Consequently, for locations with a large number of DNS clients, consider deploying additional DNS servers.

Host computers can have more than one IP address, and more than one name associated with a given IP address. Host computers also can have services, such as service (SRV) resource records that locate domain controllers. DNS servers that are running Windows Server software have round-robin turned on by default, the entire list of IP addresses associated with one name is returned in a different order each time.

- The existence of any legacy DNS servers, such as Berkeley Internet Name Domain (BIND) servers or the Microsoft Windows NT® Server 4.0 operating system DNS servers. Existing legacy DNS servers might limit the use of DNS features, such as incremental zone transfers. If your current network infrastructure relies upon earlier DNS servers, you may need to configure Windows Server 2012 DNS to support some of these legacy behaviors.
- NetBIOS name resolution requirements. If you have clients that require access to a NetBIOS-based application, you must identify whether these clients are located in the same or different broadcast domains. Because NetBIOS relies on broadcasts to register, release, and resolve names, NetBIOS clients must reside in the same local network segment, or broadcast domain. If they do not, you must provide a means by which clients can register, release, and resolve their NetBIOS names. Depending upon your needs and your network's specific configuration, this might require that you implement Windows® Internet Name Service (WINS), the DNS GlobalNames zone, or both.

The DNS server service fully loads all of its configured zones into memory at startup. If your server operates and loads a large number of zones, and if dynamic updates occur frequently for zone clients, additional memory might improve DNS server performance.

The following factors determine the overall capacity of a DNS server:

- The number of zones that the server hosts. Determining the number of zones that you need is a relatively straightforward process. You often need a single zone for each domain in your DNS namespace. For smaller networks, you can combine domains into a single zone. For example, a single zone, Contoso.com, could store the records for both the DNS domains training.contoso.com and Contoso.com.
- The number of records in each zone. Determining the number of records in each zone also is a straightforward process. However, you must include the total number of hosts within your zone, including service records (SRV record), text, and mail exchanger MX) resource records, and other specialized records.
- The number of queries for records in each zone. This is more difficult to determine. To estimate the number of records queries accurately, you should consider logging current network activity on your existing DNS infrastructure. If you are installing DNS in a new site without an existing infrastructure, after deployment, you should revisit your DNS design based on the number of queries and the loads they generate.

The resource requirements for a DNS server typically are very low, which makes DNS servers suitable virtualization candidates. However, because it is a best practice to use Active Directory–integrated DNS, server requirements must include both DNS and Active directory domain controller specifications.

Determining DNS Server Placement

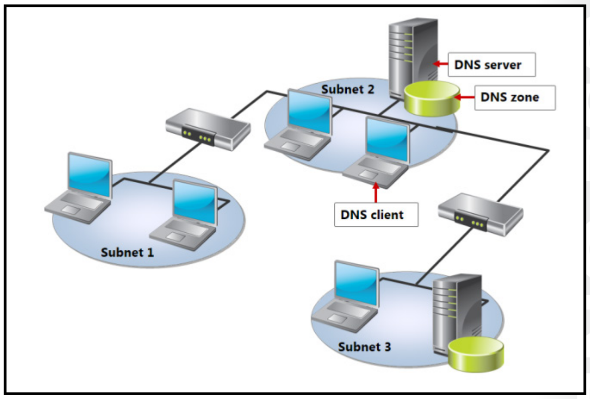

The placement of DNS servers affects application and client performance. When placing DNS servers within your physical network infrastructure, you must consider the following factors:

- DNS resolution over wide area network (WAN) links by many clients can generate significant network traffic. Additionally, resolution over WAN links is slower than local name resolution, and it may affect application performance due to latency. Where network traffic over WAN links is an issue, consider placing a DNS server at the remote network location to reduce query-based traffic. Remember that when you place a DNS server at a remote site, although it reduces the volume of query-based traffic over the WAN link, the server itself can generate additional traffic in terms of zone transfers or replication.
- If a WAN link fails, and if the current site does not cache or contain DNS information, this can affect service availability. Ensure that each location has a local DNS server so that client queries are addressed.

Note: If the WAN link fails, then all client access to applications and services across the link also fails. A local DNS server will provide name resolution to local resources, but access to remote servers will continue to fail as long as the WAN Link is down.

- If a DNS server fails, you need to have a contingency plan for DNS server redundancy. You should plan how clients will fail over from one DNS server to another. Locating multiple DNS servers at smaller sites may not be possible. Therefore, you should consider configuring clients to use the local server as a primary DNS server and a remote server as a secondary DNS server.
- If a DNS server is unavailable, you must determine which computers may experience problems. This enables you to pinpoint sources of error for troubleshooting when problems arise.
- If a DNS server or part of your internal DNS namespace is unavailable, you must determine which applications this will affect.

Your goal when placing DNS servers is to avoid—or at least reduce to an acceptable level—the general unavailability of your network service when a DNS server is unavailable.

Selecting DNS Server Roles

Depending on your requirements, you can implement a DNS server in a specific role, such as a caching-only server, a forwarding server, a nonrecursive server, or an authoritative server.

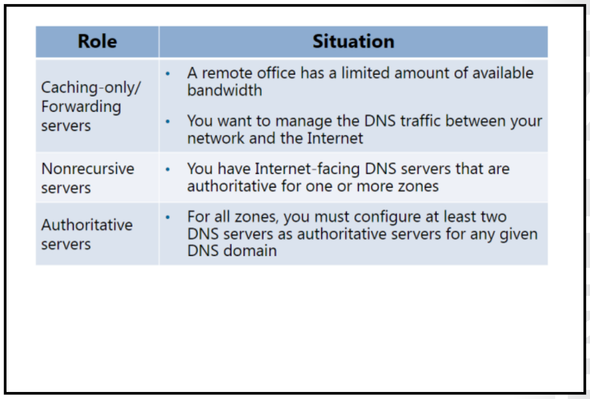

Role	Situation
Caching-only/ Forwarding servers	• A remote office has a limited amount of available bandwidth • You want to manage the DNS traffic between your network and the Internet
Nonrecursive servers	• You have Internet-facing DNS servers that are authoritative for one or more zones
Authoritative servers	• For all zones, you must configure at least two DNS servers as authoritative servers for any given DNS domain

Implementing Caching-Only Servers and Forwarding Servers

A caching-only server is a simple way to ensure that remote sites have a copy of commonly used DNS records. As the server resolves each record, it caches the record locally without having to configure zone transfers or replication. The local name server cannot resolve the initial client query because the server contains no zone data. Consequently, the caching-only server petitions another server, and caches the results. This means that the petitioned server must be available. In a branch office with a local caching-only server, if the WAN link to the site that hosts the authoritative server is unavailable, queries fail if they are not cached.

Similar to caching-only servers, forwarding servers build a local cache of resolved DNS records. However, you configure a forwarding server to forward requests to a specific server, rather than to use root hints to determine the appropriate authoritative servers.

Note: *Root hints* is the mechanism whereby a DNS server can locate records elsewhere in the Internet namespace, for which the petitioned server is not authoritative.

A *conditional forwarder* is a DNS server on a network that forwards DNS queries according to the query's DNS domain name. This is useful when you have multiple DNS namespaces in a forest.

For example, you can configure a DNS server to forward all queries that it receives for names ending with corp.contoso.com, to the IP address of a specific DNS server or to multiple DNS servers. It resolves all other queries in the usual way, which is recursively, up the DNS namespace to the root.

This enables you to create a flexible structure for name resolution. For example, you can forward internal name resolution requests to another internal DNS server, but forward request for all other domains to an Internet-facing DNS server.

Note: Use conditional forwarders if you have multiple internal namespaces, because it provides faster name resolution.

Nonrecursive Servers

A nonrecursive server can resolve only locally hosted DNS records. This is useful on Internet-facing servers to ensure that they resolve only the public records that you configure.

Authoritative Servers

Authoritative name servers are those that contain a local copy of the DNS zone data. You configure them as the point of contact for a given DNS domain, for DNS name servers elsewhere in the global DNS namespace. They also hold DNS records that identify authoritative DNS servers for subdomains in the DNS namespace.

Authoritative DNS servers can resolve queries for the local DNS zone. You configure them with root hints, which enable them to perform query recursion for DNS records for which they are not authoritative.

Security Considerations for DNS Servers

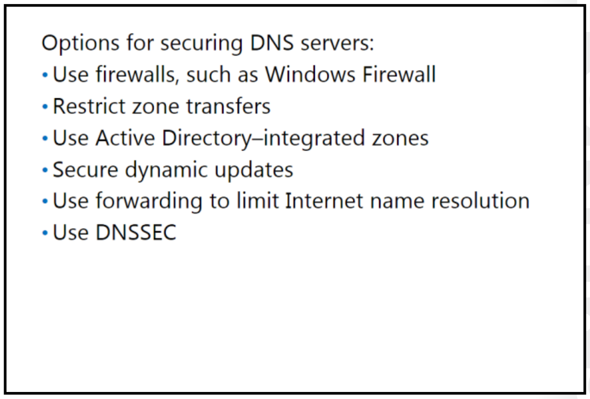

Options for securing DNS servers:

- Use firewalls, such as Windows Firewall
- Restrict zone transfers
- Use Active Directory–integrated zones
- Secure dynamic updates
- Use forwarding to limit Internet name resolution
- Use DNSSEC

DNS servers provide a foundation service for your networked computers. Without a DNS service, many applications can fail. It is important that you secure your DNS servers to ensure continuous availability. If hackers deliberately modify DNS server records, they can direct client computers to inappropriate networked application servers. Therefore, it is important to consider security carefully for your DNS servers.

You can use the following methods to secure your DNS servers:

- Use firewalls (such as Windows Firewall), to restrict communication with DNS servers to authorized IP address ranges and authorized ports only. This prevents any computer that is not part of the configured range from accessing the DNS servers.
- Restrict zone transfers to prevent network information gathering (or *footprinting)* by querying a list of internal resources from the DNS server. By default, primary name servers do not transfer zone data to secondary name servers until you configure the necessary settings. You should configure these settings properly to prevent servers from transferring data to an inappropriate location.

Note: A *zone transfer* is the mechanism whereby DNS servers that support a zone copy data between them. Later sections of this module detail zones and zone transfers.

- Use Active Directory–integrated zones. These zones transfer zone data over encrypted channels during the Active Directory replication process. This helps to ensure that the zone data remains secure while in transit between DNS servers.

Note: If you do not use Active Directory–integrated zones, consider implementing Internet Protocol security (IPsec) to encrypt zone-transfer traffic between servers.

- Enable secure dynamic updates. This helps to ensure that only the host that created the dynamic DNS record can modify it. Clients can register their own IP addresses and host names automatically with their configured DNS server, but only if the server supports the dynamic update protocol. You can configure security for these updates, which means that only the record owner can update the record. The DNS client typically is the owner of its own DNS records in the DNS zone. Use Active Directory–integrated zones for secure dynamic updates.
- Use forwarding to centralize name resolution in your organization. This enables you to restrict name resolution on the Internet to specific servers.

- DNS Security Extensions (DNSSEC). DNSSEC protects DNS clients or DNS resolvers, from receiving forged DNS data from rogue DNS servers that are masquerading as local DNS servers. Resolvers can use DNSSEC to obtain origin authentication of DNS server replies. However, that data is not confidential, since DNSSEC does not encrypt packets. However, it does sign the packet headers digitally.

A subsequent lesson in this course, "Designing DNS for High Availability and Security," provides detailed information about DNS security, including DNSSEC.

Demonstration: Installing the DNS Server Role

This demonstration shows you how to use Server Manager to install the DNS server role.

Demonstration Steps

1. On LON-SVR1, open Server Manager.
2. Use the **Add Roles and Features Wizard** to add the DNS server role.

Lesson 2
Designing the DNS Namespace

When you implement DNS, you must select a namespace design. This involves selecting the DNS domain names, and determining the relationships between these domains. Additionally, you must consider how the DNS namespace relates to your Active Directory namespace. After selecting your DNS namespace, you also must determine how to host your namespace on DNS servers.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the DNS namespace options.
- Select an appropriate namespace design.
- Describe the considerations for hosting namespaces.

DNS Namespace Scenarios

When you begin planning your DNS namespace, you must consider both the internal and external namespaces. The *internal namespace* is the namespace that the internal clients and servers use within your private network. The *external namespace* is the namespace by which the Internet references your organization. There is no requirement for you to implement the same DNS domain name internally that you have externally.

When you implement AD DS, you must use a DNS namespace for hosting Active Directory records.

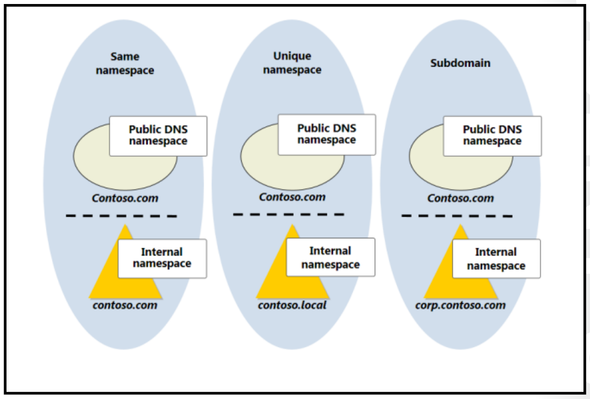

Note: Consider your options carefully before selecting a namespace design for AD DS. Although it is possible to change a namespace after implementing AD DS, this is a time-consuming and complex process.

You can choose one of the following scenarios to determine your environment's DNS namespace:

- The internal namespace matches the public namespace. In this scenario, the internal and public namespaces are the same, but contain different records. This provides simplicity, which is why it is often a suitable choice for smaller organizations. However, for larger networks, it can be difficult to manage.

Split DNS can provide a solution for larger network-management issues. In the split DNS configuration, your domain has two root-server zones that contain domain name registration information. Your internal network hosts are directed to one zone, while external hosts are directed to another for name resolution. If you use the same namespace for AD DS as for your external domain namespace, you must be careful to segregate the name servers for that namespace. External queries should be able only to resolve names such as www or ftp. External queries should not be able to resolve names such as HQDC01 or FILESERVER10. This requires that publicly accessible DNS servers host a zone for your domain that you maintain manually and that contains only the records that are appropriate for external resolution. You should point all systems within the domain to separate internal DNS servers that provide full resolution for all names in the domain.

- The internal namespace is different from the public namespace. In this scenario, the internal and public namespaces are different, with no link between them. This provides for separation in the namespace. In complex networks, with many Internet-facing applications, use of a different name introduces some clarity when configuring these applications.
- The internal namespace is a subdomain of the public namespace. In this scenario, you link the internal namespace to the public namespace, but there is no overlap between them. This provides a hybrid approach. The internal name is different, which allows for separation of the namespace. Additionally, the internal name also is related to the public name, which provides simplicity. This approach is the easiest to implement and manage. However, if you cannot use a subdomain of the public namespace for AD DS, you should use unique namespaces.

Choosing a Namespace Design

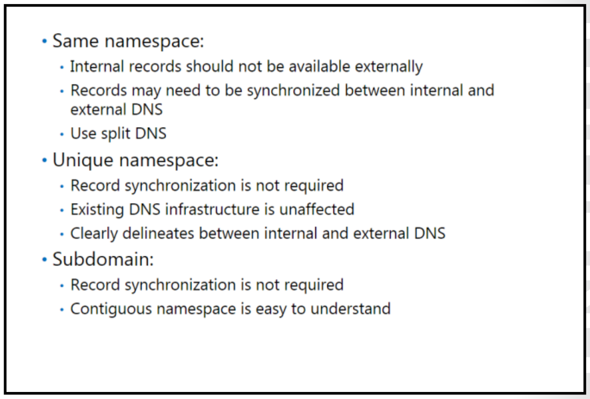

- Same namespace:
  - Internal records should not be available externally
  - Records may need to be synchronized between internal and external DNS
  - Use split DNS
- Unique namespace:
  - Record synchronization is not required
  - Existing DNS infrastructure is unaffected
  - Clearly delineates between internal and external DNS
- Subdomain:
  - Record synchronization is not required
  - Contiguous namespace is easy to understand

Using the same namespace internally and externally simplifies resource access from a user perspective. However, it increases complexity with respect to how you manage it. You should not make internal DNS records externally available, but some synchronization of records for external resources typically is necessary. For example, both your internal and external namespaces might use the name Contoso.com.

Using unique namespaces for the internal and public namespaces provides a clear delineation between internal and external DNS, and avoids the need to synchronize records between namespaces. However, having multiple namespaces may lead to user confusion. For example, you may choose the external namespace of Contoso.com and the internal namespace of Contoso.local. Note that where you implement a unique namespace configuration, you are no longer restricted to using registered domain names.

Using a subdomain of the public namespace for AD DS avoids the need to synchronize records between the internal and external DNS servers. Because the namespaces are linked, users typically find this structure easy to understand. For example, if your public namespace is Contoso.com, you might choose to implement your internal namespace as the subdomain AD or AD.Contoso.com.

The following table highlights some of the advantages and disadvantages of each namespace design method:

Type of Namespace Design	Advantages	Disadvantages
Same namespace	• Namespace already registered and owned • Easy for users; no confusion • Works well with Active Directory–integrated DNS	• Increases complexity • Requires careful management • Requires manual duplication of some records

Type of Namespace Design	Advantages	Disadvantages
Unique namespaces	- Avoids the need to synchronize records between namespaces - Provides a clear delineation between internal and external DNS - Segregates Active Directory service (SRV) resource records from publicly available records	- Can confuse users - Non-intuitive name
Subdomain of the namespace	- Avoids the need to synchronize records between namespaces	- Can confuse users - Can hamper the scalability of AD DS

Considering Split DNS

While having a matching internal and external DNS namespace can pose certain problems, split DNS can provide a solution to these problems. For example, in a nonsplit DNS configuration for the domain Contoso.com, you might have a DNS zone that looks like the example in the following table.

Host	Record type	IP address
www	host (A)	131.107.1.200
Mail	host (A)	131.107.1.201
Webserver1	host (A)	192.168.1.200
Exchange1	host (A)	192.168.0.201

When a client computer on the Internet wants to access the Simple Mail Transfer Protocol (SMTP) mail server by using the published name of mail.contoso.com, it queries the DNS server, which returns the result, 131.107.1.201. The client then establishes a connection over SMTP to that IP address.

However, the client computers on the corporate intranet also use the published name of mail.contoso.com. The DNS server returns the same result, which is a public IP address of 131.107.1.201. The client now attempts to establish a connection to the returned IP address by using the publishing computer's external interface. Depending upon the client configuration, this may not be successful.

You can avoid this problem by configuring two zones for the same domain name—one on each of the two DNS servers. The internal zone for adatum.com would now look like the example in the following table.

Host	Record type	IP address
www	alias (CNAME)	Webserver1.contoso.com
Mail	alias (CNAME)	Exchange1.contoso.com
Webserver1	host (A)	192.168.1.200
Exchange1	host (A)	192.168.0.201

The external zone for adatum.com would look like the example in the following table.

Host	Record type	IP address
www	host (A)	131.107.1.200
Mail	host (A)	131.107.1.201
Mail	mail exchanger (MX)	mail.contoso.com

Now client computers in the internal and external networks can resolve the name relay.contoso.com to the appropriate internal or external IP address.

In organizations that use Active Directory–integrated DNS zones, Internet users, and server functions outside the firewall must not use the internal Active Directory–integrated DNS servers to resolve any names. Instead, you must ensure that you confine these requests to an external non-Active Directory–integrated DNS server that resides on the perimeter network. This server is a primary zone server, and is authoritative for the same internal domain name. Therefore, no iterative queries are sent beyond this point. If the server does not find a name in this primary zone, the authoritative external DNS server declares the name invalid and not resolvable.

On the internal Active Directory–integrated DNS servers, for queries outside the firewall, Internet domain names are forwarded to the external DNS server in the perimeter network. You can make a firewall rule on the inside firewall that allows only the internal and external DNS servers to use User Datagram Protocol (UDP) port 53 packets between themselves. The firewall rule blocks all other UDP port 53 packets.

Considerations for Hosting Namespaces

Smaller organizations can use a single server to implement their complete DNS design for both the internal and external namespace. Although this design offers the benefit of simplicity, it poses security risks, including that internal records may be accessible from outside the organization.

To enhance DNS security, consider the split DNS design. This design features separate DNS servers that host the internal and external DNS records, which enhances security by preventing external users from contacting DNS servers with internal DNS records.

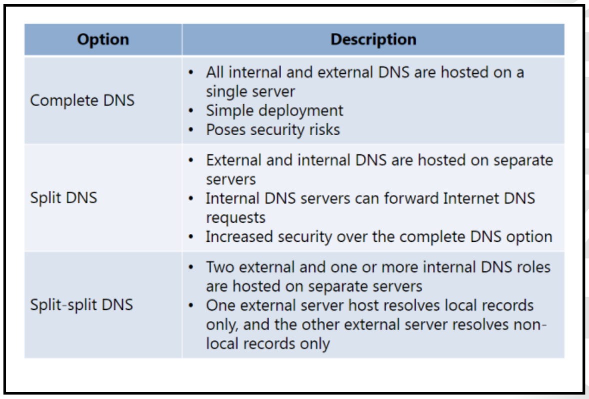

Option	Description
Complete DNS	• All internal and external DNS are hosted on a single server • Simple deployment • Poses security risks
Split DNS	• External and internal DNS are hosted on separate servers • Internal DNS servers can forward Internet DNS requests • Increased security over the complete DNS option
Split-split DNS	• Two external and one or more internal DNS roles are hosted on separate servers • One external server host resolves local records only, and the other external server resolves non-local records only

You can improve security further in your split DNS design by considering a split-split DNS. This is an enhancement to the split DNS structure because you have three DNS servers, two of which perform external name resolution. One external DNS server resolves the local names only, and the other external DNS server performs recursive Internet name resolution for internal DNS servers. You maintain the third DNS server for internal DNS queries. This can protect you from denial-of-service (DoS) attacks against your DNS infrastructure. DoS attacks against DNS servers include repeated and continuous echo requests (pings) from multiple computers. You can prevent this type of attack by turning off Internet Control Management Protocol (ICMP) at the firewall.

Lesson 3
Designing DNS Zones

A DNS zone hosts all or a portion of a domain and its subdomains. In addition to selecting DNS server locations and choosing an appropriate namespace design, you also must determine how you implement DNS zones to support these design choices. You must determine the zone types that you use, and where you store the zone data.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the different DNS zone types.
- Describe how Windows Server DNS servers store zone data.
- Select a DNS zone strategy.
- Describe NetBIOS name resolution considerations.
- Explain how to create zones in DNS.

Types of DNS Zones

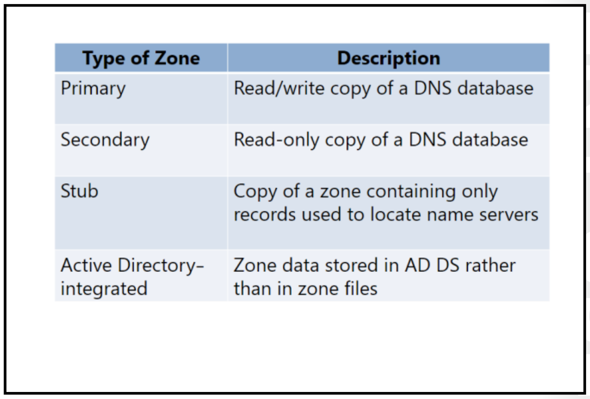

Type of Zone	Description
Primary	Read/write copy of a DNS database
Secondary	Read-only copy of a DNS database
Stub	Copy of a zone containing only records used to locate name servers
Active Directory-integrated	Zone data stored in AD DS rather than in zone files

You can replicate zone data to more than one server. This adds redundancy to a zone, because the information needed to find zone resources now exists on two servers. One reason to create zones is to ensure redundancy. If you have a zone that hosts critical server-resource records, it is likely that this zone has a higher level of redundancy than a zone that defines noncritical devices.

A DNS server maintains the zone data, and stores it in two ways:

- In a flat zone file that contains mapping lists
- Integrated into AD DS

A DNS server is authoritative for a zone if it hosts the resource records for the names and IP addresses that the clients request in the zone file.

The four DNS zone types are:

- Primary
- Secondary
- Stub
- Active Directory–integrated

Primary Zone

When a DNS server hosts a primary zone, the DNS server is the primary source for information about this zone. The DNS server stores the master copy of zone data in a local file or in AD DS. When the DNS server stores the zone in a file, it names the primary zone file zone_name.dns, by default, and stores it on the server in the %windir%\System32\Dns folder. When you do not store the zone in AD DS, this is the only DNS server that has a writable copy of the database.

Secondary Zone

When a DNS server hosts a secondary zone, the DNS server is a secondary source for the zone information. The DNS server must obtain the zone data from another remote DNS server that also hosts the zone. The DNS server that hosts the secondary zone must have network access to the remote DNS server to receive updated zone information. A secondary zone is a copy of a primary zone that another server hosts, so AD DS cannot store it. Secondary zones can be useful if you must replicate zone data to or from non-Windows DNS zones.

Stub Zone

A *stub zone* is a replicated copy of a zone that contains only those resource records necessary to identify that zone's authoritative DNS servers.

Note: Windows Server 2003 introduced stub zones, which solved several problems with large DNS namespaces and multiple tree forests. (A *multiple tree forest* is an Active Directory forest that contains two different domain names.)

A stub zone resolves names between separate DNS namespaces. This might be necessary when a corporate merger requires that the DNS servers for two separate DNS namespaces resolve names for clients in both namespaces.

Confusion can arise about when to use conditional forwarders rather than stub zones, because both DNS features allow a DNS server to respond to a query by referring or forwarding the query to a different DNS server. However, these settings have different purposes:

- A conditional forwarder setting configures the DNS server to forward a query that it receives to a DNS server, depending on the DNS name that the query contains.
- A stub zone ensures that the DNS server that is hosting a parent zone is aware of all the DNS servers that are authoritative for a child zone.

Use stub zones when you want a DNS server that is hosting a parent zone to remain aware of the authoritative DNS servers for one of its child zones. If the same DNS server that hosts the parent zone also hosts a child zone's stub zone, the DNS server will receive a list of all new authoritative DNS servers for the child zone when it requests an update from the stub zone's master server. This method of updating the DNS server that is hosting the parent zone ensures that the DNS server maintains a current list of the child zone's authoritative DNS servers.

Active Directory-Integrated Zone

If AD DS stores the zone, DNS can utilize the multimaster replication model to replicate the primary zone. This enables you to edit zone data on more than one DNS server. You can replicate Active Directory–integrated zone data to domain controllers even if you do not install the DNS role on the domain controller.

A DNS server can store zone data in the Active Directory database if the DNS server is a domain controller. When the DNS server stores zone data in this way, it stores the records in the zone file as Active Directory objects, and the various properties of these objects are Active Directory attributes. All domain controllers hosting the DNS zone in the Active Directory database are primary zone servers for the zone, and can accept changes to the DNS zone, and then replicate those changes out to all other domain controllers. Because it uses Active Directory replication, each change is sent securely via encrypted replication traffic. If a domain controller with an Active Directory–integrated DNS zone fails, DNS functionality for that zone and the domain continues to operate correctly as long as there are other domain controllers with the Active Directory–integrated zone. You can configure a non-domain controller with the DNS server role as a secondary zone server, and in this case, all domain controllers with the DNS server role assigned act as the master server and conduct zone transfers to the non-domain controller secondary server.

Discussion: Designing a DNS Zone Strategy

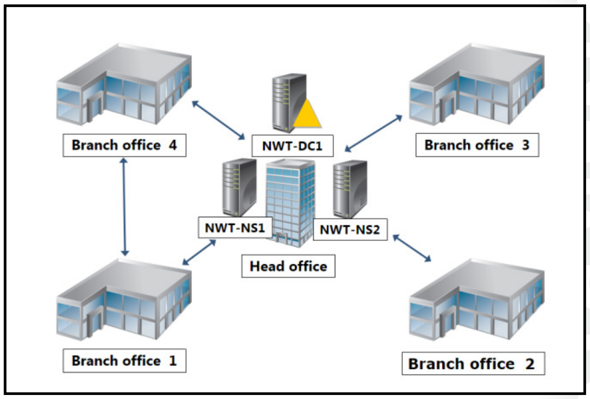

The Northwind Traders company has five locations, as shown on the slide. The Northwind Traders network:

- Implements AD DS within the network.
- Uses the name northwindtraders.local for the DNS domain name for the Active Directory forest root.
- Implements all branches as part of the same domain.
- Does not use Active Directory–integrated zones.
- Has two name servers in the head office: NWT-NS1, and NWT-NS2.
- Stores the northwindtraders.local primary zone on NWT-NS1. However, NWT-NS2 hosts a secondary zone for northwindtraders.local.
- Has no name servers elsewhere.
- Has a head office that supports more than 2,000 users.
- Has two branches, each of which support approximately 25 users.
- Has two other branches that support several hundred users each.

Consider the current configuration, and then use the information contained within the preceding topics to determine how you might design DNS. To help with your design, consider the following questions:

Question: How would you modify the DNS design for this scenario?

Question: Where would you place additional name servers, if any?

Question: What DNS server roles would you propose deploying?

Question: Assuming all Internet connectivity is through the head office, how would you design forwarding?

Question: How would you design the DNS zones?

Question: Are Active Directory–integrated zones indicated?

Question: How would you design zone transfers?

Question: Contoso, Ltd just acquired Northwind Traders. Does this affect your DNS design decisions?

Considerations for the GlobalNames Zone

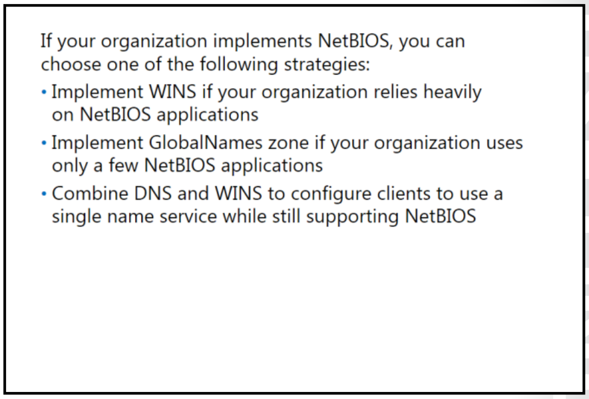

If your organization implements NetBIOS, you can choose one of the following strategies:

- Implement WINS if your organization relies heavily on NetBIOS applications
- Implement GlobalNames zone if your organization uses only a few NetBIOS applications
- Combine DNS and WINS to configure clients to use a single name service while still supporting NetBIOS

NetBIOS is a session-management protocol that earlier versions of the Microsoft network operating systems use. The network protocol, NETBios Extended User Interface (NETBeui), was the primary communications protocol for the Microsoft MS-DOS® and Windows 3.x series operating systems. Most modern applications do not rely on NetBIOS for establishing and maintaining sessions between computers, although both Windows 8 and Windows Server 2012 continue to provide support for NetBIOS. However, organizations that do rely on older, custom line-of-business (LOB) applications still may require support for NetBIOS. For these reasons, it is important that you understand when and how to provide NetBIOS name resolution as part of your design for name resolution services.

A NetBIOS name represents a single computer or a group of computers. Applications use a NetBIOS name that is up to 16 characters in length to identify the NetBIOS resource that the local network hosts.

Note: NetBIOS reserves the first 15 characters for a specific computer's name. The sixteenth character identifies a resource or service on that computer. LON-SVR2[20h] is an example of a NetBIOS name, and it represents the Server service on the LON-SVR2 computer.

Computers that support NetBIOS applications must register their NetBIOS names to avoid conflicting with other NetBIOS-based computers, and where necessary, resolve other NetBIOS names into IP addresses for lower-level communications.

WINS Considerations

Windows Internet Naming Service (WINS) provides a centralized database for registering dynamic mappings of a network's NetBIOS names. Windows Server 2012 retains WINS support to provide backward compatibility.

Note: The WINS Server feature provides WINS support.

You also can resolve NetBIOS names by using:

- Broadcast messages. These do not work well on large networks because routers typically do not propagate broadcasts.
- Lmhosts file on all computers. This is a high-maintenance solution because you must maintain the file manually on all computers.

WINS resolves NetBIOS names to IP addresses. This can reduce NetBIOS broadcast traffic and enable clients to resolve the NetBIOS names of computers that are on different subnets.

GlobalNames Zone Considerations

The GlobalNames zone provides single-label name resolution for large enterprise networks that do not deploy WINS. Some networks may require the ability to resolve static, global records with single-label names, which WINS currently provides. These single-label names refer to well-known and widely used servers with statically assigned IP addresses. You create a GlobalNames zone manually, and it is not available for dynamically registering records. GlobalNames zone helps customers migrate to DNS for all name resolution. The DNS Server role in Windows Server 2012 supports the GlobalNames zone feature.

GlobalNames zone assists in the migration from WINS. However, it is not a replacement for WINS. GlobalNames zone is not intended to support the single-label name resolution of records that are registered in WINS dynamically, or those that IT administrators typically do not manage. Support for these dynamically registered records is not scalable, particularly for larger customers with multiple domains and forests.

We recommend a GlobalNames zone deployment that uses an Active Directory–integrated zone named GlobalNames, which is distributed globally.

Whether to Use WINS or GlobalNames Zone

The method that you select to provide for NetBIOS name resolution depends largely on the number of NetBIOS clients that you have, and the consequent volume of NetBIOS name registrations, releases, and queries that you receive.

If your organization implements NetBIOS, you can use one of the following strategies:

- Implement WINS. If your organization relies heavily on NetBIOS applications, continue to use WINS.
- Implement GlobalNames zone. If your organization uses only a few NetBIOS applications, or when you have decommissioned most of your NetBIOS applications, use GlobalNames zone to manage static, single-label names.
- Combine DNS and WINS. Instead of using GlobalNames zone, you can configure DNS and WINS integration by configuring the DNS zone properties to perform WINS lookups for NetBIOS-compliant names. The advantage of this approach is that you can configure client computers to use only a single name service (DNS) and still be able to resolve NetBIOS-compliant names.

Demonstration: Creating DNS Zones

This demonstration shows you how to:

- Create a primary reverse lookup zone.
- Create a secondary forward lookup zone.

Demonstration Steps

Create a primary reverse lookup zone

1. On LON-DC1, open DNS.
2. Create a new Reverse Lookup Active Directory–integrated zone with the following properties:

- Type: **Primary**
- Name: **IPv4 Reverse Lookup Zone**
- Network ID: **172.16**
- Allow only secure dynamic updates

3. In Windows PowerShell, type the following cmdlet, and then press Enter:

```
Register-DnsClient
```

4. In DNS Manager, refresh the Reverse Lookup Zone that you just created. You should now see a pointer (PTR) resource record for Lon-dc1.adatum.com.

Create a secondary forward lookup zone

1. Switch to LON-SVR1.

2. Open DNS.

3. Create a new Forward Lookup zone with the following properties:

- Zone type: **Secondary**
- Zone name: **Adatum.com**
- Master server: **172.16.0.10**

Note: Zone transfers are shown in the next demonstration.

Lesson 4
Designing DNS Zone Replication and Delegation

You can synchronize zone information between multiple DNS servers by using zone transfer or zone replication. Which method you use depends on whether the zone is Active Directory–integrated. To design zone replication, you must consider when to use secondary zones, and how you want zone replication to occur. Finally, you should consider whether delegation is useful for your design.

Lesson Objectives

After completing this lesson, you will be able to:

- Determine when it is appropriate to implement secondary zones.
- Describe zone transfers and zone-data replication.
- Mitigate against security risks when planning how to perform zone transfers.
- Describe how to integrate namespaces.
- Explain how to configure zone transfers.

When to Implement Secondary Zones

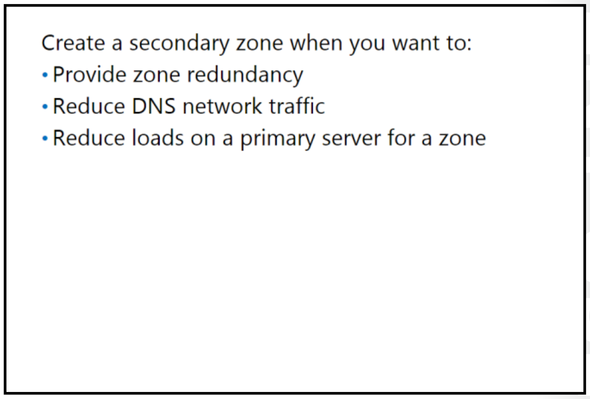

Create a secondary zone when you want to:

- Provide zone redundancy
- Reduce DNS network traffic
- Reduce loads on a primary server for a zone

Adding DNS servers provides zone redundancy, which enables DNS name resolution for clients, even if a primary server for the zone stops responding. The more servers that you have that are authoritative for a particular zone, the less likely it is that queries will go unanswered for that zone's resources.

If you place servers close to large client populations or isolated networks, you can reduce the amount of query traffic that has to flow across potentially costly and slow WAN links.

You can use additional secondary servers to reduce loads on a zone's primary server. For example, you can direct clients to secondary servers that service queries from local clients only, but not from clients from across the entire enterprise.

Note: You typically implement secondary servers only in zones that are not Active Directory–integrated. If your zone is Active Directory–integrated, zone transfers occur automatically as part of Active Directory replication between all domain controllers, depending upon how you configure them.

When you use a traditional DNS zone, the DNS server stores the zone data in a text file on the disk. In this scenario, you use a single primary zone with multiple secondary zones. By using secondary zones, you can synchronize zones from non-Windows primary zones. Additionally, you can implement disk storage on any server that is running the Windows Server operating system.

Active Directory–integrated zones store DNS information in AD DS. Zone information then replicates automatically to all domain controllers, and any domain controller can update this information. Any domain controller that has DNS installed begins servicing Active Directory–integrated zones automatically.

Active Directory–integrated zones behave as primary zones to traditional secondary zones. However, unlike traditional primary zones, there can be multiple Active Directory–integrated zones.

What Are Zone Transfers and Replication?

To provide availability and fault tolerance when resolving name queries, zones must be available from more than one DNS server on the network. Zone transfers and zone replication help accomplish this. Zone *transfers* occur in a traditional DNS zone. Zone *replication* occurs in an Active Directory–integrated zone. A zone's secondary server initiates zone transfers, which it then sends to the master that you configure.

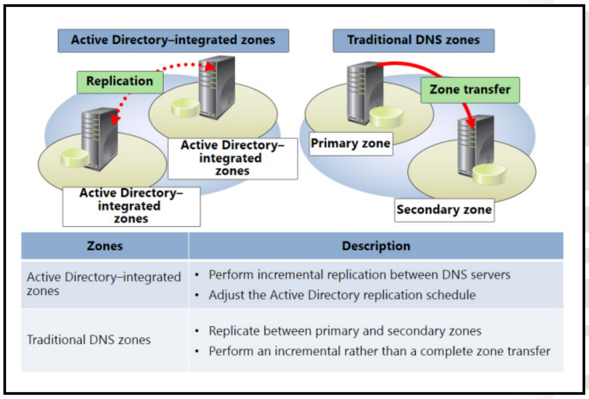
Active Directory–integrated zones
Traditional DNS zones
Replication
Zone transfer
Active Directory–integrated zones
Primary zone
Active Directory–integrated zones
Secondary zone

Zones	Description
Active Directory–integrated zones	• Perform incremental replication between DNS servers • Adjust the Active Directory replication schedule
Traditional DNS zones	• Replicate between primary and secondary zones • Perform an incremental rather than a complete zone transfer

Zone Transfers

A zone transfer occurs when you replicate the DNS zone that is on one server to another DNS server.

Zone transfers synchronize primary and secondary DNS server zones. This is how DNS builds its resilience on the Internet. DNS zones must remain updated on primary and secondary servers. Discrepancies in primary and secondary zones can cause service outages and host names that resolve incorrectly.

Zone transfers can occur in one of three ways:

- Full zone transfer. A full zone transfer occurs when you copy the entire zone from one DNS server to another. A full zone transfer is also called an *All Zone Transfer* (or *AXFR)*.
- Incremental zone transfer. An incremental zone transfer (or *IXFR*) occurs when there is an update to the DNS server and only the resource records that were changed replicate to the other server.
- Fast transfer. Windows DNS servers also perform *fast transfers*, which is a type of zone transfer that uses compression and sends multiple resource records in each transmission.

Not all DNS server implementations support incremental and fast zone transfers. When integrating a Windows Server 2012 DNS server with a BIND DNS server, you must ensure that the BIND version that is installed supports the features that you need. BIND servers are common on UNIX-based networks. You may encounter BIND servers when setting zone transfers with your ISP.

Note: BIND 8.1 was released in 1998. Therefore, most users of this DNS implementation will have updated to the free version 9, which was released in 2000.

DNS Notify

A master server uses DNS notify to alert its configured secondary servers that zone updates are available. The secondary servers then petition their master server to obtain the updates. DNS notify is an update to the original DNS protocol specification that permits notification to secondary servers when zone changes occur. This is useful in a time-sensitive environment, where data accuracy is important.

Active Directory-Integrated Zones and Active Directory Replication

Active Directory–integrated zones replicate by using multimaster Active Directory replication instead of the zone transfer process. This means that any standard domain controller that also holds the DNS role can update the DNS zone information. This information then replicates to all DNS servers that host the DNS zone. DNS record changes are treated as Active Directory–replicable values. The process to update A DNS record change does not differ from any other Active Directory replicable event.

Planning Security for Zone Transfers

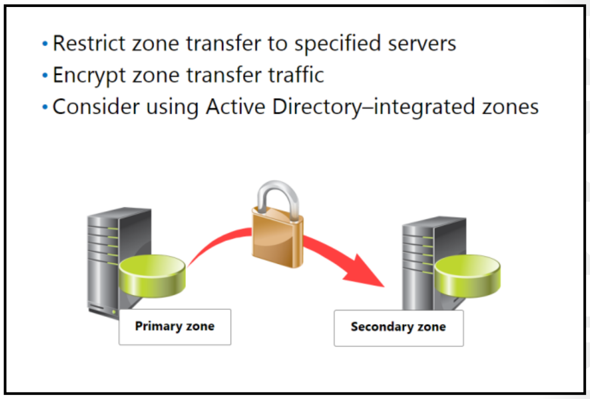

- Restrict zone transfer to specified servers
- Encrypt zone transfer traffic
- Consider using Active Directory–integrated zones

Zone information provides organizational data, so you should take precautions to ensure that you help to protect it from access by hackers. Ensure that the zone data cannot be overwritten with bad data, a process that is called *DNS poisoning*.

You can specify the list of allowed DNS servers by right-clicking a zone name on a zone's Zone Transfers tab in the Zone Properties dialog box, and then clicking Properties. You also can use these options to disallow zone transfer. By default, zone transfers are turned off. You can specify each zone transfer's secondary server by IP Address. Conversely, to allow zone transfers to any identified name server in the zone, you can use the DNS Name Servers radio button on the Zone Transfer tab.

Although these options provide security by limiting the data recipients, they do not secure that data during transmission. If the zone data is confidential, you should use an IPsec policy to help secure the transmission or replicate the zone data over a virtual private network (VPN) tunnel. This prevents packet sniffing of data in the data transmission.

Using Active Directory–integrated zones replicates the zone data as part of normal AD DS replications. The Active Directory replication process helps secure the zone transfer by using the Kerberos authentication protocol to encrypt domain controller communications. Using Active Directory–integrated zones is a security best practice.

In the zone properties sheet on the Zone Transfer tab there is a radio button that enables you to establish zone transfers to any server. However, it is unsecure and allows any entity on the Internet to copy all of your zone resource records. You would normally only use this in testing and debugging environments.

Integrating Namespaces

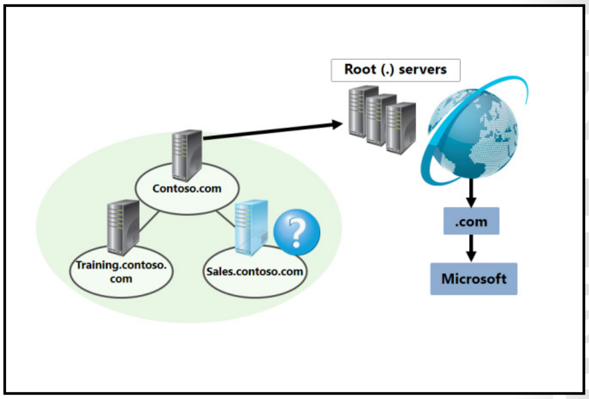

When integrating namespaces, you must consider how to configure root hints and forwarding behavior in your DNS infrastructure.

Root hints

Root hints are the list of Internet servers that your DNS server uses if it cannot resolve a DNS query by using either a DNS forwarder or its own cache. The root hints list is made up of the servers highest in the DNS hierarchy. The root hint servers can provide the necessary information for a DNS server to perform an iterative query to the next lowest layer of the DNS namespace.

When you install the DNS role, the list of root hints installs automatically. The DNS role copies them from the cache.dns file that the DNS role setup files include. You also can add root hints to a DNS server to support lookups for noncontiguous domains within a forest.

When a DNS server communicates with a root hints server, it uses only an iterative query. If you select the Disable recursion option, the server is not able to perform queries on the root hints. If you configure the server by using a forwarder, it attempts to send a recursive query to its forwarding server. If the forwarding server does not answer this query, the server responds that the host could not be found.

Note: Recursion on a DNS server and recursive queries are not the same. *Recursion on a server* means that the server will use its root hints to try to resolve a DNS query.

Forwarding

A *forwarder* is a network DNS server that forwards DNS queries for external DNS names to DNS servers outside that network. You also can use conditional forwarders to forward queries according to specific domain names.

A network DNS server is designated a forwarder when other DNS servers on the network forward to it the queries that they cannot resolve locally. By using a forwarder, you can manage name resolution for names outside your network (such as names on the Internet), and improve the efficiency of name resolution for your network's computers.

The server that is forwarding requests on the network must be able to communicate with the DNS server on the Internet. This means that you must configure it to forward requests to another DNS server, or it will use root hints to communicate.

Use a central forwarding DNS server for Internet name resolution. This can improve performance, simplify troubleshooting, and is a security best practice. You can isolate the forwarding DNS server on a perimeter network, which ensures that no server within the network is communicating directly to the Internet.

Considerations for Conditional Forwarding When Integrating Namespaces

DNS clients on separate networks can resolve the names of the other DNS clients without having to query Internet DNS servers. This can be beneficial when a company merger occurs. To enable this capability, you must configure each network's DNS servers to forward queries for names in the other network. DNS servers in one network will forward names for clients in the other network to a specific DNS server, which will build a large information cache about the other network. When forwarding in this manner, you create a direct point of contact between the DNS servers of two networks. This reduces the need for recursion.

Stub zones do not provide the same server-to-server benefit. This is because a DNS server that is hosting a stub zone in one network replies to queries for names in the other network. The reply is a list of all authoritative DNS servers for the zone with that name, rather than the specific DNS servers that you designated to handle this traffic. This configuration complicates any security settings that you want to establish between specific DNS servers that are running in each of the networks.

Considerations for Stub Zones When Integrating Namespaces

Use stub zones when you want a DNS server that is hosting a parent zone to remain aware of the authoritative DNS servers for one of its child zones. If the same DNS server that hosts the parent zone also hosts a child zone's stub zone, the DNS server will receive a list of all new authoritative DNS servers for the child zone when it requests an update from the stub zone's master server. This method of updating the DNS server that is hosting the parent zone maintains a current list of the authoritative DNS servers for the child zone.

A conditional forwarder is not an efficient way to keep a DNS server that is hosting a parent zone aware of the authoritative DNS servers for a child zone. This is because whenever the authoritative DNS servers for the child zone change, you would have to configure the conditional forwarder setting manually on the DNS server that hosts the parent zone with the IP address for each new authoritative DNS server for the child zone.

Zone Delegation

DNS is a hierarchical system, and zone delegation connects the DNS layers together. A zone delegation points to the next lower hierarchical level, and identifies the name servers that are responsible for the lower-level domain.

To decide whether to divide the DNS namespace and make additional zones, consider whether you need to do any of the following:

- Delegate management of part of the DNS namespace to another organizational location or department. In this case, implement multiple zones and delegate responsibility for the child zones to the relevant departmental administrators.
- Distribute traffic loads among multiple servers by dividing one large zone into smaller zones. This improves DNS name resolution performance and creates a more fault-tolerant DNS environment.
- Extend the namespace by adding numerous subdomains simultaneously. This accommodates the opening of a new branch or site.

For example, you can implement the DNS domain training.Contoso.com as a new zone with a matching name. This new zone would have its own name servers and administrators, together with the relevant records in the parent zone that identify the authoritative name servers. Alternatively, the Contoso.com administrator can create a subdomain record named Training in the Contoso.com zone. In this instance, no name servers exist for the child domain, and no delegation exists.

Demonstration: Configuring Zone Transfers

This demonstration shows you how to:

- Enable zone transfers on a zone.
- Perform a zone transfer.

Demonstration Steps

Enable zone transfers on a zone

1. Switch to LON-DC1, and then switch to the DNS console.
2. Modify the properties of the **Adatum.com** zone.
3. Allow zone transfers to servers that are on the **Name Servers** tab.
4. Configure the LON-SVR1 server on the **Notify** list.
5. Add the **LON-SVR1.Adatum.com** server as a named server on the **Name Servers** tab.

Perform a zone transfer

1. Switch to LON-SVR1, and then switch to the DNS console.
2. Refresh the Adatum.com zone.
3. Switch back to LON-DC1, and in the DNS console, add a new **New Alias (CNAME)** resource record.
4. Switch back to LON-SVR1, and force a zone transfer to view the new record.

Lesson 5
Optimizing DNS Servers

When DNS performs optimally, it can help improve application performance within your network. This, in turn, improves the end-user experience. When you are designing your name resolution infrastructure, you must choose the appropriate DNS configurations for your organization's particular needs.

Lesson Objectives

After completing this lesson, you will be able to:

- Optimize DNS recursion.
- Optimize DNS root hints.
- Optimize DNS server functionality.
- Optimize Active Directory–integrated zones.

Optimizing DNS Recursion

When a DNS client (also known as a *DNS resolver*) sends a request for information, its configured DNS server uses one of two query types: an iterative query or a recursive query.

Disable recursion to limit name resolution to a specific server or as a failover for another DNS server:

- Benefit: You can reduce the load on the DNS server
- Consequence: You will not be able to resolve names outside of your own zone

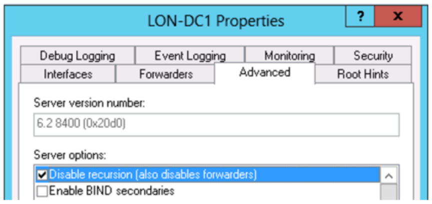

Server version number:
6.2 8400 (0x20d0)

Server options:
☑ Disable recursion (also disables forwarders)
□ Enable BIND secondaries

Iterative Query

When a DNS server receives an iterative query, it responds with the answer to the query or a referral to another DNS server that is authoritative for the queried record. All DNS servers initially support iterative queries, and they use these iterative queries to resolve names elsewhere in the Internet DNS namespace.

Recursive Query

In a recursive DNS query, the petitioned DNS server resolves a nonlocal hostname on behalf of a DNS client. When a client uses a recursive query, it expects either an answer or an error that indicates that the server cannot resolve the query. The DNS server cannot respond with a referral to another DNS server that may be authoritative for the record for which you are querying.

You do not need to configure DNS servers to support recursion because client computers typically use recursive queries of their configured DNS servers.

When you disable recursion, a DNS server does not use root hints or forwarders to resolve queries for clients. You should disable recursion on all DNS servers that do not require this functionality. Internal DNS servers typically require that you enable recursion, whereas a DNS server that is hosting an external DNS namespace should have recursion disabled. This prevents Internet clients from using that server to resolve DNS names. Note that when using split DNS and Active Directory–integrated DNS, you would expect all nonauthoritative requests to the Active Directory–integrated DNS servers to forward to the external DNS server for resolution. In this scenario, you should not disable recursion.

By preventing Internet clients from performing recursive queries on your external DNS server, you can reduce your server load and help prevent denial-of service attacks.

Note: You can disable recursion and forwarders by selecting the Disable recursion check box on the Advanced tab in the DNS server's Properties dialog box.

Optimizing DNS Root Hints

When a DNS server receives a recursive record query from a DNS domain for which it is not authoritative, the server must have a way to resolve that query. DNS servers can petition the appropriate DNS server for a given subdomain by configuring each domain in the DNS namespace with information that identifies subdomains and the authoritative servers for those subdomains.

- Delete root hints on servers that do not need to communicate with DNS servers that are authoritative for the root domain
- Modify root hints if the root domain is internal
- Update root hints when DNS servers that are authoritative for the root domain change

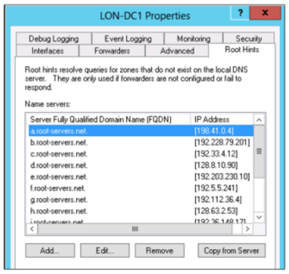

Root hints resolve queries for zones that do not exist on the local DNS server. They are only used if forwarders are not configured or fail to respond.

Name servers:

Server Fully Qualified Domain Name (FQDN)	IP Address
a.root-servers.net.	[198.41.0.4]
b.root-servers.net.	[192.228.79.201]
c.root-servers.net.	[192.33.4.12]
d.root-servers.net.	[128.8.10.90]
e.root-servers.net.	[192.203.230.10]
f.root-servers.net.	[192.5.5.241]
g.root-servers.net.	[192.112.36.4]
h.root-servers.net.	[128.63.2.53]
i.root-servers.net.	[192.36.148.17]

If you delete the root hints file, you should configure servers to forward requests to another server that has a root hints file.

However, when the query is for a domain that is elsewhere in the DNS namespace, the petitioned DNS server must determine which DNS server is authoritative for the queried domain. However, it first must look up the DNS tree to its root. The root servers contain information about subdomains, such as .com, .edu, and .uk. Elements of the DNS infrastructure, specifically DNS servers authoritative for their own zones, can provide this information to the petitioning DNS server to enable it to continue searching the namespace for the authoritative name servers.

Root hints are a list of preliminary resource records that the petitioned DNS service can use to locate other DNS servers that are authoritative for the root of the DNS domain namespace tree. Consequently, DNS servers that you configure with root hints can locate and query the root servers quickly to expedite their query.

If you delete the root-hints file from a DNS server, you remove that server's ability to contact directly a server that is authoritative for the root of the DNS infrastructure. In this case, you should configure servers to forward requests to another server that has a root hints file. This controls the path that your organization uses for Internet DNS lookups.

Note: The root hints file Cache.dns holds the root hints IP addresses. It is located in the %SystemRoot%\System32\Dns folder.

If the internal network does not connect to the Internet, you may need to create a root domain to use internally for name resolution. On servers that are authoritative for the root domain, you can remove the root hints information safely because these servers do not use the root hints file. You then should remove the default resource records on your organization's other servers and replace them with your organization's resource records.

Note: You can configure root hints on the Root Hints tab in the DNS server's Properties dialog box.

Optimizing DNS Server Functionality

You can help to optimize your DNS servers by carefully considering how you implement DNS zones. (This applies particularly to zone transfers). Additionally, you should determine how to best use caching-only DNS servers.

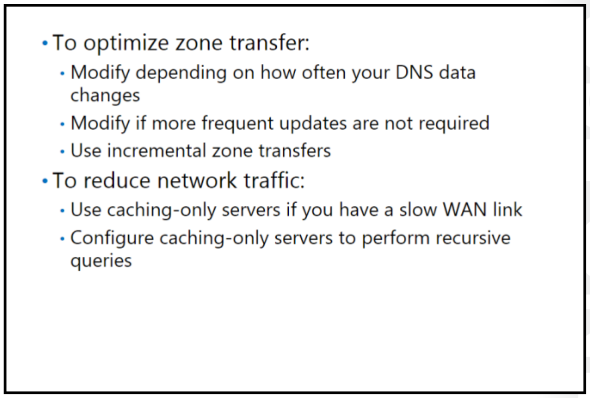

- To optimize zone transfer:
  - Modify depending on how often your DNS data changes
  - Modify if more frequent updates are not required
  - Use incremental zone transfers
- To reduce network traffic:
  - Use caching-only servers if you have a slow WAN link
  - Configure caching-only servers to perform recursive queries

Considerations for Zone Transfers

The Windows Server 2012 DNS Server role and other recent non-Microsoft implementations of DNS use incremental zone transfers, which include only changes to DNS zones, rather than a complete zone transfer. In many cases, this means that zone transfers consume very little network capacity and do not require optimization.

However, to optimize zone transfers, you can control how often zone transfers occur, how quickly the DNS server reattempts a zone transfer after a previous failed attempt, and how quickly zone data expires when it is not refreshed. By default, zone transfers occur every 15 minutes, and if a zone transfer fails, a retry occurs after 10 minutes.

If you are using Active Directory–integrated DNS, any change to a record is a replicable event, and the domain controller that writes the record change notifies its Active Directory partners that it has a replicable event for processing. Normal domain-controller replication performed at the attribute level will occur. (Note that this bypasses the zone transfer process completely.) However, it is possible to have secondary zone servers even in an Active Directory–integrated environment. These secondary servers would not be domain controllers, and the normal zone-transfer process would occur, but only to those secondary servers.

Note: By default, if the DNS server cannot contact the master server with the primary zone, the data in a secondary zone expires after 24 hours.

Considerations for Using Caching-Only Servers

Caching-only servers perform name resolution for clients, and then they cache—or *store*—the results. This type of server is not authoritative for a zone. Therefore, it does not store standard primary or standard secondary zones. Consequently, a caching-only server does not participate in zone transfers, which reduces network traffic.

The cache is populated with the most frequently requested names. Using a caching-only server is a simple way to provide some local DNS resolution capabilities without configuring zones.

When considering whether to use caching-only servers, remember that although the server does not participate in zone transfers, it must perform queries for any DNS record that is not cached. This generates network load. Therefore, it is important that you balance the benefits of avoiding DNS zone-transfer traffic against the increase in DNS queries.

Optimizing Active Directory-Integrated Zones

If you implement AD DS, you can choose to implement Active Directory–integrated zones.

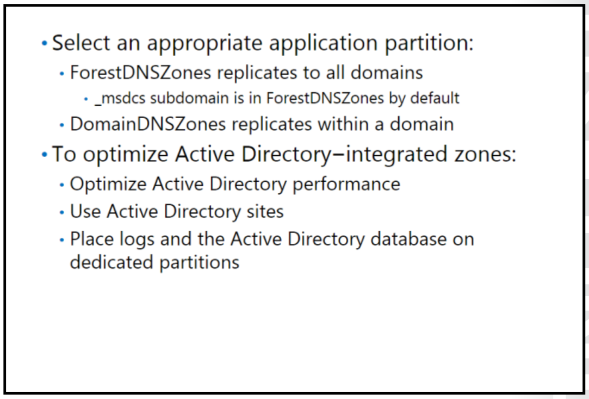

- Select an appropriate application partition:
  - ForestDNSZones replicates to all domains
    - _msdcs subdomain is in ForestDNSZones by default
  - DomainDNSZones replicates within a domain
- To optimize Active Directory–integrated zones:
  - Optimize Active Directory performance
  - Use Active Directory sites
  - Place logs and the Active Directory database on dedicated partitions

The following specialized application partitions in AD DS store these Active Directory–integrated zones:

- The ForestDNSZones partition replicates to all domain controllers in the forest.
- The DomainDNSZones partition in each domain replicates to all domain controllers that are running the DNS Server role within the domain.

Note: The _msdcs subdomain is in ForestDNSZones by default, because it contains records for all Active Directory domains. In most cases, other zones replicate only within the local domain.

When configuring your Active Directory–integrated zones, select the appropriate partition. Only use the ForestDNSZones partition when it is essential for all DNS servers within the entire forest to have copies of all other DNS zones.

Aside from selecting the appropriate partition, you also must ensure that AD DS is optimized. For example, by using Active Directory sites, you can control replication between physical locations, including the replication of DNS zone data.

Additionally, at the individual server level, place the Active Directory database and logs (which include the DNS partitions) on dedicated physical disk partitions to increase AD DS performance for queries and changes.

Lesson 6
Designing DNS for High Availability and Security

Given the importance of DNS, you must design your name resolution infrastructure to help ensure name resolution services are available to all parts of your network, even if part of your DNS name resolution infrastructure fails. You also must have a thorough understanding of the security risks that your name resolution infrastructure faces, and design your DNS implementation to help mitigate those risks.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the best practices for making DNS highly available.
- Describe the common DNS security risks.
- Select a DNS security strategy.
- Select additional security settings.
- Describe how DNSSEC works.
- Determine a suitable DNS security model.

Best Practices for Making DNS Highly Available

To optimize your DNS name resolution services for high availability, consider how clients would resolve names if various network infrastructure elements fail, including the name servers. For example, in a branch office, how would name resolution occur if a link to the head office were unavailable? For Active Directory clients, failure of DNS means they cannot reach a domain controller. This is because the Service Locator records enumerate the domain controllers for any particular Active Directory domain. If a client cannot reach a DNS server that can provide a domain controller address for the requested domain, Active Directory authentication fails. You can further optimize DNS on the client by ensuring that the Preferred DNS server is a local DNS server, such as one in the branch office, and the alternate is set to a remote DNS server, such as one at the head office.

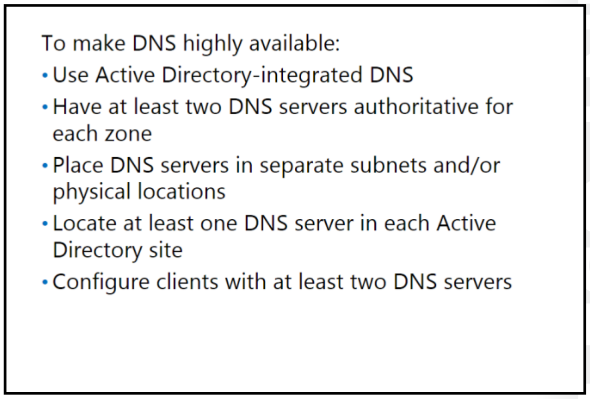

To make DNS highly available:

- Use Active Directory-integrated DNS
- Have at least two DNS servers authoritative for each zone
- Place DNS servers in separate subnets and/or physical locations
- Locate at least one DNS server in each Active Directory site
- Configure clients with at least two DNS servers

If a query to the Preferred DNS server fails to reach that server because the server is down, not responding, or is having network issues, the client attempts to send the query to the next or Alternate DNS server. This process repeats until either a DNS server responds to the query, or the list of available DNS servers on the client is exhausted. At this point, all name resolution beyond what is in the resolver cache fails. If the client successfully establishes a connection to an alternate or next DNS server, it will continue to use the responding DNS server for 15 minutes. After this, it will attempt to reach the Preferred DNS server again. If the Preferred DNS server still does not respond, the process repeats. If the Preferred DNS server responds but cannot find a name resolution, the client does not use the alternate or next DNS servers. The client resolver will list that name as not found—that is, the name does not exist.

When planning DNS availability, consider the following guidelines:

- Use Active Directory–integrated DNS. Because it is an AD DS best practice to have two or more domain controllers per domain or site, Active Directory-integrated DNS provides fault tolerance for domain authentication and DNS. As a result of Active Directory–integrated DNS using Active Directory replication, changes replicate quickly and securely.

- Have at least two DNS authoritative servers for each DNS zone. This ensures that if one DNS server is unavailable, the remaining DNS server can continue to service the DNS zone. Depending on your design, more than two DNS servers may be desirable.

- Place at least one DNS server in each separate physical location, depending upon the specific configuration. By placing a DNS server in each separate physical location, you can be sure that a WAN link failure does not affect name resolution. Placing DNS servers in separate subnets also reduces the possibility that routing problems will affect name resolution.

- Place at least one DNS server in each Active Directory site. Placing a DNS server at each Active Directory site helps to ensure that a WAN link failure does not affect AD DS. In most cases, you should configure two DNS servers for each Active Directory site. If the site resides in the same physical location and is not separated by WAN links, then you would not have to provide a DNS server for each site.

- Configure clients with at least two DNS servers. To be fault tolerant, you must configure clients with at least two DNS servers. If clients cannot contact the first DNS server, they contact the second DNS server automatically. You can set more than two DNS servers. You might do this when there is an issue with network connectivity, such as intermittent dropouts or limited access.

Common DNS Security Attacks

Your DNS infrastructure is prone to attack due to both the nature of the data that it contains and the importance of the name resolution process. The data that a DNS zone contains is attractive to hackers because they can capture the zone data and then use it for malicious purposes. For example, knowing the name and IP address of a given server can potentially identify that server's function, and it provides an IP address to which to send attack or denial-of-service packets. Additionally, because of the critical nature of name resolution services, if attackers can disrupt DNS, other important applications such as Active Directory authentication, may fail.

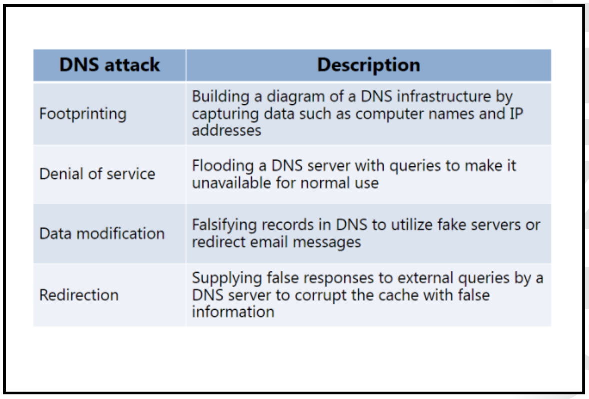

DNS attack	Description
Footprinting	Building a diagram of a DNS infrastructure by capturing data such as computer names and IP addresses
Denial of service	Flooding a DNS server with queries to make it unavailable for normal use
Data modification	Falsifying records in DNS to utilize fake servers or redirect email messages
Redirection	Supplying false responses to external queries by a DNS server to corrupt the cache with false information

Therefore, when you design your DNS infrastructure, you must give careful attention to security issues to help safeguard your zone data, and to help ensure DNS availability.

Common DNS attacks include:

- Footprinting. This is the process of building a diagram or footprint of a DNS infrastructure by capturing DNS zone data. Attackers who can retrieve a list of your network's hosts and IP addresses can gain valuable information that they can then use to launch attacks against specific services.

- Denial of service. This attack attempts to make network services unavailable by flooding one or more DNS servers in the network with recursive queries. Targeted servers' central processing unit (CPU) usage can reach the maximum of 100 percent, and as a result, the DNS Server service becomes unavailable due to it no longer being able to process legitimate client requests.

- Data modification. This type of attack uses IP spoofing to modify zone data. By changing zone data, attackers can redirect users to fake web servers or redirect email to a server that attackers control.
- Redirection attack. Hackers redirect DNS name queries to servers that they control. One method of redirection involves corrupting a server's DNS cache with erroneous data that directs future queries to servers that the attacker controls.

Selecting a DNS Security Strategy

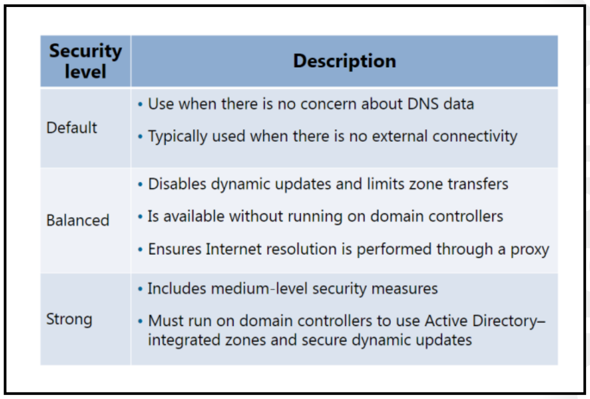

Security level	Description
Default	• Use when there is no concern about DNS data • Typically used when there is no external connectivity
Balanced	• Disables dynamic updates and limits zone transfers • Is available without running on domain controllers • Ensures Internet resolution is performed through a proxy
Strong	• Includes medium-level security measures • Must run on domain controllers to use Active Directory–integrated zones and secure dynamic updates

Security configurations can have unintended consequences. The more secure you make a component, the less usable or the more difficult to administer it becomes, or both.

When considering the potential threats that your name resolution infrastructure faces and the solutions that you propose to mitigate these threats, you should consider the potential unintended consequences carefully. Often, you must design security settings that involve a compromise. Settings must be secure enough to help protect against most perceived threats, but not so secure that the service becomes unusable or less manageable.

Selecting a DNS Security Strategy

By implementing certain features and configuring specific DNS settings, you can configure DNS security for one of the following security levels:

- Default security. The default DNS security policy is appropriate when there are no concerns about data integrity or when a private network has no external connectivity. When you use the default security settings, all DNS servers on your network:
  - Perform standard DNS resolution.
  - Are configured with root hints that point to the root servers for the Internet.
  - Permit zone transfers to any server.
  - Are configured to listen on all of their IP addresses.

Additionally, remember that the default DNS security setting:

  - Disables secure cache against pollution on all DNS servers.
  - Allows dynamic updates for all DNS zones.
  - Opens UDP and TCP/IP port 53 on the firewall for your network for both source and destination addresses.

- Balanced security. To protect zone data, this configuration disables dynamic updates and limits zone transfers, but does not require Active Directory–integrated zones. This configuration isolates internal DNS servers from the Internet by using a proxy. When you implement the balanced security settings, your organization's DNS infrastructure has limited exposure to the Internet. Additionally, all DNS servers:
  - Use forwarders to point to a specific list of internal DNS servers when they cannot resolve names locally.

- Limit zone transfers to servers listed in the name server (NS) resource records in their zones.
- Listen on specified IP addresses.

Additionally, remember that the balanced security setting:

- Enables secure cache against pollution on all DNS servers.
- Enables secure dynamic updates for all DNS zones.
- Enables internal DNS servers to communicate with external DNS servers through the firewall that contains a limited list of allowed source and destination addresses.
- Configures your external DNS servers in front of your firewall with root hints that point to the Internet's root servers.
- Ensures that proxy servers and gateways perform all Internet name resolution.

- Strong security. This configuration requires the use of Active Directory–integrated zones to implement secure dynamic updates. You can configure only domain controllers as DNS servers. Furthermore, you must configure security to restrict DNS modification to specified individuals.

  When you implement strong DNS security settings, your organization's DNS infrastructure has no Internet communication by means of internal DNS servers. In addition:

  - Your network uses an internal DNS root and namespace, where all authority for DNS zones is internal.
  - DNS servers that you do not configure with forwarders use internal DNS server IP addresses only.
  - All DNS servers limit zone transfers to the IP addresses that you specify.
  - DNS servers are configured to listen on specified IP addresses.
  - Secure cache against pollution option is enabled on all DNS servers.
  - Internal DNS servers are configured with root hints that point to the internal DNS servers that are hosting the root zone for your internal namespace. If you do not require external name resolution, you can remove root hints from all DNS servers.
  - Secure dynamic update is configured for all DNS zones, except for the top-level and root zones, which do not allow dynamic updates at all.
  - All DNS servers run on domain controllers. You configure an access control list (ACL) on the DNS Server service to allow only specific individuals to perform administrative tasks on DNS servers.
  - All DNS zones are stored in AD DS. You configure an ACL to allow only specific individuals to create, delete, or modify DNS zones.
  - You configure ACLs on a DNS resource record to allow only specific individuals to create, delete, or modify DNS data.
  - You can use IP Address Management (IPAM) to:
    - Discover DNS servers automatically across an Active Directory forest
    - Monitor DNS services and DNS zones
    - Manage DNS servers running Windows 2008 or later operating systems
    - Provide automatic and on-demand retrieval of server data from managed DNS servers and DNS zone status monitoring based on DNS zone events.

Note: These security levels do not represent a single, configurable option, but rather an approach to helping to secure your DNS security settings.

Selecting Additional Security Settings

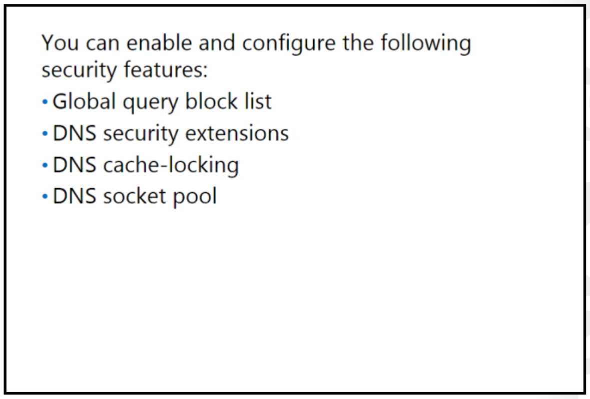

You can enable and configure the following security features:

- Global query block list
- DNS security extensions
- DNS cache-locking
- DNS socket pool

You can enable and configure several security features in addition to the generic security settings discussed in the preceding topic. These additional security features include:

- The global query block list. The Windows Server DNS dynamic update feature enables DNS resolver computers to register and dynamically update their resource records with their configured DNS server, whenever necessary. However, this convenient behavior can allow a malicious user to take over a special name and then divert certain types of network traffic to that malicious user's computer. To help prevent such a takeover, the DNS Server role in Windows Server 2012 includes a global query block list that can help prevent a malicious user from taking over DNS names that have special significance.
- DNSSEC. Windows Server 2012 DNS zones support DNSSEC. This means that you can sign and host DNSSEC-signed zones to provide additional security for your name resolution infrastructure. DNSSEC are extensions to the DNS protocol. These extensions add origin authority, data integrity, and authenticated denial of existence to DNS. These changes enable your DNS zones and the records that they contain to be signed digitally.
- DNS cache-locking. Windows Server 2012 DNS servers support DNS cache locking. When you enable cache locking, the DNS server does not allow the overwriting of cached records for the duration of the Time to Live (TTL) value. Cache locking provides improved security against cache poisoning attacks.
- DNS socket pool. Windows Server 2012 DNS servers support the DNS socket pool. Instead of using a predetermined source port (TCP or UDP 53) when issuing queries, the DNS server uses a random port number that it selects from a *socket pool*. The socket pool makes cache poisoning attacks more difficult, because a hacker must guess the DNS query source port and the random transaction ID correctly.

Note: You can configure these security options on the Advanced tab of the DNS server’s Properties dialog box.

DNSSEC in Windows Server 2012

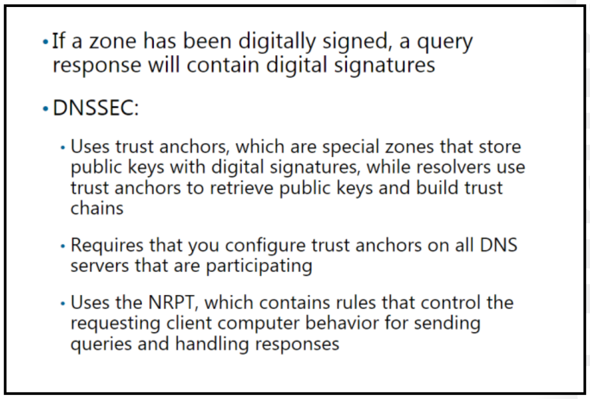

- If a zone has been digitally signed, a query response will contain digital signatures
- DNSSEC:
  - Uses trust anchors, which are special zones that store public keys with digital signatures, while resolvers use trust anchors to retrieve public keys and build trust chains
  - Requires that you configure trust anchors on all DNS servers that are participating
  - Uses the NRPT, which contains rules that control the requesting client computer behavior for sending queries and handling responses

Intercepting and tampering with an organization’s DNS query response is a common attack method. When hackers 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 performing DNS queries from accepting false DNS responses.

When a DNS server receives a query while hosting a digitally signed zone, it returns the digital signatures and 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 modified. To do this, you must configure the resolver or server 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 Delegation Signer (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, Group Policy is the preferred method of configuring the NRPT. If there is no NRPT present, the client computer accepts responses without validating them.

Deploying DNSSEC

To deploy DNSSEC, perform the following procedure:

1. Install Windows Server 2012, and then 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 also can add this role when you add the AD DS role. Then, configure the primary zones on the DNS server. After the zone signing completes, any new DNS servers in Windows Server 2012 receive the DNSSEC parameters automatically.

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 as another signed zone.
- Use recommended settings. This option signs the zone by using the default values.

Note: You can remove zone signatures by using the DNSSEC management user interface.

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 distribution points to distribute the trust anchors to all of the forest's servers. 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 you require trust anchors on computers that are not domain-joined, such as a DNS server on the perimeter network (or *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 for which 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 Features in DNSSEC for Windows Server 2012

Although Windows Server 2008 R2 supports DNSSEC, you have to perform most configuration and administration tasks manually, and zones are signed while offline. Windows Server 2012 simplifies DNSSEC implementation through several new features.

DNSSEC Zone Signing Wizard

Windows Server 2012 includes a DNSSEC Zone Signing Wizard, which helps simplify the configuration and signing process, and enables online signing. You can use the wizard to choose the zone-signing parameters. If you choose to configure the zone-signing settings rather than by using parameters from an existing zone or by using default values, you can use the wizard to configure settings, such as:

- KSK options
- ZSK options
- Trust anchor distribution options
- Signing and polling parameters

New resource records

You achieve DNS response validation by associating a private/public key pair (which the administrator generates), with a DNS zone. You then define 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 that the DNS server holds. These keys require periodic replacement through key rollovers. Windows Server 2012 supports automated key rollovers. Every zone has multiple DNSKEYs that then are broken down to the ZSK and KSK.
Delegation Signer	This delegation record contains the hash of a child zone's public key. The parent zone's private key signed this record. If a child zone of a signed parent also is signed, then you must add the DS records manually from the child to the parent to create a chain of trust.

Resource record	Purpose
Resource Record Signature	This record holds a signature for a set of DNS records, and is for checking the authority of a response.
Next Secure	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 Next Secure record, which prevents 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.
- The Windows PowerShell® command-line interface for management and scripting.

Discussion: Guidelines for Designing DNS Security

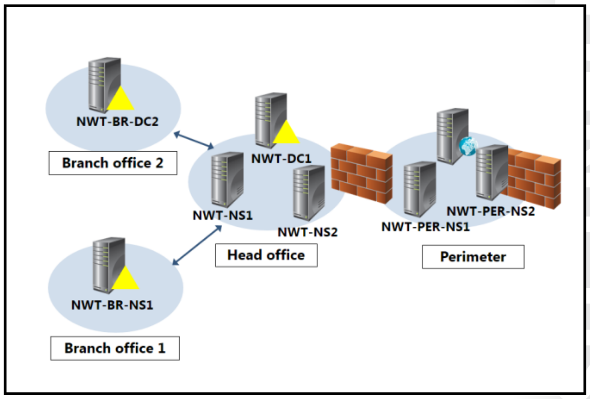

The Northwind Traders company has three locations, as shown on the slide. The Northwind Traders network has the following configuration:

- There are two domain controllers currently: NWT-DC1 in the head office, and NWT-BR-DC2 in the larger branch office. IT personnel believe that the second, smaller branch office does not need a domain controller.
- There are two DNS name servers, NWT-NS1 and NWT-NS2, which are located in the head office.
- There is one name server, NWT-BR-NS1, which is installed at the smaller branch office.
- The perimeter network, which is separated from the head office network by a suitably configured firewall, contains a web server and two DNS name servers: NWT-PER-NS1, and NWT-PER-NS2.
- High-speed WAN links connect all locations.
- The branches and the head office locations are configured as separate Active Directory sites by using the default Active Directory replication configuration.
- Neither domain controller has the DNS Server role installed.
- The head office supports several hundred users, while Branch Office 2 supports 100 users. Branch Office 1 is smaller, with about 25 users.
- A primary zone, northwindtraders.local, is configured on the NWT-NS1 name server. A secondary zone is configured on both the NWT-NS2 and the NWT-BR-NS1 name servers. In both instances, the master server is configured as NWT-NS1.
- NWT-PER-NS1 hosts a primary zone, northwindtraders.com, for external users. NWT-PER-NS2 forwards outbound queries.

Consider the current configuration, and then use the information from the preceding topics to determine how you might optimize the DNS configuration. Additionally, consider the following questions:

Question: What is the first step you might take to make the internal DNS infrastructure more secure?

Question: What configuration changes would be necessary to support your proposals?

Question: How would you recommend configuring updates on your DNS server?

Question: What DNS security policy level have you selected?

Question: Are there any other security considerations that relate to the DNS design?

Lab: Designing and Implementing Name Resolution

Scenario

A. Datum Corporation has experienced rapid growth, and the migration to Windows Server 2012 provides an ideal opportunity to validate, and where necessary, redesign and reconfigure the DNS infrastructure. As part of this process, you must review the DNS infrastructure in Contoso, Ltd, a current partner and imminent acquisition. You must examine, and where necessary, suggest changes to the DNS design.

The following tables provide some additional information about the various network locations at Contoso.

Head Office and Regional Hubs

Location	Function	Characteristics
Paris, France	Head office. Planned role is sales, marketing, and distribution center for Europe	Current employees: 1,800 (Planned employees: 4,700)
Rome, Italy	Regional hub office	Employees: 250
Barcelona, Spain	Regional hub office	Employees: 200
Munich, Germany	Regional hub office	Employees: 200
Athens, Greece	Regional hub office	Employees: 200

Regional Branches and Distribution Centers

Location	Number of servers	Number of users (total across all branches)	Branches
Germany	1 at each branch	100	3
Spain	1 at each branch	250	5
Italy	1 at each branch	250	5
Greece	1 at each branch	75	2

Objectives

After completing this lab, you will be able to:

- Design a DNS name resolution strategy.
- Implement the DNS name resolution strategy.
- Design DNS zones and zone replication.
- Implement DNS.

Lab Setup

Estimated Time: 80 minutes

Virtual machines	20413C-LON-DC1 20413C-LON-SVR1 20413C-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 **20413C-LON-DC1**, and then in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on by using the following credentials:

- User name: **Administrator**
- Password: **Pa$$w0rd**
- Domain: **Adatum**

5. Repeat steps 2 through 4 for 20413C-LON-SVR1.
6. Start 20413C-LON-CL1.

Exercise 1: Designing a Strategy for DNS Name Resolution

Scenario

Contoso, Ltd, a partner organization and an imminent acquisition, presently has no Windows Servers installed. Currently, UNIX hosts and a mix of Windows client computers support their multisite network. UNIX is providing name resolution services. The organization has a single registered domain name, Contoso.com, which Contoso uses internally and externally.

A single DNS zone manages this structure, with secondary servers distributed to the regional hub offices. Branch offices and distribution centers currently have no local DNS servers, and all clients at these offices are configured to use a DNS server at the nearest regional hub.

Contoso has external DNS records that an administrator synchronizes manually with the internal DNS structure. These records change on an average less than once per year. The table in the following "Supplemental Documentation" section lists these records.

Supplemental Documentation

Note: The following email is reproduced as a real email message with the initial messages at the end of the message thread. You must read the email thread from the bottom up.

Email from Charlotte Weiss at Contoso:

From:	Charlotte Weiss [Charlotte@contoso.com]

Sent: 04 Aug 09:05

To: Brad@Adatum.com

Subject: Re: Contoso DNS Design

Attachments: Contoso.vsd

Brad,

Yes, I am responsible for some of the IT infrastructure in Contoso, and that includes DNS. Here's what I can tell you:

- We have a single registered domain name, Contoso.com, which is used both internally and externally.
- This is managed as a single DNS zone, with secondary servers distributed to the regional hub offices.
- Branch offices and distribution centers (15 in total across Europe) have no local DNS servers at present.
- All clients at these offices are configured to use a DNS server at the nearest regional hub.
- Contoso has external DNS records that are synchronized manually with the internal DNS structure.
- These records change on an average less than once per year. The following table lists these records.

External DNS records	Purpose
www.contoso.com	Public website
Customer.contoso.com	Secure website for customers
Vpn.contoso.com	VPN server used by roaming staff
Mail.contoso.com	Internet mail server
Dns1.contoso.com	External DNS server
Dns2.contoso.com	External DNS server

I have attached a schematic of the current network infrastructure. Hope it helps.

Regards,

Charlotte

----- Original Message -----

From: Brad Sutton [Brad@Adatum.com]

Sent: 03 Aug 08:45

To: Charlotte@contoso.com

Subject: Contoso DNS Design

Hi Charlotte,

I've been assigned the role of putting together a DNS infrastructure for Contoso. I know that you've been involved in the existing DNS deployment over in Paris. I'm wondering if you could send over any information that you have that relates to that existing deployment.

Thanks,

Brad

Network Diagram of Contoso Locations (Contoso.vsd)

The network diagram of the Contoso locations is as follows:

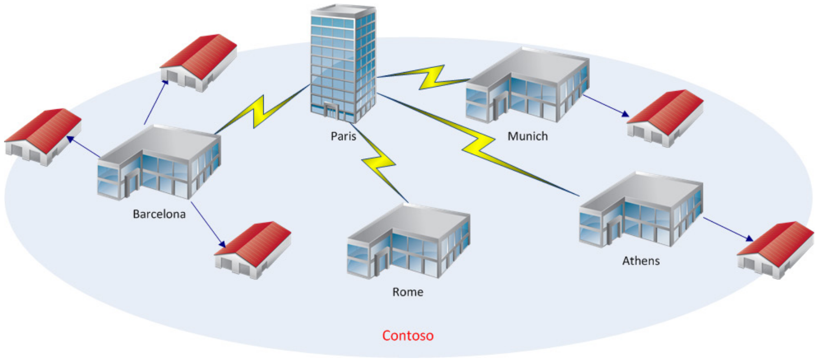

Proposal Document

Contoso DNS Name Resolution Strategy

Document Reference Number: BS00806/1

Document Author	Brad Sutton
Date	6th Aug

Requirements Overview

Design a new DNS name resolution strategy:

- Fault tolerance and performance of name resolution are important.
- The plan should support any anticipated growth or changes in the network infrastructure.
- It is highly likely that Contoso will implement AD DS when the organization is acquired in the near future. Any DNS design should accommodate this change.
- Even though the expected merger between A Datum, Trey Research, and Contoso is yet to occur, users are sharing resources already between the two organizations. Recent network monitoring has shown that the volume of name resolution queries for adatum.com and treyresearch.net has increased significantly. The design must accommodate this fact.

Proposals

1. If you create a new design, what would be your preferred namespace for AD DS?
2. What additional factor should you consider when modifying an existing design?
3. What DNS namespace do you recommend that Contoso use for AD DS?

The main tasks for this exercise are as follows:

1. Read the supporting documentation

2. Update the proposal document with your planned course of action

3. Examine the suggested proposals in the Lab Answer Key

4. Discuss your proposed solution with the class, as guided by your instructor

▶ Task 1: Read the supporting documentation

- Read the documentation provided.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the Contoso DNS Name Resolution Strategy document.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with those in the Lab Answer Key.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you will have created a DNS name resolution design.

Exercise 2: Designing a Strategy for DNS Server Placement

Scenario

When considering the new DNS design, planning the placement of DNS servers is important because it helps you minimize WAN traffic and ensure high availability of the name resolution service. You must determine which locations will have DNS servers, based on the network infrastructure and number of users. Additionally, the failure of a WAN link should not cause a failure in name resolution.

Supporting Documentation

Proposal Document

Contoso DNS Server Placement Strategy	
Document Reference Number: BS00810/2	
Document Author Date	Brad Sutton 10th Aug
Requirements Overview Plan a new DNS server placement strategy: • Fault tolerance and performance of name resolution are important. • The plan should support any anticipated growth or changes in the network infrastructure. • It is highly likely that Contoso will implement AD DS when the organization is acquired in the near future. Any DNS plan should accommodate this change.	
Additional Information • There are approximately 30 to 50 client computers and one server at each of the smaller regional branch offices. • Significant growth is anticipated at the head office. • UNIX hosts that provide DNS are running BIND 9.6. • Contoso has not deployed AD DS, but does have imminent plans to do so.	

Contoso DNS Server Placement Strategy

- Each regional hub has a single DNS server. However, the head office has two DNS servers for internal resolution.
- Small branches do not have a local DNS server, and rely instead on name resolution through their nearest regional hub.
- Contoso uses the same domain name (Contoso.com) internally and externally.

Proposals

1. How many DNS servers do you require at the head office in Paris?
2. Do the branch locations require DNS servers?
3. Are additional DNS servers required at each regional hub site?

The main tasks for this exercise are as follows:

1. Read the supporting documentation

2. Update the proposal document with your planned course of action

3. Examine the suggested proposals in the Lab Answer Key

4. Discuss your proposed solution with the class, as guided by your instructor

▶ Task 1: Read the supporting documentation

Read the documentation provided.

▶ Task 2: Update the proposal document with your planned course of action

Answer the questions in the proposals section of the Contoso DNS Server Placement Strategy document.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

Compare your proposals with those in the Lab Answer Key.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you will have determined where to place Contoso DNS servers to support your initial DNS design.

Exercise 3: Designing DNS Zones and DNS Zone Replication

Scenario

After determining the location of DNS servers, you must determine how to divide the DNS namespace and how to perform replication. You should manage DNS separately for each of the Active Directory domains. Each DNS zone should be capable of performing secure dynamic updates for computers in the local domain.

Supporting Documentation

Proposal Document

Contoso DNS Zones and Zones Replication Strategy	
Document Reference Number: BS00812/2	
Document Author Date	Brad Sutton 12th Aug
Requirements Overview Plan a new DNS zone and zone-transfer strategy: • Fault tolerance and name resolution performance are important. • The plan should support any anticipated growth or changes in the network infrastructure. • It is highly likely that Contoso will implement AD DS when the organization is acquired in the near future. Any DNS plan should accommodate this change.	
Additional Information • You can assume that management accepted your DNS design and server-placement plans.	
Proposals 1. Which zones do you need to create on internal DNS servers? 2. Which zones do you need to create on external DNS servers? 3. In which regional hub sites will you place each DNS zone? 4. How will you configure replication or zone transfers for each zone? 5. How would implementing AD DS affect your design?	

The main tasks for this exercise are as follows:

1. Read the supporting documentation

2. Update the proposal document with your planned course of action

3. Examine the suggested proposals in the Lab Answer Key

4. Discuss your proposed solution with the class, as guided by your instructor

▶ Task 1: Read the supporting documentation

7. Read the documentation provided.

▶ Task 2: Update the proposal document with your planned course of action

Answer the questions in the proposals section of the Contoso DNS Zones and Zones Replication Strategy document.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

Compare your proposals with the ones shown above.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you will have a DNS zone design that you can use to implement DNS.

Exercise 4: Implementing DNS

Scenario

In preparation for deploying and configuring DNS at Contoso, you have been assigned to a project team that is deploying additional DNS servers at A. Datum.

The main tasks for this exercise are as follows:

1. Install the DNS Server role

2. Create and configure secondary zones

3. Enable and configure zone transfers

4. Test DNS resolution from a client

5. To prepare for the next module

▶ Task 1: Install the DNS Server role

Switch to LON-SVR1.
If necessary, sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
From Server Manager, add the **DNS Server** role with the default values.
Leave both Server Manager and the DNS console open.

▶ Task 2: Create and configure secondary zones

1. On LON-SVR1, open Windows PowerShell as an administrator.
2. At the command prompt, type the following cmdlets, pressing Enter at the end of each row:

```
add-dnsserversecondaryzone –masterservers 172.16.0.10 –Name Adatum.com –Zonefile
”Adatum.com.dns”
Register-DnsClient
```

3. Switch to LON-DC1. If necessary, sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
4. Open DNS, and modify the properties of the **Adatum.com** zone.
5. Allow zone transfers to servers that are listed on the **Name Servers** tab.
6. Configure the LON-SVR1 server in the **Notify** list by using the IP address: **172.16.0.11**.

7. Add the **LON-SVR1.Adatum.com** server as a named server on the **Name Servers** tab.

8. Switch to LON-SVR1, and open DNS.

9. Verify that the **Adatum.com** zone exists and contains records.

Task 3: Enable and configure zone transfers

1. On LON-SVR1, in **DNS Manager**, perform the following tasks:

    a. Configure LON-DC1 as this server's DNS forwarder.

    b. Disable round-robin DNS.

    c. Delete all records in Root Hints.

2. Switch to LON-DC1.

3. Open Windows PowerShell as an administrator.

4. At the command prompt, type the following command, and then press Enter:

```
set-dnsserverglobalnamezone
–enable $true
```

Note: This command enables support for the GlobalNames zone.

5. In DNS Manager, create a new primary zone on LON-DC1 with the following properties:

- Zone type: **Primary**
- Active Directory Zone Replication Scope: **To all DNS server running on domain controllers in this forest: Adatum.com**
- Zone name: **GlobalNames**
- Dynamic Update: Accept default values

Task 4: Test DNS resolution from a client

1. Sign in on LON-CL1 as **Adatum\Administrator** with the password **Pa$$w0rd**.

2. Run Windows PowerShell as administrator, and then use the **Get-DnsClientServerAddress** cmdlet to see the assigned DNS servers for **LON-CL1**.

3. Use Network Connections to add the LON-SVR1 IP Address as an Alternate DNS server.

4. Run Windows PowerShell, and then observe the output of the following cmdlets:

- **DnsClientServerAddress**
- **Get-DnsClientCache**
- **Resolve-DnsName LON-SVR1**

Task 5: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.

2. In the **Virtual Machines** list, right-click **20413C-LON-DC1**, and then click **Revert**.

3. In the **Revert Virtual Machine** dialog box, click **Revert**.

4. Repeat steps 2 and 3 for 20413C-LON-SVR1 and 20413C-LON-CL1.

Results: After completing this exercise, you will have implemented DNS successfully.

Question: What was your approach to the DNS design exercises?

Question: Did your design differ from the suggested solution?

Question: How does the DNS design for Contoso compare with your organization's DNS implementation?

Module Review and Takeaways

Best Practice: Whenever possible, use Active Directory–integrated DNS. This provides you with more fault tolerance, because it enables you to have more than one domain controller. This provides better security, because Kerberos version 5 (V5) protocol encrypts all domain replication, and dynamic registrations can occur in secure mode. Furthermore, Active Directory–integrated DNS is more efficient, because all domain replication, including DNS record changes, use attribute-level replication. This is faster and uses fewer bits than typical full-zone transfers (AXFRs), or even incremental zone transfers (IXFRs).

Review Questions

Question: What is the difference between a subdomain in a DNS zone, and a delegated zone?

Question: Contoso has created a regional sales department. Some sales staff is located at regional sales centers, where there are approximately 10 computers. These computers should be able to access the same applications and resources as the rest of the Contoso staff. How would you implement DNS at these smaller branches?

Question: True or false? You should disable recursion on all internal DNS servers.

Question: Why is it not good practice to disable round-robin rotation on all DNS servers?

Question: When would you configure a caching-only server?

Question: When considering NetBIOS name resolution, when would you choose WINS over the GlobalNames zone?

Question: You are concerned about the security of zone data while it travels across the network during a zone transfer. All of your DNS servers also are domain controllers. What two strategies could you implement to mitigate your perceived security threats?

Module 6

Designing and Implementing an Active Directory Domain Services Forest and Domain Infrastructure

Contents:

Module Overview

To design the infrastructure of an Active Directory® forest for your organization, you must first gather organizational and administrative requirements, and then you must decide which design to use. There are several possible designs that you can choose from, and each design requires some degree of compromise.

Based on your organization’s needs, you must determine the type of Active Directory forest and forest root domain that you need. You also must determine whether your organization requires multiple forests, and whether you will need trusts between forests.

After selecting an appropriate Active Directory forest design, you must then design the domain mode. This involves determining your organization’s needs, selecting the location of domain controllers, and then deploying the domain controllers. After designing the Active Directory domain infrastructure, you then integrate the internal and external Domain Name System (DNS) namespaces with the Active Directory domain by using DNS servers. Finally, if your design consists of multiple domains, you create domain trusts to enable communication from one domain to another.

Objectives

After completing this module, you will be able to:

- Design an Active Directory forest.
- Design and implement Active Directory forest trusts.

- Design Active Directory integration with Windows Azure™ Active Directory.
- Design and implement Active Directory domains.
- Design DNS namespaces in an Active Directory environment.
- Design and implement Active Directory domain trusts.

Lesson 1
Designing an Active Directory Forest

When designing your Active Directory forest, a forest design based on a single forest provides for the most straightforward integration and greatest simplicity. This method also supports features such as a single global address list (GAL) in Microsoft® Exchange Server. However, there are circumstances where organizations require multiple forests to address security concerns, governance issues, or administrative isolation requirements. In this lesson, you will explore the criteria that can help you to decide how best to implement an Active Directory forest in your organization.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe Active Directory forests.
- Describe the design models of Active Directory forests.
- Describe the benefits and disadvantages of a single forest model.
- Implement multiple forests.
- Design an Active Directory forest infrastructure.
- Select an appropriate forest design.

What Is an Active Directory Forest?

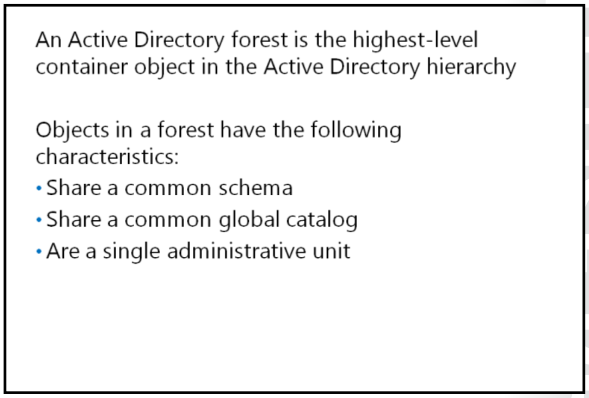

An Active Directory forest is the highest-level container object in the Active Directory hierarchy

Objects in a forest have the following characteristics:
- Share a common schema
- Share a common global catalog
- Are a single administrative unit

An *Active Directory forest* is the highest-level container object in the Active Directory Domain Services (AD DS) hierarchy. It is a container for any domain within your organization.

An Active Directory forest has the following characteristics:

- All objects within a forest share a common Active Directory schema. The schema is the collection of object types and their respective attributes.
- The forest has a single global catalog. All domain controllers that have been assigned the global catalog role share this global catalog.
- The forest has a single Enterprise Admins group. Members of this group belong, indirectly, to the Domain Admins groups in the member domains. By default, Enterprise Admins can perform administrative tasks anywhere in the forest.

Because of these characteristics, you can consider a forest to be a single administrative unit with a common global catalog and schema. A forest is often referred to as a *security boundary* because it provides the most complete and secure separation of Active Directory domains.

Note: An Active Directory forest has a single root domain. This root domain contains two forest-wide operations master roles: the schema operations master, and the domain naming operations master.

Active Directory Forest Models

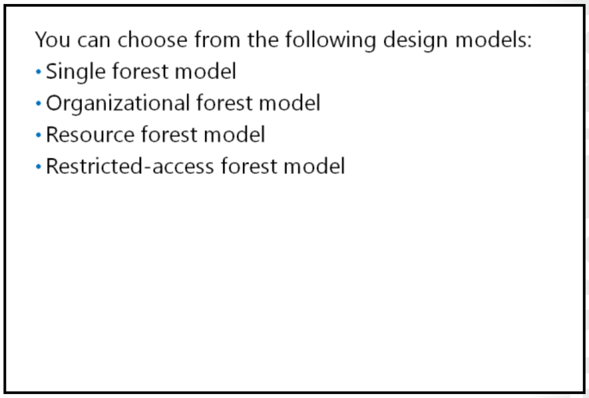
You can choose from the following design models:

- Single forest model
- Organizational forest model
- Resource forest model
- Restricted-access forest model

Choosing an appropriate forest design model at an early stage is an important part of your infrastructure design and planning. This is because trying to change the design later can be difficult and time-consuming. When planning your forest design model, you can choose from the following designs:

- Single forest model. In this design model, all user accounts and resources are contained in one or more domains that are in a single Active Directory forest. This is the most common forest design model, and it is the easiest model to implement and support. The single forest model satisfies the requirements of most organizations.
- Organizational forest model. In this design model, administrators design a forest to accommodate a number of factors based on an organization's needs. These factors may include departmental structure, physical locations, and any business criteria. Often, these same organizational needs take precedence over other design considerations.

  In this design model, the forest includes user accounts and resources, which administrators manage independently. This means that the administrator usually creates multiple forests in a single organization, each with administrative separation and autonomy. If users in one of the organizational forests require access to resources in another forest, you can create trust relationships between the organizational forest and the other forest.
- Resource forest model. Sometimes, in an environment that has an application, shared folder, or other system resource that is particularly critical or secure, administrators create a forest specifically for users who must access that resource. Usually, this type of forest—known as a *resource forest*—is a new or additional forest in an organization, and you then establish trusts to access this forest.

  Resource forests contain only those user accounts that are required to provide service administration and alternate access to resources in the forest. You can establish forest trusts so that users from other forests can access the resources in the resource forest.
  Resource forests also provide service isolation, which helps protect areas of the network that need to be highly available. For example, if your company has a critical department that needs to operate even if the rest of the network has problems, you can create a separate resource forest for that department's services.
- Restricted-access forest model. A separate forest includes user accounts and data that you must isolate from the rest of the organization. Users from other forests cannot be granted access to the restricted data, because no trust exists.

  In this model, restricted access users have one account in an organizational forest so they can access general organizational resources, and they have a separate account in the restricted-access forest so that they can access the classified data. These users must use two separate workstations—one that connects to the organizational forest, and the other that connects to the restricted-access forest. This helps protect against the possibility that a service administrator from one forest can access a workstation in the restricted forest.
  In extreme cases, you might maintain the restricted-access forest on a separate physical network. Organizations that work on classified government projects sometimes maintain restricted-access forests on separate networks to meet security requirements.

Benefits of a Single Forest Model

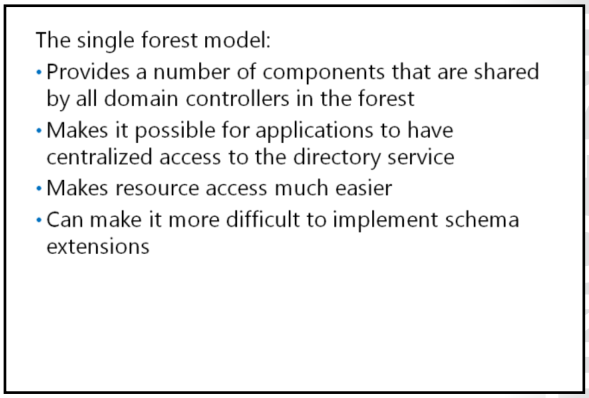

The single forest model:

- Provides a number of components that are shared by all domain controllers in the forest
- Makes it possible for applications to have centralized access to the directory service
- Makes resource access much easier
- Can make it more difficult to implement schema extensions

When choosing between the four forest models, as a best practice you should consider the single forest model first. This is because the single forest model offers several advantages:

- A common set of components and a simple design. A single Active Directory forest has a number of components that all domain controllers in the forest share. Some of these components include a single schema directory partition, a single configuration directory partition, and global catalog information. In addition, each Active Directory forest has a set of administrators (the built-in Enterprise Admins group) that can perform forest-level administration, such as adding or removing domains from the forest.

- Centralized access to the directory service. In the single forest model, applications such as Exchange Server can have centralized access to the directory service. Exchange Server uses forest-wide data in the schema and configuration partitions, so you can implement a single Exchange Server organization to manage all domains in a forest and to serve the users in those domains. This enables easier collaboration, because everyone can more easily see the GAL and the free/busy information.

- Quick resource access. With a single forest, users can access resources easily because by default all domains in the forest trust each other. This means that the administrator of one domain can easily provide resource access to users from other domains in the same forest.

Disadvantages of a Single Forest Model

The single forest model has some disadvantages:

- Security risk. From a security perspective, be aware that no real security boundaries exist inside the forest. Consequently, you cannot isolate any domain from any other domain. However, you can isolate some forest-related accounts such as the Enterprise Admin and Schema Admin accounts, by creating an empty, dedicated forest root domain. (A *dedicated forest root domain* is an Active Directory domain that you create exclusively to function as the forest root domain.) Beneath the dedicated forest root domain, you create resource and account-holding domains. The forest root domain does not contain any end-user accounts. If you use a dedicated forest root domain, you can separate the accounts of forest-level service administrators from those of domain-level service administrators.

- Difficulty in implementing schema extensions. In large organizations, implementing forest-wide changes such as schema extensions can be difficult. It becomes even more difficult if you use a single forest. Each department within the organization might be concerned about how any schema extensions might affect their department. Moreover, it might be the case that only a single department in the organization requires a forest-wide change.

- Governance process causing a delay in implementing changes. If an organization deploys a single forest, the organization must create a proper governance process that involves all stakeholders. This governance process can cause forest-wide changes to take longer to implement.

Considerations for Implementing Multiple Forests

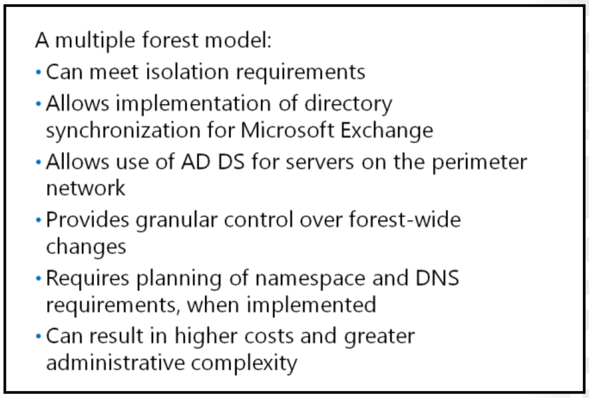

A multiple forest model:

- Can meet isolation requirements
- Allows implementation of directory synchronization for Microsoft Exchange
- Allows use of AD DS for servers on the perimeter network
- Provides granular control over forest-wide changes
- Requires planning of namespace and DNS requirements, when implemented
- Can result in higher costs and greater administrative complexity

Although a single Active Directory forest typically offers the best design compromise, in some cases business, administrative, or security requirements may lead you to implement multiple forests.

The following are some reasons to implement multiple forests:

- Your company is merging with another organization, and merging the two forests into a single entity can be prohibitively expensive.
- Your company has a department that must be isolated for security or administrative reasons. An isolation requirement is one of the most common reasons for using multiple forests.
- You need to implement directory synchronization between multiple forests. Multiple forests allow you to use a directory synchronization product such as Microsoft Forefront® Identity Manager to synchronize objects between two or more forests. One of the most common uses for directory synchronization is to synchronize GALs in Microsoft Exchange by populating contact objects in both forests.
- You want to place domain controllers on a perimeter network. In this case, you need a separate Active Directory forest so that you can isolate data and separate authentication and authorization for internal resources from perimeter resources. This scenario is a special case of a general isolation requirement.

Note: In this situation, consider alternatives such as implementing Active Directory Lightweight Directory Services (AD LDS) on the perimeter network to support your requirements.

- You need granular control over forest-wide changes. Implementing schema changes in the single forest model can be difficult or even impossible, and if you require granular control over this process, you must implement multiple forests.

Considerations for Implementing Multiple Forests

If you implement multiple forests, you must consider other factors that relate to administration and naming. Each forest must have a unique forest namespace that any user who requires forest access can resolve. This means that you also must plan for name resolution and DNS design. In addition, if you assume that users in each forest require resource access to the other forest, you will need to implement forest trusts.

Disadvantages of Multiple Forests

Implementing multiple forests has the following disadvantages:

- Implementing multiple forests has much higher design, implementation, hardware, and administrative costs than a single forest implementation.
- You must establish trusts if you want to share network resources.
- Global catalog queries include only objects in the local forest.

Guidelines for Designing an Active Directory Forest Infrastructure

When you design your Active Directory forest infrastructure, consider the following guidelines:

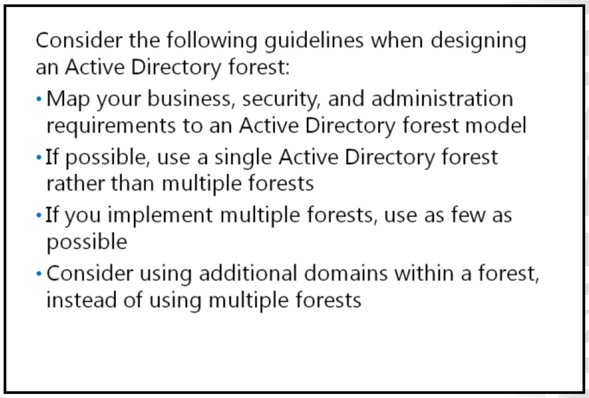

Consider the following guidelines when designing an Active Directory forest:

- Map your business, security, and administration requirements to an Active Directory forest model
- If possible, use a single Active Directory forest rather than multiple forests
- If you implement multiple forests, use as few as possible
- Consider using additional domains within a forest, instead of using multiple forests

- Gather and document the business, security, and administrative requirements. These requirements typically determine the most appropriate forest model. They also define features that you must implement. After you understand all requirements completely, it is a good idea to map the requirements to an Active Directory forest model. In some cases, you might not be able to find an ideal solution, so you should select the forest model that addresses most of your requirements, or the most important requirements.
- Use a single Active Directory forest rather than multiple forests whenever possible. The infrastructure will be less complex to support, and will provide all of the benefits previously discussed for multiple forests. Additionally, the cost will likely be lower, and you will need fewer IT staff to maintain and administer the system.
- Use the fewest forests possible. The more forests that you deploy, the more independent Active Directory infrastructures you must support and maintain. In addition, the trust infrastructure and resource access management will be more complex.
- Consider using additional domains within a forest rather than multiple forests to meet your business and security requirements. Multiple domains in a single forest require much less administrative effort than multiple forests. If you need administrative isolation but you still want to share some parts of the infrastructure, you should consider using multiple domains and use a dedicated forest root domain.

Discussion: Selecting a Suitable Forest Design

Wingtip Toys is planning to merge with Tailspin Toys, a former competitor. The following information has been gathered about the respective IT infrastructures at these two organizations:

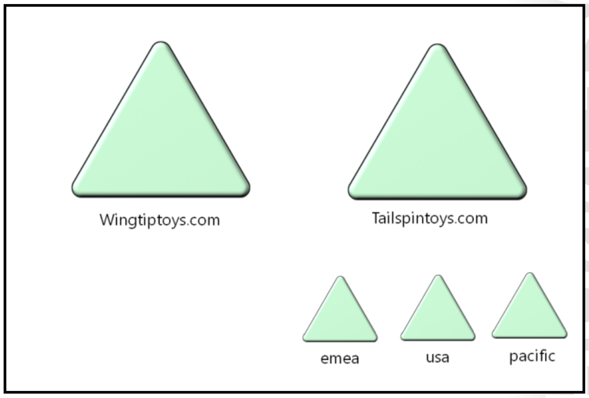

- Wingtip Toys has implemented an Active Directory forest consisting of a single domain in Windows Server® 2012.
- Tailspin Toys has a Windows Server 2008 R2–based AD DS single forest comprising multiple domains: tailspintoys.com (the forest root domain), emea.tailspintoys.com, usa.tailspintoys.com, and pacific.tailspintoys.com.
- High bandwidth links connect the respective sites of each organization. Planning calls for deploying additional links between the two organizations' locations.
- Tailspin Toys has a custom application that required schema modifications for deployment. This application manages critical manufacturing processes.

- Both organizations implement on-premises messaging systems.
- Both organizations produce toys that have well-known worldwide brand names. These brands must remain unaffected by any IT changes.
- The cost of any restructuring must be kept to a minimum.

Consider the current configuration, and then using this information, determine how you might design the Active Directory forest infrastructure for this merged organization. For help with your design process, consider the following questions.

Question: How many forests are required to integrate the two organizations?

Question: How would you recommend integrating the two organizations?

Question: What is the relevance of the schema changes in the Tailspin Toys forest, to any design that you might consider?

Question: How do the existing external domain names used affect your design?

Lesson 2
Designing and Implementing Active Directory Forest Trusts

When you deploy multiple forests, for purposes of collaboration you might need to implement trusts between those forests. When designing trusts, you can use either external trusts or forest trusts, according to your business requirements. You can also link two disjoined forests by using a transitive trust. Additionally, you need to decide on a method such as user principal name (UPN) suffix routing, or selective authentication to secure the forest trusts.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe characteristics of a forest trust.
- Describe the security considerations to take into account while planning forest trusts.
- Explain how resource access works in forest trusts.
- Design forest trusts.
- Create a forest trust.

Characteristics of Forest Trusts

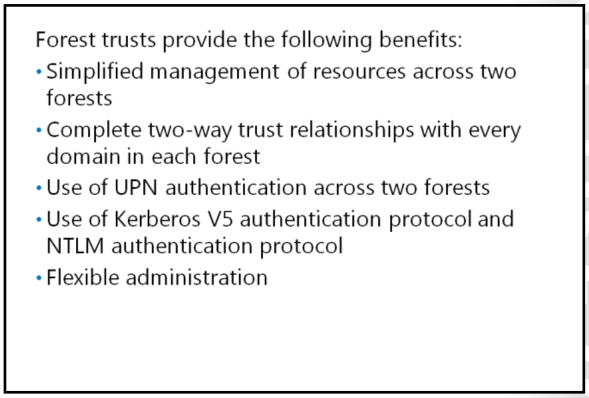

Forest trusts provide the following benefits:

- Simplified management of resources across two forests
- Complete two-way trust relationships with every domain in each forest
- Use of UPN authentication across two forests
- Use of Kerberos V5 authentication protocol and NTLM authentication protocol
- Flexible administration

When you require collaboration between two separate organizations with two separate forests, consider implementing a forest trust. A *forest trust* is a one-way or two-way trust relationship between the forest root domains of two forests. A single forest trust relationship allows users who are authenticated by a domain in one forest to access resources that are in another 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 useful particularly in scenarios that involve cross-organization collaboration or mergers and acquisitions. Forest trusts are also useful within a single organization that has more than one forest in which to isolate Active Directory data and services.

In AD DS, you can link two forests together to form a one-way or two-way trust relationship. You can use a two-way forest trust to form a transitive trust relationship between every domain in both forests, although 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 then create a forest trust between Forest 2 and Forest 3, Forest 1 does not have an implicit trust with Forest 3.

Forest trusts are useful for:

- Application service providers.
- Companies undergoing mergers or acquisitions.

- Collaborative business extranets.
- Companies seeking a solution for administrative autonomy.

Benefits of Forest Trusts

Forest trusts provide the following benefits:

- Simplified management of resources across two Active Directory 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 both the Kerberos version 5 (V5) authentication protocol and NTLM authentication protocols to improve the trustworthiness of authorization data that is transferred between forests
- Flexibility of administration that addresses the need for administrative tasks that can be unique to each forest

You must address several requirements before you can implement a forest trust. For example, the forest functional level must be configured to Windows Server 2003 or newer, and you must have functional name resolution between the forests.

Forest Trust Security Considerations

In some cases, forest trusts can cause security issues. Additionally, if you do not configure a forest trust properly, users who belong to another forest can gain unwanted access to some resources. However, several technologies exist that you can use to help control and manage security in a forest trust.

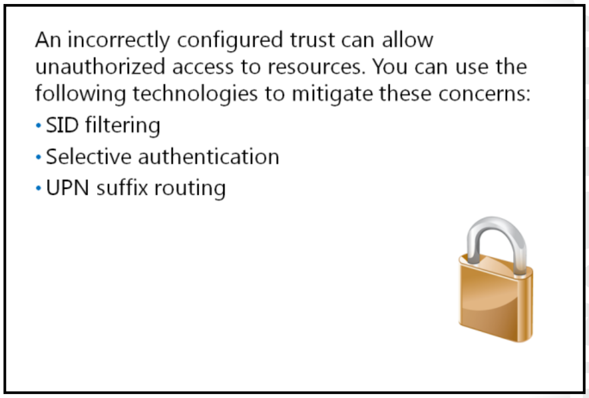

An incorrectly configured trust can allow unauthorized access to resources. You can use the following technologies to mitigate these concerns:

- SID filtering
- Selective authentication
- UPN suffix routing

SID Filtering

By default, when you establish a forest trust, you enable a domain quarantine, which also is known as *security identifier (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 security identifiers from other attributes of the user and the user's groups.

AD DS sets SID filtering by default to help prevent unauthorized users who have access at the domain or enterprise administrator level in a trusted forest from granting (to themselves or to other user accounts in their forest) elevated user rights to a trusting forest. The act of maliciously elevating privileges is often referred to as a *privilege escalation attack*.

SID filtering prevents misuse of the attributes that contain SIDs on security principals (including inetOrgPerson) in the trusted forest. One common example of an attribute that contains a SID is the SID history attribute (**sIDHistory**) of a user account object. Domain administrators typically use the SID history attribute to migrate the user and group accounts that are stored by a security principal from one domain to another. All SID filtering activities occur in the background, and administrators do not need to configure anything explicitly unless they want to disable SID filtering.

When security principals are created in a domain, the SID of the principal includes the domain SID. The domain SID is important because the Windows security subsystem uses it to verify the identity of the security principal, which in turn determines which domain resources the principal can access. You also can use the domain SID to identify in which domain it was created.

Selective Authentication

When you create an external trust or a forest trust, you can control the scope of authentication of trusted security principals. There are two modes of authentication for an external or forest trust:

- Selective authentication
- Domain-wide authentication (for an external trust) or forest-wide authentication (for a forest trust)

If you choose domain-wide or forest-wide authentication, remember that this authentication type enables all trusted users to authenticate for services access on all computers in the trusting domain. Trusted users, therefore, can be given permission to access resources anywhere in the trusting domain.

If you use this authentication mode, you must have confidence in your enterprise's security procedures and in the administrators who implement those procedures, so that trusted users will not receive inappropriate access to services. Remember, for example, that users from a trusted domain or forest are considered Authenticated Users in the trusting domain. Therefore, 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, the users are allowed to authenticate only for services on computers that you specify. 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 Active Directory forests that are joined by forest trusts. To simplify the administration of authentication requests, when you create a forest trust, by default AD DS routes all unique name suffixes.

A *unique name suffix* is any one of the following:

- A name suffix within a forest (such as a UPN suffix)
- Service principal name (SPN) suffix
- DNS forest or domain tree name that is not subordinate to any other name suffix

For example, the DNS forest name contoso.com is a unique name suffix within the contoso.com forest.

AD DS routes all names that are subordinate to unique name suffixes implicitly. For example, if your forest uses contoso.com as a unique name suffix, authentication requests for all child domains of contoso.com (childcomain.contoso.com) are routed, because the child domains are part of the contoso.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.

Resource Access

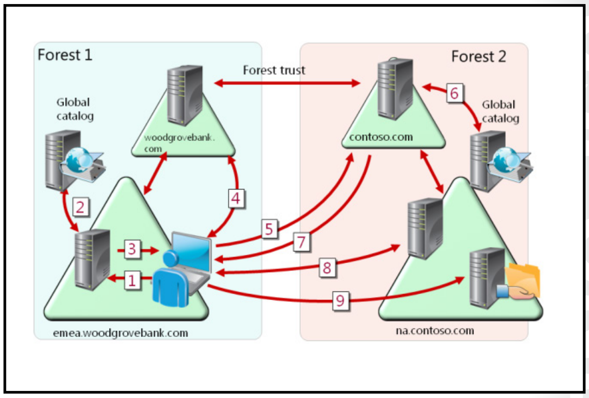

If you implement forest trusts, you can enable users from one forest to access resources in another forest by adding them to the access control list (ACL) of the target resource. If a user attempts to access a resource in a trusted forest, AD DS must first locate the resource. After it does, the user can be authenticated and allowed to access the resource.

The following example is a description of how a client computer locates and accesses a resource in another Active Directory forest:

1. A user who is signed into the domain emea.woodgrovebank.com attempts to access a shared folder in the contoso.com forest. The user's computer contacts the domain controller in emea.woodgrovebank.com and requests a service ticket by using the SPN of the computer on which the resource resides. An SPN can be the DNS name of a host or a domain, or it can be the distinguished name of a service connection point object.

2. The resource is not located in emea.woodgrovebank.com, so the domain controller for emea.woodgrovebank.com queries the global catalog to see if the resource is in another domain in the forest. Because a global catalog contains information only about its own forest, it does not find the SPN. The global catalog then checks its database for information about any forest trusts that are established with its forest. If the global catalog is successful, it compares the name suffixes that are listed in the forest trust's trusted domain object to the suffix of the target SPN. After it finds a match, the global catalog provides routing information to the domain controller in emea.woodgrovebank.com about how to locate the resource.

3. The domain controller in emea.woodgrovebank.com sends a referral for its parent domain, woodgrovebank.com, to the user's computer.

4. The user's computer contacts a domain controller in woodgrovebank.com for a referral to a domain controller in the forest root domain of the contoso.com forest.

5. Using the referral that the domain controller in the woodgrovebank.com domain returns, the user's computer contacts a domain controller in the contoso.com forest for a service ticket to the requested service.

6. The resource is not located in the forest root domain of the contoso.com forest, so the domain controller contacts a global catalog to find the SPN. The global catalog finds a match for the SPN, and then it sends the SPN to the domain controller.

7. The domain controller sends the user's computer a referral to na.contoso.com.

8. The user's computer contacts the Key Distribution Center (KDC) on the domain controller in na.contoso.com, and then it negotiates a ticket that enables the user to gain access to the resource in the na.contoso.com domain.

9. The user's computer sends the server service ticket to the computer on which the shared resource is located. The computer reads the user's security credentials, and then it constructs an access token, which gives the user access to the resource.

Note: The user's computer must be able to communicate with all of the domain controllers in the trust path. Without the ability to communicate with all of the domain controllers in the trust path, the user will not have access to the resource.

Guidelines for Designing Forest Trusts

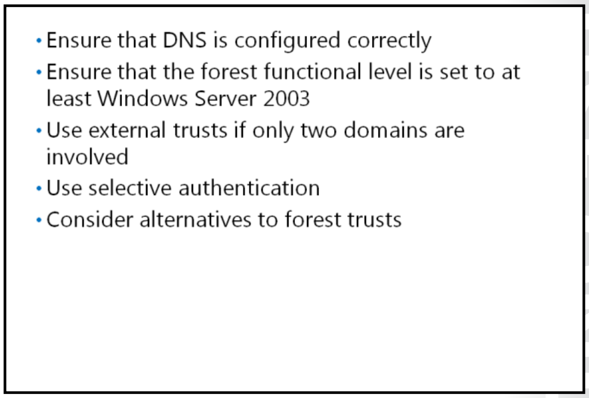

- Ensure that DNS is configured correctly
- Ensure that the forest functional level is set to at least Windows Server 2003
- Use external trusts if only two domains are involved
- Use selective authentication
- Consider alternatives to forest trusts

When designing your forest trust, consider the following guidelines:

- Ensure that you configure DNS properly so that name resolution between the forests functions correctly. Name resolution is required so that DNS can resolve the IP addresses for the domain controllers in both forests. Each forest DNS must have a conditional forwarder or a stub zone that points to another forest namespace.
- Ensure that the forest functional level is set to at least Windows Server 2003. Earlier forest functional levels do not support forest trusts. If you do not have the Windows Server 2003 or newer forest functional level, you can use an external trust.
- Do not create a forest trust if you simply need to establish a trust relationship between two specific domains in two forests.
- Use selective authentication to specify the scope of authentication for users who are authenticating through forest trusts. By using selective authentication, you can specify the users and groups from a trusted forest who are allowed to authenticate to specific resource servers in a trusting forest.
- Consider alternatives to a forest trust. If a client only needs to access one application or web service in another forest, you can implement AD LDS or Active Directory Federation Services (AD FS). This enables you to authenticate clients and provide access without having to establish a forest trust.

Demonstration: Creating a Forest Trust

In this demonstration, you will see how to:

- Configure the prerequisites for a forest trust.
- Create a forest trust.

Demonstration Steps

Configure the prerequisites for a forest trust

1. On LON-DC1, in Server Manager, open the DNS management console, and then create a new conditional forwarder with the following settings:
   - DNS Domain: **treyresearch.net**
   - IP address: **172.16.10.10**
2. Open a command prompt, and verify that you can resolve names in the target DNS domain by using the following command:

```
Nslookup trey-dc1.treyresearch.net
```

3. Switch to TREY-DC1.
4. If necessary, sign in as **Treyresearch\Administrator** with the password **Pa$$w0rd**.

5. Open the DNS management console, and create a conditional forwarder with the following settings:
    - Domain name: **Adatum.com**
    - IP address: **172.16.0.10**
6. Open a command prompt and verify that you can resolve names in the target DNS domain by using the following command:

```
Nslookup lon-svr1.adatum.com
```

Create a forest trust

1. Switch to LON-DC1.
2. From Server Manager, open Active Directory Domains and Trusts.
3. View the **Adatum.com** domain properties.
4. Create a new forest trust with the following properties:
    - Trust name: **treyresearch.net**
    - Trust type: **Forest trust**
    - Direction: **Two-way**
    - Sides of trusts: **Both this domain and the specified domain**
    - Authentication: **Forest-wide authentication**
    - Confirm Outgoing Trust: **Yes**
    - Confirm Incoming Trust: **Yes**

Lesson 3
Designing Active Directory Integration with Windows Azure Active Directory

Implementing Windows Azure Active Directory is not the same as deploying virtual machines in Windows Azure, adding AD DS, and deploying some domain controllers for a new forest and domain. Windows Azure Active Directory is a small subset of AD DS that is built into Windows Azure, and provides authentication and authorization services in a private cloud. Windows Azure Active Directory only authenticates users, not computers. As such, it is not meant to replace an on-premises deployment of AD DS.

Understanding the benefits of Windows Azure Active Directory is an important part of designing and maintaining an identity infrastructure. In this lesson, you will learn about what Windows Azure Active Directory is, where it fits in a typical environment, what services it offers, and how it integrates with AD DS.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe Windows Azure Active Directory and its advantages.
- Describe the authentication protocols that Windows Azure Active Directory supports.
- Integrate authentication with on-premises AD DS.
- Plan directory synchronization between an on-premises Active Directory environment and Windows Azure Active Directory.
- Manage Windows Azure Active Directory user accounts.

Windows Azure Active Directory Overview

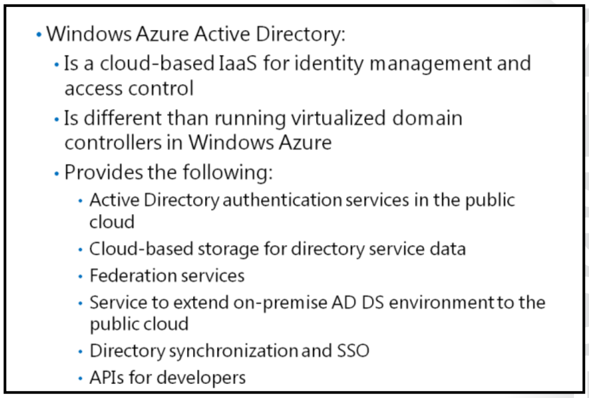

- Windows Azure Active Directory:
  - Is a cloud-based IaaS for identity management and access control
  - Is different than running virtualized domain controllers in Windows Azure
  - Provides the following:
    - Active Directory authentication services in the public cloud
    - Cloud-based storage for directory service data
    - Federation services
    - Service to extend on-premise AD DS environment to the public cloud
    - Directory synchronization and SSO
    - APIs for developers

Windows Azure Active Directory is a cloud-based infrastructure as a service (IaaS) that you can use for identity management and access control. Windows Azure Active Directory has been the identity provider for Office 365™ since before Windows Azure Active Directory was made available to the public.

Implementing Windows Azure Active Directory is different than running virtualized domain controllers in Windows Azure. Windows Azure Active Directory is a service that enables administrators to manage their applications and identity services without having to manage the underlying hardware and software. Running virtualized domain controllers in Windows Azure is similar to running virtualized domain controllers in another data center or offsite location. Although both methods have their specific use cases, this lesson focuses on Windows Azure Active Directory.

Windows Azure Active Directory provides the following features:

- Active Directory authentication services in public or private clouds
- Cloud-based storage for directory service data
- Federation services
- A service for extending an on-premises Active Directory environment to cloud services

Windows Azure Active Directory provides the following benefits:

- High availability. Generally, the Windows Azure environment is much more highly available than what a typical company can deliver with on-premises services.
- Scalability. The Windows Azure environment can scale to meet very large workloads on demand.
- Integration with on-premises AD DS, including directory synchronization and single sign-on (SSO). This includes the ability to limit the data that synchronizes to Windows Azure Active Directory.
- Application programming interfaces (APIs). The Representational State Transfer (REST) API of Windows Azure Service Management provides developers the ability to perform management portal tasks programmatically. The Graph API allows developers to query directory data from their applications.

Windows Azure Active Directory Authentication

Windows Azure Active Directory authentication is different from Active Directory authentication. The protocols used are different for Windows Azure Active Directory than for AD DS. Often, Active Directory administrators are less experienced with the web-based authentication protocols that are used for Windows Azure Active Directory.

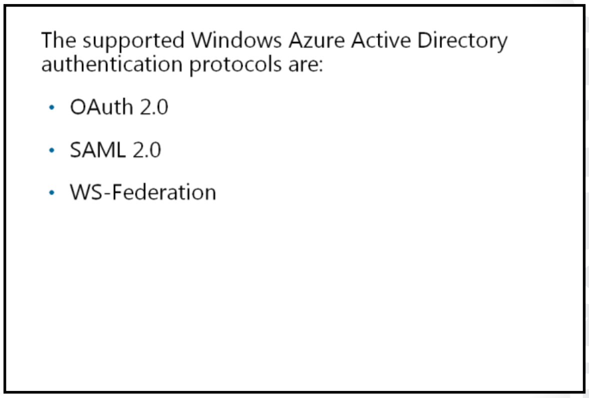

The supported Windows Azure Active Directory authentication protocols are:

- OAuth 2.0
- SAML 2.0
- WS-Federation

Windows Azure Active Directory supports a few different authentication protocols:

- OAuth 2.0. Based on RFC 6749, “The OAuth 2.0 Authorization Framework,” OAuth 2.0 is an open standard for authorization that provides granular access control to destination services. Access can be provided on a temporary basis. OAuth allows decoupling of authentication credentials, which prevents credentials from being passed to the destination.
- Security Assertion Markup Language 2.0 (SAML 2.0). SAML is an open standard XML protocol made up of security tokens and claims, which was originally introduced in the year 2002. A security token contains claims, which are typically Active Directory attributes that the workflow application uses to make decisions for authorization and access.
- WS-Federation. Web Services Federation (WS-Federation) is a security mechanism that allows identity federation so that users in one realm (or directory) can access resources in another realm. A consortium of technology companies, including Microsoft, created this protocol for use as an open standard on the Internet.

These supported protocols are web-based protocols intended for use on the Internet. However, Active Directory authentication protocols were designed for use on a private network, and initially without a need for open standards for authentication.

Options for Integrating Authentication with On-Premises AD DS

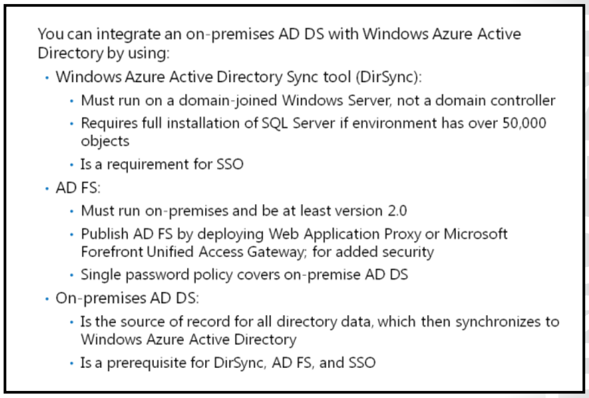

You can integrate an on-premises AD DS with Windows Azure Active Directory by using:

- Windows Azure Active Directory Sync tool (DirSync):
  - Must run on a domain-joined Windows Server, not a domain controller
  - Requires full installation of SQL Server if environment has over 50,000 objects
  - Is a requirement for SSO
- AD FS:
  - Must run on-premises and be at least version 2.0
  - Publish AD FS by deploying Web Application Proxy or Microsoft Forefront Unified Access Gateway; for added security
  - Single password policy covers on-premise AD DS
- On-premises AD DS:
  - Is the source of record for all directory data, which then synchronizes to Windows Azure Active Directory
  - Is a prerequisite for DirSync, AD FS, and SSO

There are numerous deployment scenarios for Windows Azure Active Directory that do not involve an on-premises Active Directory environment. However, for many organizations that have some services on the corporate network and some services in cloud services, integrating Windows Azure Active Directory and an on-premises AD DS provides a good end user experience.

You can integrate an on-premises Active Directory environment with Windows Azure Active Directory by using the following technologies:

- Windows Azure Active Directory Sync tool. This tool, commonly referred to as *DirSync*, has been in use since before the introduction of Windows Azure. DirSync provides directory synchronization from the on-premises AD DS to Windows Azure Active Directory. You use directory synchronization primarily to synchronize user objects and user attributes. It is a requirement for SSO. The tool is run on an on-premises domain-joined computer, and performs Windows Azure Active Directory synchronizations every three hours by default. DirSync is available for download on the Windows Azure portal. If you are synchronizing more than 50,000 objects, then a full installation of Microsoft SQL Server® is required.

- AD FS. AD FS is deployed on-premises and provides SSO for applications and services that reside on-premises or in Windows Azure. In addition, AD FS enables all authentications to take place in the on-premises Active Directory, and it offers multi-factor authentication. A Web Application Proxy or Microsoft Forefront Unified Access Gateway can securely extend AD FS to the perimeter network. AD FS must at least version 2.0.

- On-premises AD DS. The foundation for integration with Windows Azure Active Directory is the on-premises implementation of AD DS. AD DS is the authentication provider and the source of directory data. AD DS is a requirement for DirSync, AD FS, and SSO.

Designing Directory Synchronization

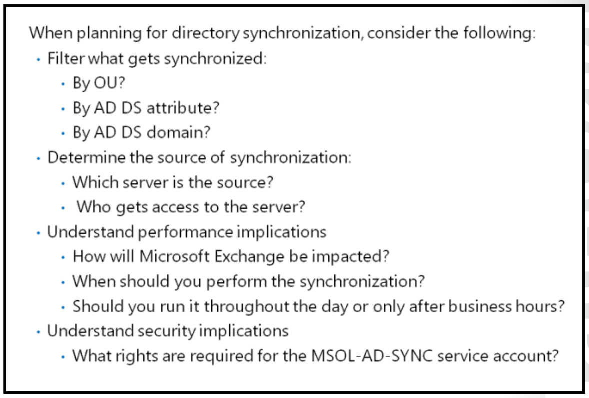

When planning for directory synchronization, consider the following:

- Filter what gets synchronized:
  - By OU?
  - By AD DS attribute?
  - By AD DS domain?
- Determine the source of synchronization:
  - Which server is the source?
  - Who gets access to the server?
- Understand performance implications
  - How will Microsoft Exchange be impacted?
  - When should you perform the synchronization?
  - Should you run it throughout the day or only after business hours?
- Understand security implications
  - What rights are required for the MSOL-AD-SYNC service account?

You must plan carefully for the directory synchronization between an on-premises Active Directory environment and Windows Azure Active Directory. Without proper planning, an organization could cause unnecessary security exposure, reduce performance, and create too much administrative overhead for ongoing support staff.

Although the DirSync tool is straightforward, there are some important design considerations to work through before using the tool for synchronization:

- Filtering what to synchronize. There are three ways to filter the Active Directory user objects that synchronize from the on-premises Active Directory domain to Windows Azure Active Directory:

- Filtering by organizational unit (OU). This is a straightforward way to filter. However, many organizations have an existing OU structure that is not suited for directory synchronization. This is often the case when you want a subset of user objects synchronized, but the OU in which the subset of user objects are located contains other Active Directory user objects that you do not want synchronized. Administrators need to ascertain whether an OU restructure is the best course of action, or whether another filtering method would be a better option.
- Filtering by domain. This is a lesser-used filtering method because synchronizing is per domain, and it is not common to synchronize all of the user objects from a domain to Windows Azure Active Directory. However, some organizations that have resource domains and user domains may be able to utilize domain filtering.
- Filtering by user object attributes. This is the most granular method of filtering. It allows for precise targeting of user objects. Some of the benefits of this method are being able to maintain the existing OU structure, being able to quickly synchronize a user object (whether existing or newly created), and being able to maintain strict control over which user objects are synchronized. You can populate a custom attribute with the attribute value **Sync to Windows Azure Active Directory** for all user objects that you want to synchronize. Then, you can set up a filter only to synchronize user objects that have the custom attribute populated with the **Sync to Windows Azure Active Directory** attribute value.

- Ascertaining which server will perform the synchronizing. You cannot install the DirSync tool on a domain controller. In addition, the operations that the DirSync tool performs are highly sensitive. Therefore, you should install the tool only on a highly secure server with access provided only to domain administrators or a similar small group of trusted administrators.
- Understanding performance implications. You need to decide the following: If your existing directory has a large number of user objects to synchronize, when should you perform the synchronization, and how will you monitor the performance of your on-premises environment during the initial synchronization and subsequent synchronizations? In general, you should perform the initial synchronization after business hours and only after successful testing of the synchronization in a pre-production environment.
- Understanding security requirements. When you configure directory synchronization, a service account named MSOL_AD_SYNC is created. You should have a good understanding of the service account, including the location in the on-premises directory, the rights that it requires, and how the password is managed.

Options for Managing Windows Azure Active Directory Accounts

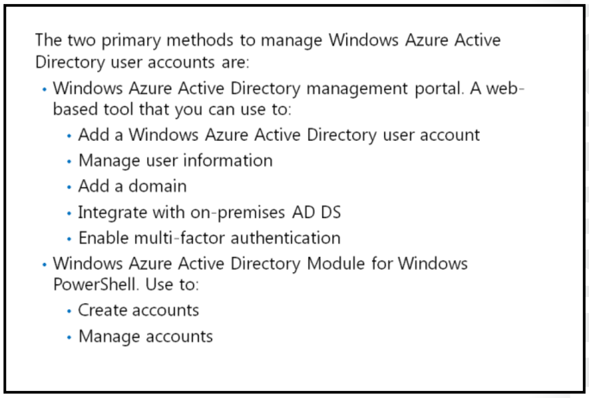

The two primary methods to manage Windows Azure Active Directory user accounts are:

- Windows Azure Active Directory management portal. A web-based tool that you can use to:
  - Add a Windows Azure Active Directory user account
  - Manage user information
  - Add a domain
  - Integrate with on-premises AD DS
  - Enable multi-factor authentication
- Windows Azure Active Directory Module for Windows PowerShell. Use to:
  - Create accounts
  - Manage accounts

You can manage Windows Azure Active Directory accounts in two different ways:

- Directory synchronization. When using directory synchronization, you manage user accounts by using your on premises management tools in the same manner as you would when you manage on premises Active Directory environments.
- Windows Azure. When not using directory synchronization, you manage your accounts in Windows Azure.

To manage user accounts that are only in the public cloud, Windows Azure Active Directory offers you the following two methods:

- Windows Azure Active Directory management portal. The Windows Azure Active Directory management portal offers a web-based tool that you can use to add and manage Windows Azure Active Directory user accounts. In this portal, you can add a user, manage user information, add a domain, integrate with on-premises AD DS, or enable multi-factor authentication.
- Windows Azure Active Directory Module for Windows PowerShell®. Windows PowerShell allows you to create and manage accounts using Windows PowerShell cmdlets.

Windows Azure Active Directory user accounts are simplistic when compared to on-premises Active Directory accounts. This is because Windows Azure Active Directory user accounts are limited to a core set of attributes, which you can only manage using a few select tools. When directory synchronization is in use, an indirect form of management uses DirSync to add or remove Active Directory accounts from Windows Azure Active Directory.

Lab A: Designing and Implementing an Active Directory Domain Services Forest Infrastructure

Scenario

The current Active Directory environment at A. Datum Corporation consists of a single Active Directory forest, which has only domain controllers running Windows Server 2008 R2. All domain controllers are deployed in a single AD DS site at the London head office's data center.

A. Datum wants to integrate its newly acquired companies, Contoso, Ltd, and Trey Research, into their organization. Contoso currently is running UNIX; however, Trey Research is currently running AD DS in Windows Server 2008 R2.

A. Datum also is planning to expand the number of employees who are located at the Contoso office in Paris. This is because the Paris office will become the primary sales, marketing, and delivery office for the company's aggressive expansion into European markets.

Additionally, A. Datum plans to deploy some applications and services for external clients. These resources will be located on a perimeter network, and should be independent from the company's Active Directory forest. However, the same administrators will administer the perimeter network resources.

Objectives

After completing this lab, you will be able to:

- Design an Active Directory forest infrastructure.
- Implement an Active Directory forest trust.

Lab Setup

Estimated Time: 40 minutes

Virtual machines	20413C-LON-DC1 20413C-TREY-DC1
User name	Adatum\Administrator TreyResearch\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 **20413C-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in using the following credentials:
   - User name: **Administrator**
   - Password: **Pa$$w0rd**
   - Domain: **Adatum**

5. In Hyper-V Manager, click **20413C-TREY-DC1**, and in the Actions pane, click **Start**.

6. In the Actions pane, click **Connect**. Wait until the virtual machine starts.

7. Sign in using the following credentials:

   - User name: **Administrator**
   - Password: **Pa$$w0rd**
   - Domain: **TreyResearch**

Exercise 1: Designing an Active Directory Forest Infrastructure

Scenario

The most important business requirement for integrating the new corporate acquisitions is the ease of collaboration. Trey Research currently has a single forest environment, and both Contoso and Trey Research currently have hosted email systems. As much as possible, all of the users within the merged companies should be able to interact with each other by using a wide variety of technologies. A second requirement in integrating the three companies is cost savings.

Supplemental documentation

Note: The following email is reproduced as a real email message with the initial messages at the end of the message thread. You must read the email thread from the bottom up.

Email from Patricia Doyle at Trey Research:

Brad Sutton

From:	Patricia Doyle [Patricia@TreyResearch.net]
Sent:	12 Sep 10:15
To:	Brad@Adatum.com
Subject:	Re: Trey Research Forest Design

Brad,

I'm glad to help. Here's the background.

Trey Research is based in London, Cambridge, and Boston in the US. You might find the following table useful. It gives a breakdown of user distribution across our three locations:

Location	Function	Characteristics
London, England	Head office	Current employees: 1,200 (Planned employees: 2,000)
Cambridge, England	Research facility	Employees: 750
Boston, USA	Research facility	Employees: 1,250

We specialize in pharmaceutical research. Some of our research is extremely confidential, and it is critical (and legally required) that the research data is not accessible to anyone but authorized users within Trey Research. As a result, the research department at Trey Research has its own subdomain: research.treyresearch.net.

We're using Windows Server 2008 R2 to provide standard file and print services, and our existing forest and domain infrastructure is based on Windows Server 2008 R2 domain controllers.

There is one other consideration. We have a research application that required some Active Directory schema changes. I am not sure if that would impact any decisions you might make in regards to the technical side of the merger.

Looking forward to your visit next week,

Patricia

----- Original Message -----

From:	Brad Sutton [Brad@Adatum.com]
Sent:	11 Sep 17:05
To:	Patricia@TreyResearch.net
Subject:	Trey Research Forest Design

Patricia,

As you know, I mentioned on the telephone that we're just starting the early design phase of the integration project between ourselves and Trey Research. I'm doing a bit of fact-finding. Can you please provide me some details about the Trey Research business, the way it organizes its resources, and anything else you think might be helpful? I'm looking forward to my visit to your Cambridge offices next week. Hopefully, by then, I shall be in a position to provide some more information on how we see this merger moving forwards – at least from a technical standpoint.

Thanks,

Brad

Email from Charlotte Weiss at Contoso:

Brad Sutton

From:	Charlotte Weiss [Charlotte@Contoso.com]
Sent:	11 Sep 12:02
To:	Brad@Adatum.com
Subject:	Re: Contoso Active Directory Implementation

Brad,

Yes, I was aware of the imminent Active Directory rollout. You should be aware of the following facts regarding Contoso:

- We're still using the single domain name contoso.com, but as per your DNS design, we are going to be using contoso.com externally and ad.contoso.com internally.
- We have mostly managed our IT infrastructure centrally from the Paris offices, but sometimes also with some localized IT staff at the regional hubs in Rome, Barcelona, Munich, and Athens.
- Our IT staff is beginning on some classroom-based and online training for both the Windows Server operating system and Active Directory, but it will take them some time before they have the necessary skills.
- The suggestion by the management that we provide some email services internally and not exclusively hosted is gaining momentum, and given that A. Datum has an investment in an on-premises Exchange Server environment, I imagine that's the way we would seek to go.

I can't really tell you much more, but if you have any other questions, please give me a call.

Regards,

Charlotte

----- Original Message -----

From: Brad Sutton [Brad@Adatum.com]

Sent: 10 Sep 17:40

To: Charlotte@Contoso.com

Subject: Contoso Active Directory Implementation

Hi Charlotte,

Thanks for your help with the DNS project. As you know, things are moving forward with the mergers, and we must devise an Active Directory design for your part of the organization. I just wanted to give you a heads up that this project was starting.

Brad

A. Datum, Trey Research, and Contoso Forest Integration Strategy

Document Reference Number: BS00913/1

Document Author	Brad Sutton
Date	13th Sep

Requirements Overview

To design a forest infrastructure that addresses the needs of A. Datum and its two imminent acquisitions, Trey Research, and Contoso, Ltd:

- The design should be simple.
- The design should reflect the IT support needs of the respective parts of the merged A. Datum organization.
- The cost of the restructuring should be minimized.
- Trey Research security must not be compromised. It is imperative that only authorized research staff can access research resources.

Summary of Information

- A. Datum:
    - Has an Active Directory deployment based on a single forest (Windows Server 2012).
    - Has a hybrid Exchange Server deployment with elements both hosted and deployed to the various locations around A. Datum.
    - Perimeter network contains Exchange Server Edge Transport server roles that require some knowledge of Active Directory objects, such as user email addresses and other Exchange Server–related properties.
- Trey Research:
    - Has an Active Directory deployment based on a single forest (Windows Server 2008 R2).
    - Schema has been modified to suit the needs of a critical line-of-business (LOB) research application.
    - Has in-house Active Directory expertise.
    - Is vital that Trey Research be administratively isolated.

A. Datum, Trey Research, and Contoso Forest Integration Strategy

- Contoso:
  - Has no Active Directory deployment in place.
  - Has no Active Directory expertise, although staff are being trained; the environment would need to be managed initially from London.
  - Has no Active Directory–aware applications, although there is some likelihood of Exchange Server being deployed in the future (email currently is exclusively hosted).

Proposals

1. How many forests does the current deployment have?
2. Is that number of forests sufficient?
3. What forest design and forest trust design will enable collaboration between A. Datum, Contoso, and Trey Research? Are there any special requirements for this scenario?
4. How can you address the requirement to protect confidential data in Trey Research from unauthorized access?
5. How should you plan for the A. Datum perimeter server requirements?
6. Are there any alternatives to the forest design that you would consider?
7. What are the benefits and drawbacks of the alternative design, if any?
8. Draw a simple diagram of your proposed Active Directory forest design in the space provided.

The main tasks for this exercise are as follows:

1. Read the supporting documentation.

2. Update the Proposals section with your planned course of action.

3. Examine the suggested proposals in the Lab Answer Key.

4. Discuss your proposed solution with the class, as guided by your instructor.

▶ Task 1: Read the supporting documentation

- Read the documentation provided.

▶ Task 2: Update the Proposals section with your planned course of action

- Answer the questions in the Proposals section of the A. Datum, Trey Research, and Contoso Forest Integration Strategy document.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones in the Lab Answer Key.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you should have created a forest design that incorporates A. Datum Corporation, Trey Research, and Contoso, Ltd.

Exercise 2: Implementing Active Directory Forest Trusts

Scenario

You must now implement a portion of your Active Directory forest design, specifically, the design that integrates A. Datum and Trey Research forests. You must configure the type of trust, the direction of that trust, and the type of authentication for trust relationships between the two companies. First, you must create conditional forwarders between the treyresearch.net and adatum.com domains. Then, you must configure the forest trust as per your design.

The main tasks for this exercise are as follows:

1. Configure DNS to support the forest trusts.

2. Create the required forest trusts between A. Datum and Trey Research.

▶ Task 1: Configure DNS to support the forest trusts

1. Switch to LON-DC1, and, if necessary, sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.

2. Open the DNS management console, and create a new conditional forwarder with the following settings:

   o DNS Domain: **treyresearch.net**

   o IP address: **172.16.10.10**

3. Open a command prompt and verify that you can resolve names in the target DNS domain by typing the following command:

```
Nslookup trey-dc1.treyresearch.net
```

Note: The query should be successful, returning the IP address of 172.16.10.10.

4. Switch to TREY-DC1.

5. If necessary, sign in as **Treyresearch\Administrator** with the password **Pa$$w0rd**.

6. Open the DNS management console, and create a conditional forwarder with the following settings:

   o Domain name: **Adatum.com**

   o IP address: **172.16.0.10**

7. Open a command prompt and verify that you can resolve names in the target DNS domain by typing the following command:

```
Nslookup lon-svr1.adatum.com
```

Note: The query should be successful, returning the IP address of 172.16.0.11.

▶ Task 2: Create the required forest trusts between A. Datum and Trey Research

1. Switch to LON-DC1.

2. From Server Manager, open Active Directory Domains and Trusts.

3. View the Adatum.com domain’s properties.

4. Create a new forest trust with the following properties:
   - Trust name: **treyresearch.net**
   - Trust type: **Forest trust**
   - Direction: **Two-way**
   - Sides of trusts: **Both this domain and the specified domain**
   - Authentication: **Forest-wide authentication**
   - Confirm Outgoing Trust: **Yes**
   - Confirm Incoming Trust: **Yes**

Results: After completing this exercise, you should have successfully implemented part of the forest infrastructure strategy that you designed.

▶ Task: To prepare for the next lab

- Leave all the virtual machines running.

Lesson 4
Designing and Implementing Active Directory Domains

When designing your Active Directory domains, you must decide whether a single domain will support your organization's requirements. A single domain model is easier to implement and requires less administrative overhead than multiple domains. However, you can deploy multiple domains to minimize replication traffic and enable a higher degree of administration separation. You also need to consider whether your organization will benefit from deploying a forest root domain.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the design models for Active Directory domains.
- Describe the reasons for deploying multiple Active Directory domains.
- Deploy dedicated forest root domains.
- Design Active Directory domains.
- Implement an Active Directory domain.

Active Directory Domain Models

You must choose the best Active Directory domain model for your organization as soon as you begin the design process. When choosing the design model that you want to use for your Active Directory domain, you must consider several factors, including:

- The number of users that your organization has.
- Administrative needs.
- Your company's organizational structure.
- The network bandwidth capacity between locations, if applicable.

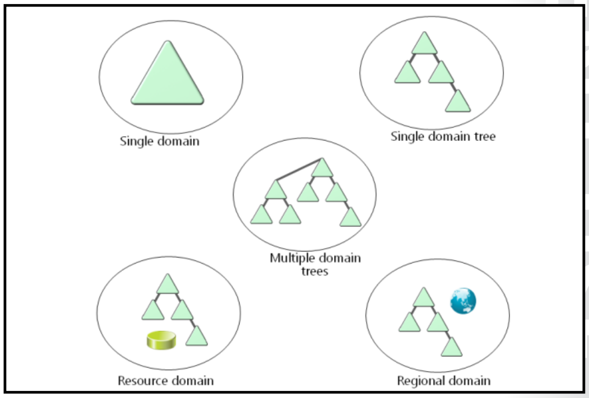

You also must be aware of the boundaries that Active Directory domains provide:

- A domain acts as a boundary for replication of Active Directory domain partitions.
- A domain acts also as a boundary for the application of Group Policy.
- Domains provide administrative boundaries, because a domain administrator for one domain in a forest cannot administer another domain in the forest without being elevated to the Enterprise Admins group.

Note: Domains do not provide security boundaries, because all domains in a forest trust each other. Additionally, the enterprise administrator account that resides at the forest root domain has administration privileges for all forest domains.

Note: Although you can rename domains, this is not an easy process. Therefore, when you start to deploy your domains, be sure that you specify the correct domain names.

Domain Models

The design models for domains that organizations most commonly use include the following:

- *Single domain*. The single domain model consists of a forest with a single domain. Any domain controller can authenticate any user in the forest, and all domain controllers can be global catalog servers. In this model, all Active Directory data replicates to all locations that host domain controllers. The single domain model is the simplest domain design, because it is easier to administer and less expensive to maintain. However, the single domain model creates the most replication traffic, particularly for organizations that locate domain controllers at multiple sites.
- *Single domain tree*. This model consists of a forest with one or more Active Directory domains that share the same namespace. All domains that you add to the domain tree become child domains to the forest root domain. For example, assuming you have a parent domain named contoso.com, you can add domains with DNS names such as sales.contoso.com and hr.contoso.com.
- *Multiple domain trees*. This model consists of a forest with multiple domain trees that have their own Active Directory namespaces. Using additional domain trees does not provide additional functionality, because domains behave in the same way regardless of whether you use a single namespace or multiple namespaces. A domain tree simply provides a way to organize and name domains in the forest. For example, you can have both the contoso.com and adatum.com domain trees in the same forest. An Active Directory domain tree contains at least one domain. Domain trees can have additional child domains.
- *Regional domain*. The regional domain model consists of a forest with one or more regional domain trees, each of which represents geographic locations within an organization. The regions that define each domain in the regional domain model typically represent static elements, such as countries or continents. Network connectivity is a key factor when planning to use a regional domain model. The regional domain model is more complex to design, and it requires a thorough analysis of the wide area network (WAN) connectivity and each region's number of users. However, because all object data within a domain replicates to all domain controllers in that domain, using regional domains can reduce network traffic over the WAN link.
- *Resource domain*. Resource domains store company resources, such as servers and printers. Resource domains typically do not contain user accounts.

Note: Fine-grained password and account lockout policies also can affect the domain design model that you select. In versions of the Windows Server operating system prior to Windows Server 2008, you could apply only one password and account lockout policy to all domain users by specifying it in the domain's Default Domain Policy. As a result, you could create different password and account lockout settings for different sets of users, but you had to either create a password filter or deploy multiple domains. Since the launch of Windows Server 2008, you can use fine-grained password policies to specify multiple password policies, and to apply different password restrictions and account lockout policies to different sets of users within a single domain.

Reasons for Deploying Multiple Domains

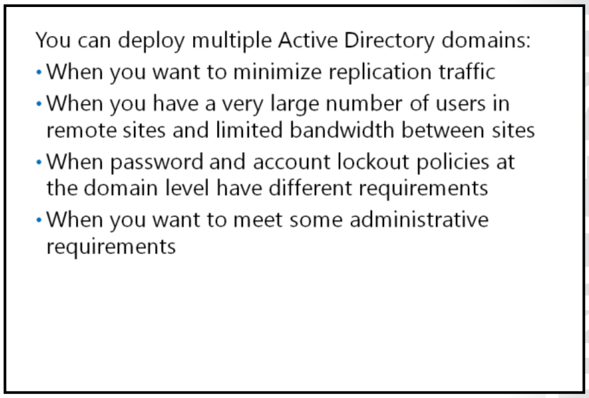

You can deploy multiple Active Directory domains:

- When you want to minimize replication traffic
- When you have a very large number of users in remote sites and limited bandwidth between sites
- When password and account lockout policies at the domain level have different requirements
- When you want to meet some administrative requirements

There are several reasons for deploying multiple Active Directory domains in your organization. For example, you may want to use multiple Active Directory domains when:

- You want to minimize replication traffic between physical locations. The domain directory partition replicates to all domain controllers in a domain. If you have users in separate geographic regions with WAN links to the network, it is a good idea to deploy a separate domain for these users. This helps to reduce replication traffic because domain controllers only replicate the content in the Active Directory schema and configuration partitions. Changes within these partitions are relatively rare, so replication traffic is reduced.
- You have an extremely large number of users (more than 100,000 users) in a remote site, and have very limited bandwidth between sites. However, this is a unique situation. Typically, administrators create domains for business purposes rather than purely for technical limitations.
- You want to meet the requirements for diverse password policies and account lockout policies at the domain level. Even though Windows Server 2008 and Windows Server 2012 support multiple password policies and account lockout policies for users in the same Active Directory domain, there may be situations that require diverse policies at the domain level. An example of this is a non-Microsoft password filter that only a subset of the company should use.
- You want to meet specific administrative requirements. For example, an organization may have a set of Active Directory administrators for a specific region, and may need to limit their scope to those regions only. Using multiple domains, you can add these administrators to the Domain Admins group for their particular domains.

Considerations for Determining the Number of Domains to Use

Unless you are implementing the single domain design model, you must determine the number of domains that your organization requires. This will vary depending on the domain design model that you select. Commonly, business reasons such as organizational layout, legal requirements, and political boundaries determine the required number of domains. When deciding on the number of domains that you will need, you also should consider the number of users and the available bandwidth between your company's physical locations. These factors can influence your decision greatly regarding whether you deploy one or more domains. The following factors determine the maximum number of users that a single domain forest can contain:

- The slowest link that must accommodate replication between domain controllers
- The available bandwidth that you want to allocate to AD DS

Considerations for Deploying Dedicated Forest Root Domains

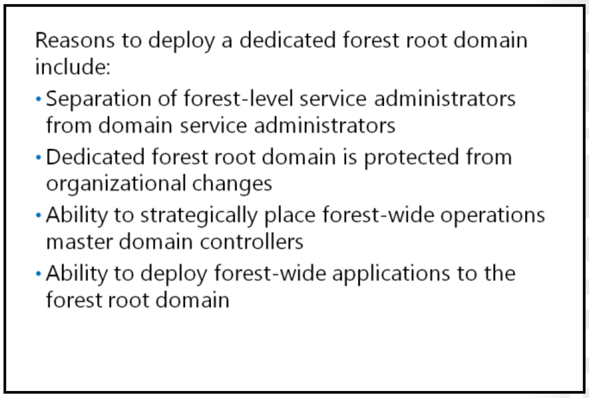

Reasons to deploy a dedicated forest root domain include:

- Separation of forest-level service administrators from domain service administrators
- Dedicated forest root domain is protected from organizational changes
- Ability to strategically place forest-wide operations master domain controllers
- Ability to deploy forest-wide applications to the forest root domain

A *dedicated forest root domain* is an Active Directory domain that you create exclusively to function as the forest root domain. A dedicated forest root domain typically does not contain end-user accounts, except some administrative accounts that are necessary to perform forest-wide tasks such as modifying schema and adding new domains.

The first domain that you deploy in an Active Directory forest is the forest root domain. After you deploy the forest root domain, this domain remains the forest root domain for the life of that Active Directory deployment. You cannot change the forest root domain, and you cannot designate any other domain in AD DS as the forest root domain. This is why it is particularly important to consider the design of your forest root domain during your project's design and planning stages.

Designing the forest root domain involves determining whether you need to deploy a dedicated forest root domain. That is, you need to decide if you will deploy an additional domain to accommodate users and resources while keeping the forest root domain for only forest-wide tasks and resources.

The benefits of deploying a dedicated forest root domain include the following:

- Forest-level service administrators are kept separate from domain service administrators.
- Organizational changes typically will not affect dedicated forest root domains.
- You have the ability to place domain controllers that have forest-wide operations master roles in more than one data center. Because these roles are unique and critical, you can distribute them in multiple locations. This leads to a simpler disaster recovery. Additionally, because the dedicated forest root domain does not contain end users, groups, or computers, replicating domain partitions will not consume much bandwidth.
- You can configure forest-wide applications (such as Exchange Server) to the forest root domain. Specifically, you create the administrative groups for Exchange Server or Microsoft Lync® Server in the forest root domain.

Using a dedicated forest root domain, however, introduces additional management overhead. The additional overhead includes tasks such as installing updates on domain controllers, troubleshooting problems, maintaining the hardware or virtual machines, and monitoring the environment.

If you choose not to use a dedicated forest root domain, then you must select a regional domain to function as the forest root domain. This regional domain will be the first domain in the forest that you deploy. Using a regional domain as a forest root domain does not generate additional management overhead, because you do not have to support an additional domain as a forest root domain.

Guidelines for Designing Active Directory Domains

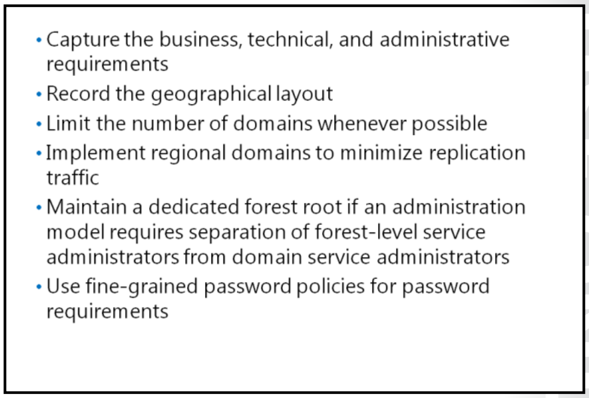

- Capture the business, technical, and administrative requirements
- Record the geographical layout
- Limit the number of domains whenever possible
- Implement regional domains to minimize replication traffic
- Maintain a dedicated forest root if an administration model requires separation of forest-level service administrators from domain service administrators
- Use fine-grained password policies for password requirements

When you are designing your infrastructure, the design of your Active Directory domains is an important step. Many factors can influence the Active Directory domain design, so you should complete a detailed analysis and plan before you begin your deployment.

Use the following guidelines when you design Active Directory domains:

- Investigate the business, technical, and administrative requirements for your organization. These requirements can greatly influence the domain design, so, you must include them in the early phases.
- Record the geographical layout of the network and the number of users at each location. You should have detailed information about the physical network layout, available bandwidth between locations, and the number of users in each location. These factors determine whether you need to deploy separate domains for each location, and how you will place the domain controllers.
- If you deploy multiple domains, limit the number of domains as much as possible. More domains cause additional administrative overhead, so you should try to minimize the total number of domains that require support.
- Use regional domains when the network topology includes diverse geographical locations. If you have enough IT staff in each geographical location, and if you have limited bandwidth between locations, deploy the regional domain model.
- Deploy a dedicated forest root domain if the administration model requires the separation of forest-level service administrators from domain service administrators.
- Use fine-grained password policies to define different password requirements. Instead of deploying another domain for this purpose, it is easier to use the fine-grained password policy feature to define password requirements for various Active Directory groups or users.

Demonstration: Implementing an Active Directory Domain

This demonstration shows how to:

- Add the Active Directory server role.
- Create a new domain in an existing forest.

Demonstration Steps

Add the Active Directory server role

1. Switch to CON-SVR and if necessary, sign in as **Administrator** with a password of **Pa$$w0rd**. This computer is a stand-alone server running Windows Server 2012.
2. Add the DNS and AD DS roles by using default options.

Create a new domain in an existing forest

1. In Server Manager, click the **Promote this server to a domain controller** option.
2. In the Active Directory Domain Services Configuration Wizard, on the **Deployment Configuration** page, click **Add a new domain to an existing forest**. Specify the following settings:
   - Domain type: **Tree Domain**
   - Forest name: **Adatum.com**
   - New domain name: **contoso.com**
   - User name: **Adatum\Administrator**
   - Password: **Pa$$w0rd**
   - Recovery password: **Pa$$w0rd**
3. CON-SVR will restart. After restarting, sign in as **Contoso\Administrator** with the password **Pa$$w0rd**.

Lesson 5
Designing DNS Namespaces in Active Directory DS Environments

As part of your Active Directory domain design, you must consider how your domain design will integrate with the public DNS namespace. You need to determine whether the internal and external namespaces will be the same or different for the Active Directory domain. Additionally, if you implement DNS servers in your organization, you should ensure that the internal and external namespaces are hosted on separate DNS servers.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe AD DS and DNS integration.
- Design an Active Directory namespace.
- Design DNS application partitions.
- Implement DNS servers into Active Directory environments.

AD DS and DNS Integration

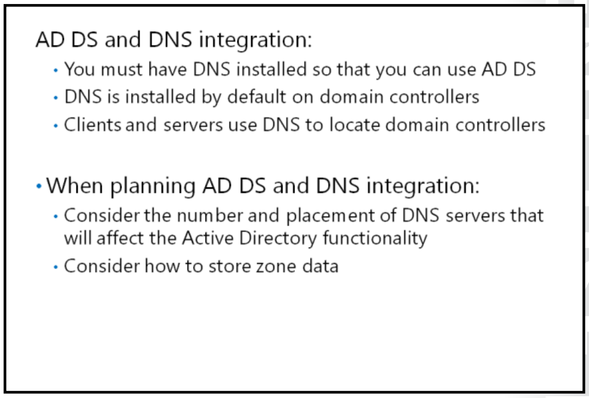

AD DS and DNS integration:

- You must have DNS installed so that you can use AD DS
- DNS is installed by default on domain controllers
- Clients and servers use DNS to locate domain controllers

- When planning AD DS and DNS integration:
  - Consider the number and placement of DNS servers that will affect the Active Directory functionality
  - Consider how to store zone data

AD DS is closely associated with DNS. This means that you must have DNS services available so that you can install and use AD DS. Integrating DNS and AD DS is essential because all clients and servers must use DNS to locate a domain controller so that users can log on to a domain and use the Active Directory services. Computers locate domain controllers and services by using:

- Host (A) resource records. The host (A) resource record contains the fully qualified domain name (FQDN) and IP address for the domain controller.
- Service (SRV) resource records. The service (SRV) resource record contains the FQDN for the domain controller and the name of the service that the domain controller provides.

By default, when you install an Active Directory domain controller, DNS also installs automatically. This integration between AD DS and DNS means that you must plan for both the DNS design and the Active Directory domain design. The number and placement of DNS servers can greatly influence AD DS functionality and performance.

One of the most important decisions that you must make when planning for DNS is how to store DNS zone data. After installing AD DS, use one of the following two methods for storing and replicating your zones when operating the DNS server on the new domain controller:

- Use a text-based file for standard zone storage
- Use the Active Directory database for directory-integrated zone storage

For networks deploying DNS to support AD DS, as best practice use a directory-integrated primary zone, which offers the following benefits:

- Provides multimaster updates and enhances security based on available Active Directory capabilities.
- Enables you to replicate and synchronize zones to new domain controllers automatically, whenever you add a new domain controller to an Active Directory domain.
- Enables you to streamline database replication planning for your network.
- Provides faster and more efficient directory replication than standard DNS replication.

Options for Designing an Active Directory Namespace

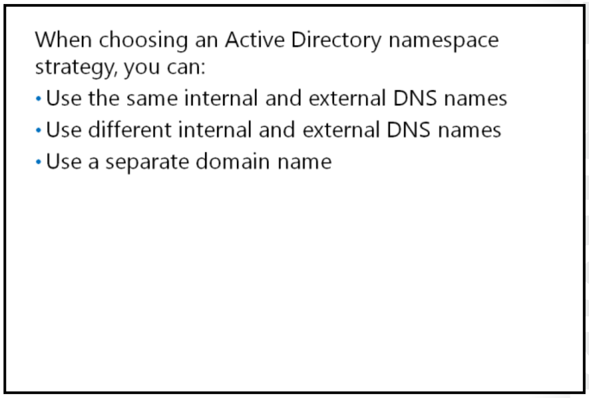

When choosing an Active Directory namespace strategy, you can:

- Use the same internal and external DNS names
- Use different internal and external DNS names
- Use a separate domain name

An Active Directory domain must have a DNS domain name. Because you also use DNS as a globally available, standards-based namespace, you should consider carefully where in the namespace you will position your Active Directory domain.

Currently, the separation between internal and external networks is relatively insubstantial. Because of this, it is often difficult to maintain namespace separation, and the separation provides less value than in previous iterations of AD DS. For this reason, many organizations use the most familiar domain name—the public domain name—for both the internal and external namespaces. The public domain name is most closely associated with an organization, and it typically is easiest to type. There are steps that you must take to support this configuration, but the cost is typically far less than the benefits provided by this configuration.

Regardless of the namespace that you choose, you must manage name resolution, perimeter protection, and security. This ensures equivalent levels of administrative effort to support any of these namespace choices. Therefore, use a DNS name that is easy for your users to type and remember.

In early Active Directory versions, organizations sometimes used a custom top-level domain, such as .msft, or the .local top-level domain for the Active Directory domain. Due to changes in the networked world (including IPv6 and increased interconnectivity), only explore these options after carefully considering the benefits that they might provide, their cost in terms of administration and user support, and their ability to support your business requirements.

Designing DNS Application Partitions

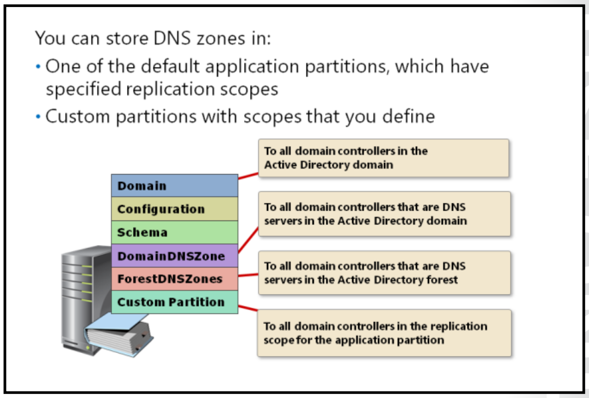

You can store DNS zones in:

- One of the default application partitions, which have specified replication scopes
- Custom partitions with scopes that you define

Domain — To all domain controllers in the Active Directory domain

Configuration

Schema

DomainDNSZone — To all domain controllers that are DNS servers in the Active Directory domain

ForestDNSZones — To all domain controllers that are DNS servers in the Active Directory forest

Custom Partition — To all domain controllers in the replication scope for the application partition

The DNS installation process creates two default application partitions: the domainDNSzones application partition, and the forestDNSzones application partition. Domain controllers within a domain that have the DNS service installed automatically receive a copy of the domainDNSzones application partition. All domain controllers within the forest—if they have the DNS service installed—receive a copy of the forestDNSzones application partition. However, if you have implemented DNS in your environment already, and if you use the existing DNS servers for AD DS, the Active Directory installation will not create the default application partitions.

You can create additional application partitions to store information. When you create an application partition, you must define which of the forest's domain controllers will participate in its replication. To create application partitions and enlist servers to replicate application partitions, use the Dnscmd.exe tool or the Ntdsutil.exe command-line tool.

When using Active Directory–integrated zones, you can control which domain controllers receive a zone by using AD DS partitions. You can also define which domain controllers within your Active Directory forest receive a copy of a given application partition. This helps reduce replication traffic by allowing AD DS to replicate the zone data only to domain controllers that require the information.

Specifying the Replication Scope

You can specify the replication scope when you create an Active Directory–integrated zone, or you can change the scope later. You can replicate to the following locations:

- All DNS servers in the Active Directory forest. The forestDNSzones application partition stores the zone. All domain controllers in the forest—if they have DNS installed—receive a copy of the zone. We recommend this configuration for zones that all clients need access to throughout the Active Directory forest. For example, the _msdcs zone includes information about global catalog servers and domain controllers, to which hosts anywhere in the forest may require access. You can store this zone in the forestDNSzones partition if your forest includes multiple domains and locations.
- All DNS servers in the Active Directory domain. The domainDNSzones application partition stores this zone. Only domain controllers in the same domain on which you install the DNS service receive a copy of the zone.
- All domain controllers in the Active Directory domain. The domain partition stores this zone, and all domain controllers in the domain receive a copy of it, even if you do not install the DNS service on them. This may cause unwanted replication traffic.
- All domain controllers that you specify in the replication scope of the specified application directory partition. The domain controllers that receive a copy of the application partition will receive a copy of the zone. You must create the application partition in advance.

Guidelines for Implementing DNS Servers into Active Directory Environments

As a best practice, use the Windows Server–based DNS servers whenever possible, and configure the AD DS and DNS zones as Active Directory–integrated zones. This enables you to avoid configuring a separate DNS replication topology that uses ordinary DNS zone transfers, because all zone data replicates automatically by means of AD DS replication. Using Active Directory–integrated zones also simplifies the process of deploying DNS, because multiple masters are created for DNS replication, which provides the following advantages:

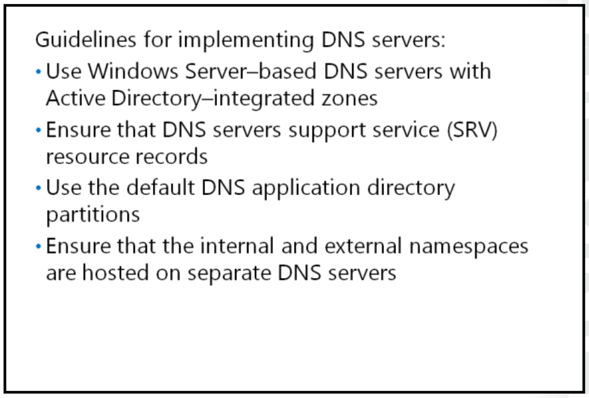

Guidelines for implementing DNS servers:

- Use Windows Server–based DNS servers with Active Directory–integrated zones
- Ensure that DNS servers support service (SRV) resource records
- Use the default DNS application directory partitions
- Ensure that the internal and external namespaces are hosted on separate DNS servers

- Any domain controller in the domain that is running the DNS server service can write updates to the Active Directory–integrated zones for the domain name for which they have authority. You do not need a separate DNS zone transfer topology.
- Active Directory–integrated zones support secure dynamic updates. These dynamic updates prevent unauthorized computers from overwriting existing names in DNS.
- Active Directory–integrated zones include the ForestDNSZones and DomainDNSZones zones. You should ensure that you use these partitions, because this will reduce replication traffic and the amount of data that the global catalog has to store.

If you have an existing DNS infrastructure in place and you must use that infrastructure, ensure that the existing infrastructure supports service (SRV) resource records, which are required by both AD DS and related services and applications.

You should host internal DNS namespaces and external DNS namespaces on separate DNS servers. This ensures that your internal DNS records are not visible on the Internet and are not subject to external compromise.

Note: Exposing internal DNS records on the Internet is a significant security risk. Potential attackers can use this information to determine which internal servers they will try to attack. For example, attackers may use this information to identify database servers or domain controllers, and then stage their attack on these servers.

Lesson 6
Designing Active Directory Domain Trusts

Active Directory domain trusts enable users in one domain to access resources in another domain. When you install multiple domains in the same forest, a default trust configuration links all domains. You can configure additional trust relationships by implementing shortcut trusts within a forest, and external trusts between domains in separate forests. For your Active Directory domain design, you should choose the domain trust that satisfies your organizational requirements.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe trust relationships.
- Describe shortcut trusts.
- Describe external trusts and realm trusts.
- Design Active Directory domain trusts.

Trust Relationships

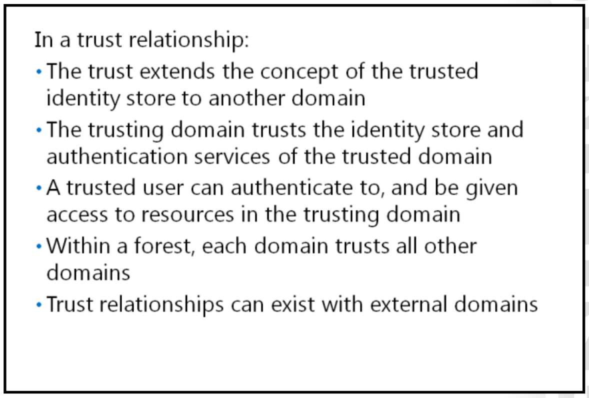

In a trust relationship:

- The trust extends the concept of the trusted identity store to another domain
- The trusting domain trusts the identity store and authentication services of the trusted domain
- A trusted user can authenticate to, and be given access to resources in the trusting domain
- Within a forest, each domain trusts all other domains
- Trust relationships can exist with external domains

When you implement a scenario involving two or more Active Directory domains, you typically are working with trust relationships, or *trusts*. Before designing trusts between domains in the same forest or between domains in separate forests, you should understand the purpose, functionality, and configuration of trust relationships.

Trust Relationships Within a Domain

You do not necessarily have to establish a trust relationship between two domains or two forests. Frequently, organizations establish trusts only between computers and domains. In a workgroup, the computer maintains an identity store in the Security Accounts Manager (SAM) database. The computer then authenticates users against that identity store, and secures system resources only with identities from the SAM database. When the computer joins a domain, it forms a trust relationship with the domain. As a result, the computer allows the authentication services and identity store of the Active Directory domain to authenticate users, rather than by the local system and its local identity store.

The domain member also allows you to use domain identities to secure system resources. For example, you can add Domain Users to the local Users group, which provides Domain Users with the right to log on locally to the system. In addition, you can add domain user and group accounts to ACLs on files, folders, registry keys, and printers. All domain members have similar trust relationships with the domain, enabling the domain to be a central identity store and a centralized service that provides authentication.

Trust Relationships Between Domains

You can extend the concept of trust relationships to other domains. A trust relationship between two domains enables one domain to trust the authentication service and the identity store of another domain, and to use those identities to secure resources. In effect, a trust relationship is a logical link established between domains to enable pass-through authentication.

There are two domains in every trust relationship: a trusting domain, and a trusted domain. The trusted domain contains the identity store, and provides authentication for users in that identity store. The *trusting domain* is a domain that trusts the identity store and authentication services of the trusted domain. When a user in the directory of the trusted domain logs on to, or connects to a system in the trusting domain, the trusting domain cannot authenticate that user because the user is not in its data store. Therefore, it passes the authentication to a domain controller in the trusted domain. The trusting domain, therefore, trusts the trusted domain to authenticate the user's identity. The trusting domain extends trust to the authentication services and the identity store of the trusted domain.

Because the trusting domain trusts the identities in the trusted domain, the trusting domain can use the trusted identities to grant access to resources. It can grant users in a trusted domain user rights such as the right to log on to workstations in the trusting domain. Additionally, it can add users and global groups in the trusted domain, to domain local groups in the trusting domain, and then grant them permissions to shared folders by adding the identities to ACLs in the trusting domain.

As an example, if Domain A trusts Domain B, then Domain A becomes the trusting domain and Domain B becomes the trusted domain. If a user in Domain B connects to or logs on to a computer in Domain A, Domain A will pass the authentication request to a domain controller in Domain B. Domain A also can use the identities from Domain B—users and groups, for example—to grant user rights and resource access in Domain A. Therefore, you can add a user or group in Domain B to an ACL on a shared folder in Domain A, or you can add a user or group in Domain B to a domain local group in Domain A.

Within a forest, all domains trust each other because each tree's root domain in a forest trusts the forest root domain, and each child domain trusts its parent domain. You should never delete trusts that are created automatically, as they are transitive and two-way. The net result is that a domain trusts the identity stores and authentication services of all other domains in its forest. The domain can:

- Add users, global groups, and universal groups from any domain in the forest to domain local groups.
- Grant user rights to users, global groups, and universal groups from any domain in the forest.
- Add users, global groups, and universal groups from any domain in the forest to ACLs on resources in any other domain in the forest.

You must establish trusts manually to other forests and domains outside the forest.

Shortcut Trusts

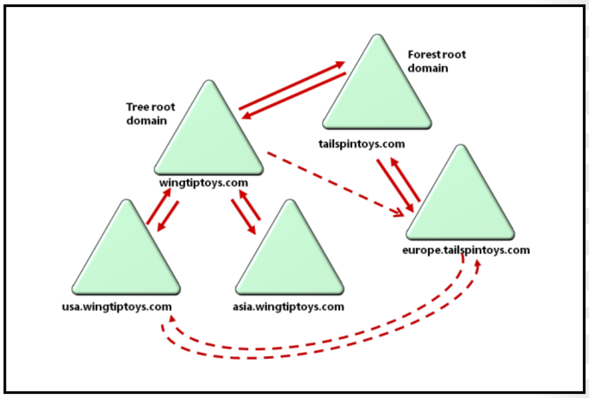

When designing trusts in a multidomain environment, you must understand the trust path between two domains, especially when you have two or more domain trees in a same forest.

When a user from a domain in one domain tree within a forest wants to access resources in a domain in another domain tree within the same forest, many steps occur to grant the session ticket that allows access to the target domain's resources. Most of these steps involve referrals to domains on the trust path between the user's domain and the shared folder's domain. When a user from a domain logs on to a computer in another domain, the authentication request must also traverse the trust path. This can affect performance, and if a domain controller is not available in a domain along the trust path, the client will not be able to authenticate to, or access the service.

To overcome these problems, you can use shortcut trusts to create a trust relationship directly between child domains in the forest trust path. However, you should note that shortcut trusts do not increase responsiveness unless the underlying physical network supports the shortcut path.

Shortcut trusts eliminate the forest trust path and therefore eliminate the time necessary for ticket requests to traverse that path. This significantly improves the performance of session ticket requests.

One-Way or Two-Way Shortcut Trusts

Shortcut trusts are transitive trusts, in that they can be one-way trusts or two-way trusts. In the illustration on the slide, a one-way shortcut trust exists, whereby wingtiptoys.com trusts europe.tailspintoys.com.

When a user from europe.tailspintoys.com logs on to a computer in wingtiptoys.com or requests a resource in wingtiptoys.com, the request is referred directly to a domain controller in the trusted domain, asia.wingitiptoys.com. However, the reverse is not true. If a user in wingtiptoys.com logs on to a computer in europe.tailspintoys.com, the authentication request traverses the trust path up to tailspintoys.com and down to wingtiptoys.com.

The slide also illustrates a two-way shortcut trust between usa.wingtiptoys.com and europe.tailspintoys.com. Because users in both domains can be authenticated by, and can request resources from computers in the other domain, the shortcut trust path is used.

External Trusts and Realm Trusts

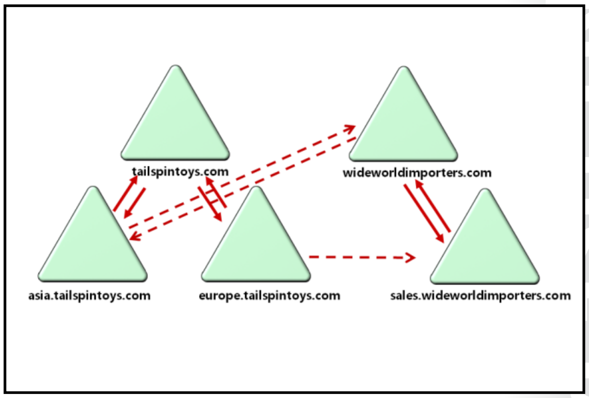

External Trusts

You create an external trust when you need to access a domain that is not in your forest. An *external trust* is a trust relationship between a domain in your forest and a Windows Server domain outside of your forest.

The slide displays a graphic of a one-way trust between the sales.wideworldimporters.com domain and the europe.tailspintoys.com domain. The Europe domain trusts the Sales domain, so users in the Sales domain can log on to computers in the Europe domain or connect to resources in the Europe domain.

The illustration also demonstrates a two-way trust between the wideworldimporters.com domain and the asia.tailspintoys.com domain. Users in each domain can access resources in the other domain. All external trusts are nontransitive, one-way trusts. If you need to create a two-way external trust, you will effectively create two one-way trusts, one in each direction.

When you create an outgoing external trust, AD DS creates a foreign security principal object for each security principal in the trusted domain. You then can add those users, groups, and computers to domain local groups or ACLs on resources in the trusting domain.

Realm Trusts

You create a realm trust between your domain and a realm. For example, if you need cross-platform interoperability with security services based on other Kerberos V5 implementations, you will create a realm trust between your domain and the Kerberos V5 realm. Like external trusts, realm trusts are one-way, and to create a two-way realm trust, you create two one-way trusts. By default, realm trusts are nontransitive, but you can make them transitive.

If a non-Windows Server Kerberos V5 realm trusts your domain, then the realm trusts all security principals in your domain. If your domain trusts a non-Windows Server Kerberos V5 realm, users in the realm can receive access to resources in your domain. However, the process is indirect. When a non-Windows Kerberos realm authenticates users, Kerberos tickets do not contain all the authorization data that Windows Server requires. Therefore, you should use an account mapping system. With an account mapping system, you create security principals in the Windows domain, and then map them to a foreign Kerberos identity in the trusted non-Windows Kerberos realm. The Windows domain uses only these proxy accounts to evaluate access to domain objects that have security descriptors. You can use all Windows Server proxy accounts in groups and on ACLs to control access on behalf of the non-Windows Server security principal.

Guidelines for Designing Active Directory Domain Trusts

When administering domain and forest trusts, the following best practices will enable you to increase availability, reduce the number of incidents, and make administrative tasks easier:

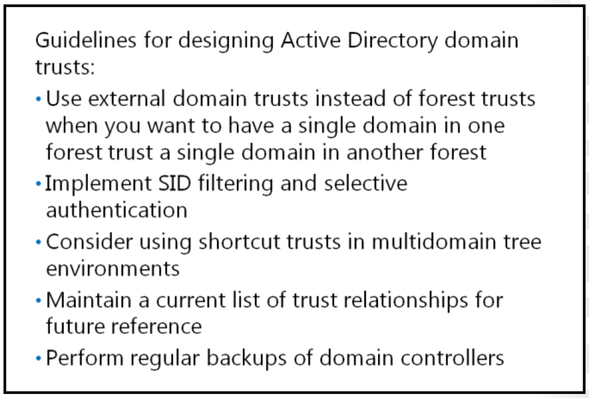

Guidelines for designing Active Directory domain trusts:

- Use external domain trusts instead of forest trusts when you want to have a single domain in one forest trust a single domain in another forest
- Implement SID filtering and selective authentication
- Consider using shortcut trusts in multidomain tree environments
- Maintain a current list of trust relationships for future reference
- Perform regular backups of domain controllers

- Use external domain trusts instead of forest trusts when you want to establish a trust relationship only between two specific domains in various forests. This approach guarantees a much lower security risk, because it focuses on only two specific domains in two separate forests. Also, use selective authentication to secure trust relationships.

  Creating an external or forest trust between two forests essentially provides a pathway for authentications to travel from the trusted forest to the trusting forest. Because the forest trust allows all secured communications to occur over this pathway, this action by itself does not necessarily create a threat to either forest. However, it does create a larger attack surface for malicious users in a trusted forest. You can set selective authentication on interforest trusts to help minimize this attack surface area.

  If you choose selective authentication, all users in the trusted domain are trusted identities. However, they can authenticate only for services on computers that you specify. For example, suppose you have an external trust with a partner organization's domain, and you want to ensure that only users from the marketing group in the partner organization 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.

- Implement SID filtering (also known as *Domain Quarantine*). In a trusted domain scenario, a rogue administrator could use administrative credentials in the trusted domain to load SIDs into a user's **sIDHistory** attribute that are the same as the SIDs of privileged accounts in your domain. That user would then have inappropriate levels of access to your domain's resources. SID filtering prevents this from occurring 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. Therefore, 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. Domain quarantine is enabled by default for all outgoing trusts to external domains and forests.

- When your forest contains domain trees with many child domains, if noticeable user authentication delays between the child domains are occurring, you can optimize the user authentication process between the child domains. You can do this by creating shortcut trusts to mid-level domains in the domain tree hierarchy.

Some other best practices include:

- Maintain a current list of trust relationships for future reference.
- Perform regular backups of domain controllers to preserve all trust relationships within a particular domain.

Lab B: Designing and Implementing an Active Directory Domain Infrastructure

Scenario

During the Active Directory forest design process at A. Datum Corporation, the design team members decided that they will need to maintain a separate forest for the treyresearch.net domain to fulfill the research department's isolation requirements. However, the design team currently is considering how best to integrate the Contoso, Ltd organization into the A. Datum network infrastructure. At this time, Contoso has not deployed AD DS.

Objectives

After completing this lab, you will be able to:

- Design an Active Directory domain infrastructure.
- Implement an Active Directory domain infrastructure.

Lab Setup

Estimated Time: 45 minutes

Virtual machines	20413C-LON-DC1
User name	Adatum\Administrator
Password	Pa$$w0rd

Virtual machines	20413C-TREY-DC1
User name	TreyResearch\Administrator
Password	Pa$$w0rd

Virtual machines	20413C-CON-SVR
User name	.\administrator
Password	Pa$$w0rd

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

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, click **20413C-CON-SVR**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.

4. Sign in using the following credentials:
   - User name: **Administrator**
   - Password: **Pa$$w0rd**

Exercise 1: Designing an Active Directory Domain Infrastructure

Scenario

Contoso, Ltd is based solely in Europe with its head offices in Paris, and is a former partner of A. Datum Corporation. The primary goal in purchasing Contoso is to integrate the companies' two product lines. You must consider the scenario and decide how best to deploy AD DS to Contoso, and how to integrate Contoso into the A. Datum organizational structure.

Contoso AD DS Integration Strategy	
Document Reference Number: BS00915/1	
Document Author Date	Brad Sutton 15th Sep
Requirements Overview Design an Active Directory domain infrastructure to support the following objectives: • The sales, marketing, and production groups in the two companies will be working closely together, and they must be able to share information easily. • The users in both the adatum.com domain and the Contoso organization must be able to access some information in the adatum.com forest. In addition, they must be able to access user mailboxes on Exchange servers that will be deployed in London, and access files on file share server LON-SVR1, which is located in London. The users should not be required to sign in with multiple accounts to access the files. • Some users in the Contoso organization require access to resources in the Trey Research organization. Your plan must facilitate this. The solution must not compromise the security of the treyresearch.net forest.	
Additional Information • A. Datum is planning to hire a large number of additional staff in Paris. These new employees will be working in the sales, marketing, and distribution departments in Paris. • Contoso has no Active Directory administrators, although the company's staff is being trained for this. • The Paris office is connected to London by one 6 Megabits per second (Mbps) link that is used for all communication and data sharing. The network team at A. Datum is concerned about bandwidth utilization between Paris and London. • A. Datum also has implemented a Voice over Internet Protocol (VoIP) and conferencing solution based on Lync Server 2010, and is planning to expand that deployment to include its London servers. The network team wants to ensure that the new Active Directory deployment uses as little bandwidth as possible for Active Directory–specific traffic.	
Proposals 1. Should you create a separate forest to accommodate the Contoso organization?	

Contoso AD DS Integration Strategy
2. If you decide to implement Contoso as part of the existing A. Datum forest, is it better to implement Contoso as a separate domain or as an OU in the existing A. Datum domain?
3. Assuming you choose a separate domain for Contoso, what Active Directory domain name will you use for Contoso? (Contoso already uses contoso.com for DNS purposes.)
4. Which is the forest root domain?
5. Is it a good idea to deploy additional domain controllers from the adatum.com domain in Paris? Why or why not?
6. How do you plan to address the requirement that users in Contoso need to access resources in the Trey Research organization?

The main tasks for this exercise are as follows:

1. Read the supporting documentation.

2. Update the proposal document with your planned course of action.

3. Examine the suggested proposals in the Lab Answer Key.

4. Discuss your proposed solution with the class, as guided by your instructor.

▶ Task 1: Read the supporting documentation

- Read the documentation provided.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the Proposals section of the Contoso AD DS Integration Strategy document.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones in the Lab Answer Key.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you should have successfully designed a domain infrastructure strategy for the integration of Contoso into the A. Datum organization.

Exercise 2: Implementing an Active Directory Domain Infrastructure

Scenario

The management team at A. Datum Corporation has approved your Contoso, Ltd integration strategy. You must now deploy a domain controller in Contoso, and add the contoso.com domain as a new domain in an existing forest.

The main tasks for this exercise are as follows:

1. Verify that the prerequisites for adding a new domain are satisfied.

2. Add CON-SVR as a domain controller in a new domain in an existing forest.

▶ Task 1: Verify that the prerequisites for adding a new domain are satisfied

1. Switch to CON-SVR, and if necessary, sign in as **Administrator** with a password of **Pa$$w0rd**. This computer is a stand-alone server running Windows Server 2012.

2. From Server Manager, add the DNS and AD DS roles by using default settings.

▶ Task 2: Add CON-SVR as a domain controller in a new domain in an existing forest

1. In Server Manager, in AD DS, run the **Promote this server to a domain controller** option.

2. In the Active Directory Domain Services Configuration Wizard, on the **Deployment Configuration** page, click **Add a new domain to an existing forest**.

3. Specify the following settings:

   - Domain type: **Tree Domain**

   - Forest name: **adatum.com**

   - New domain name: **contoso.com**

- User name: **Adatum\Administrator**
- Password: **Pa$$w0rd**
- Recovery password: **Pa$$w0rd**

4. CON-SVR will restart. When prompted, sign in as **Contoso\Administrator** with the password **Pa$$w0rd**.

Results: After completing this exercise, you should have successfully implemented a part of the domain infrastructure strategy that you devised.

▶ Task: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 through 4 for **20413C-TREY-DC1** and **20413C-CON-SVR**.

Module Review and Takeaways

Review Question(s)

Question: What is the purpose of the resource forest model?

Question: What forest functional level must you set in AD DS to be able to establish a forest trust?

Question: If you want to integrate multiple internal namespaces, which technologies would you use?

Question: A user from Contoso attempts to access a shared folder in the Tailspin Toys domain and receives an Access Denied error. A trust relationship between these two domains exists. What must you do to provide the user with access?

Module 7

Designing and Implementing an AD DS Organizational Unit Infrastructure

Contents:

Module Overview

In an Active Directory® domain you can create an organizational unit (OU) structure in which you can create a management infrastructure. This management infrastructure enables you to separate administrative tasks, delegate administrative permissions, and apply Group Policy Objects (GPOs). However, designing an OU structure and an Active Directory administrative tasks delegation model can be challenging. There are many possible design options, and often many administrative departments—such as domain administration, server and application operators, virtualization admins, user helpdesk, and decentralized services. In addition, you might have to modify your OU structure as business requirements or organizational structures change.

In addition to creating the correct OU structure, you must also determine which administrative departments or users should have control over which attributes and objects, and at which level. If you are redesigning the OU structure, you also must decide how to move the objects from the current structure into the new structure. Then you must plan the administrative task delegation model, decide how to implement it, and determine which groups to create to delegate the tasks. Finally, you must ensure everything is working as expected, and then define processes for future changes.

Objectives

After completing this module, you will be able to:

- Plan an Active Directory administrative tasks delegation model.
- Design an OU structure.
- Design and implement an Active Directory Domain Services (AD DS) group strategy.

Lesson 1
Planning the Active Directory Administrative Tasks Delegation Model

When designing your OU structure, planning the Active Directory administrative tasks delegation model is essential. Although you can use OUs to apply GPOs, partition objects, or represent your organization's business structure, the most important design consideration for the OU structure is the administrative tasks delegation model, which depends on the processes and administrative requirements in your organization.

In this lesson, you will learn which components are important in an Active Directory administrative tasks delegation model. You will also learn what information you must gather when designing an Active Directory OU structure and an administrative tasks delegation model.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe a typical Active Directory administrative tasks delegation model.
- Describe a typical IT administrative model.
- Describe the information that you must collect to design the OU structure and Active Directory administrative tasks delegation model.
- Explain how to design and implement an Active Directory administrative tasks delegation model.
- Describe considerations for branch office delegation.

What Is an Active Directory Administrative Tasks Delegation Model?

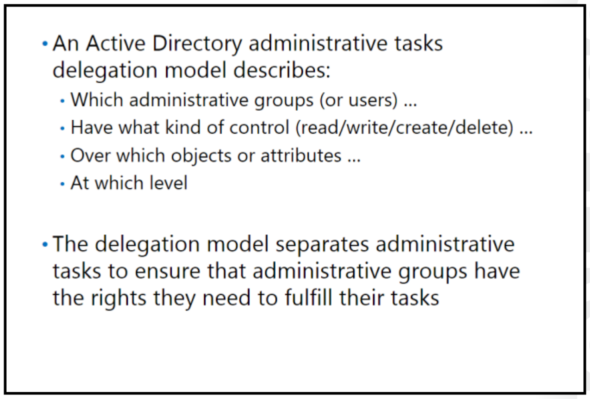

- An Active Directory administrative tasks delegation model describes:
  - Which administrative groups (or users) …
  - Have what kind of control (read/write/create/delete) …
  - Over which objects or attributes …
  - At which level
- The delegation model separates administrative tasks to ensure that administrative groups have the rights they need to fulfill their tasks

The Active Directory administrative tasks delegation model describes which groups can perform which tasks at which level. Unlike the forest and domain structure, the OU structure can separate administrative responsibilities within the same domain. For example, you can configure administrative groups in the following ways:

- Separate the administration of objects (such as users, groups, and computers) from the infrastructure administration (such as ensuring that AD DS is healthy and replication is functioning properly).
- Provide Service Accounts with the minimum required rights.
- Allow multi-tenancy in a single Active Directory domain.
- Allow application administrators to self-maintain group memberships for their applications.
- Allow owners of distribution lists to self-maintain the members.
- Delegate location information or phone numbers to local administrative staff in the branch office.
- Allow Branch Office Administrators to reinstall clients and rejoin them to the domain.

In an Active Directory domain, you can delegate administrative tasks and responsibilities. Because every OU contains security settings, you must determine carefully to whom you grant administrative permissions (whether to single users, or preferably to groups), which will allow them to administer objects or attributes of objects within a specific OU. The users or groups to whom you grant permissions do not have to be in the Domain Admins group, Server Operators groups, or any other built-in, high-privileged security group.

Examples of tasks that you can delegate include the following:

- Exercise full control over a certain type of object (for example, users or groups).
- Create certain objects, such as users, computers, or groups.
- Change a specific attribute of specific objects (for example, only the **telephone number** attribute of user objects, or the **members** attribute of group objects).
- Delete certain objects, such as computers.
- Link or unlink GPOs from a specific OU.
- Set send-as rights for email-enabled user accounts.
- Set Read permissions on Active Directory attributes that cannot be read by default, such as BitLocker® Drive Encryption recovery keys.
- Set Self-permissions, which define what information users or computers can update on their own accounts (for example, location or mobile phone numbers).

You should monitor and validate the Active Directory administrative tasks model regularly, and adjust them as your organization’s business requirements and processes change.

Additional Reading: For more information about Best Practices for Delegating Active Directory Administration, see the following:

- Best Practices for Delegating Active Directory Administration
  http://go.microsoft.com/fwlink/?linkid=279914
- Best Practices for Delegating Active Directory Administration Appendices
  http://go.microsoft.com/fwlink/?linkid=279915

Typical IT Administrative Models

When designing an Active Directory administrative tasks model and OU structure, you should consider several IT administrative models, depending on your corporate strategies. It is common to find mixtures of these models, and often two companies do not have the same model. The following are the most common types of IT administrative models:

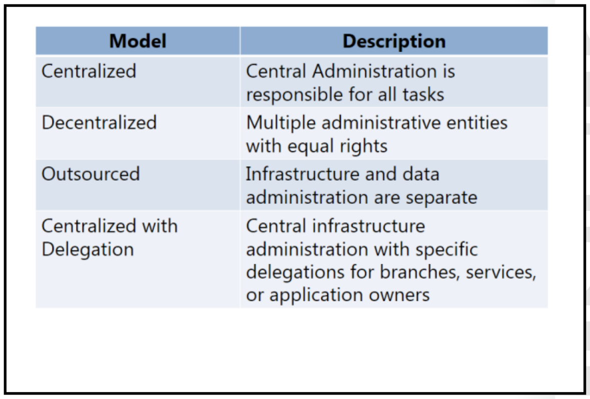

Model	Description
Centralized	Central Administration is responsible for all tasks
Decentralized	Multiple administrative entities with equal rights
Outsourced	Infrastructure and data administration are separate
Centralized with Delegation	Central infrastructure administration with specific delegations for branches, services, or application owners

- Centralized IT administrative model. In this model, all administration is centralized as much as possible. Decentralized sites often have no administrator, or only limited administration. Every change request or support issue passes through the central administration.

- Decentralized IT administrative model. In this model, companies with autonomous sites share a common Active Directory domain. Each site has its own administrative staff and works separately in their assigned OUs. Global design decisions, such as schema changes or other major infrastructure updates, require agreement across these entities.
- Outsourced IT administrative model. When Active Directory administration is outsourced, or when a company hosts multiple organizations in a single Active Directory domain, you should distinguish administrative responsibilities between the outsourcer and the outsourcing company. Sometimes the outsourcer is responsible for running the Active Directory infrastructure, while the outsourcing company is responsible for managing the data in AD DS (such as users, computers and groups). Sometimes it is the other way around, or sometimes both the infrastructure and data administration are outsourced.
- Centralized IT with Delegation administrative model. This is a very common model wherein centralized IT is responsible for maintaining the infrastructure and managing several of the administrative tasks. Delegated administrative groups fulfill the remaining administrative tasks, which may include the following:
    - Decentralized administrators are responsible for access to local file services for specific groups.
    - Decentralized administrators can unlock or reset passwords for users in their branch.
    - Decentralized administrators can create computer objects in their OUs.
    - The central group Virtualization Admins is responsible for joining virtual machines to the domain, and then assigning rights to specific groups to use the virtual machines.
    - Operators of applications can create and manage groups to control access to their resources—such as Microsoft® SharePoint® sites.

In addition to choosing the type of administrative model to use, when designing the Active Directory administrative tasks model you must also consider additional business or technical requirements. These requirements include applications such as fax or mobile device applications that require Send-As permissions, Voice over Internet Protocol (VoIP) applications (which must update certain information in AD DS), and service accounts for other applications that require rights in the Active Directory domain.

Gathering Information on Current Administrative Structures

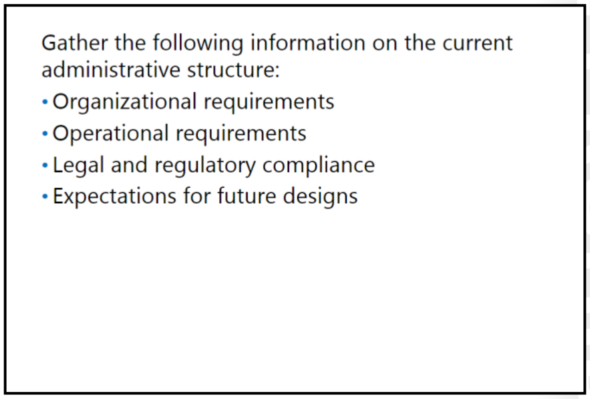
Gather the following information on the current administrative structure:

- Organizational requirements
- Operational requirements
- Legal and regulatory compliance
- Expectations for future designs

Prior to designing your Active Directory administrative tasks model, you must gather the required information. Sometimes, you must first define what this information is. Project sponsors often can authoritatively decide which administrative entities should be responsible for which tasks.

In the current administrative structure, you must gather the following information:

- Organizational requirements. Your organization's structure affects the design of your domain administration. As such, you must gather information about each geographic location, including the number of users and computers, whether there are hosting servers, which departments reside there, and whether there is local IT staff.

- Operational requirements. Your organization might have specific business requirements for security and for the isolation of specific resources. You might have a perimeter network (also called *DMZ* or *demilitarized zone*), and a screened subnet, where applications and servers provide services to users, customers, or vendors outside your corporate network. Certain applications in the environment might require specific rights, or you might allow certain users or groups to manage their applications or distribution lists by using self-service.
- Legal and regulatory compliance requirements. Depending on the type of business, your company might have legal or regulatory compliance requirements, such as the Health Insurance Portability and Accountability Act (HIPAA). These requirements regulate access to certain information, or the time required to recover from failures. Legal requirements apply to most companies, especially in the health, financial, and governmental sectors, but also to all companies that work with external data.
- Expectations for future designs. When you are tasked with designing or redesigning an administrative structure, you should be aware of any expectations about what the new design must accomplish or fix, and the requirements that it should fulfill. These requirements may include organizational growth, mergers, changing business requirements, administrative functions, or business reorganizations.

Gathering Information on Organizational Resources

When designing the administrative structure for the Active Directory domain, you must gather information about organizational resources. These resources can directly affect Active Directory objects.

The following table lists key resources about which you should gather information.

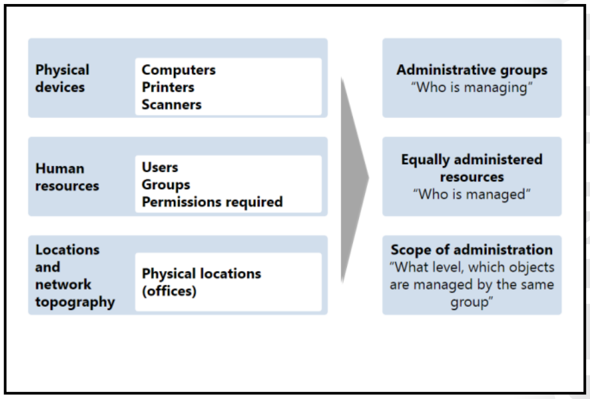

Resources	Active Directory objects
Physical devices including servers, workstations, and printers	Computer objects and printers
Employees, organization groups, project teams, security, and self-service requirements	OUs, users, groups, and permissions
Geographic locations and network topology	OUs, sites, subnets, and replication objects

After gathering the information about organizational resources, you must group them based on their administrative relevance:

- Administrative groups. You must identify the different groups that fulfill administrative tasks. These are not necessarily Active Directory groups at this stage, but groups of people who should have the same administrative rights over the same Active Directory objects at the same level.
- Equally administered resources. You must identify which objects are equally administered, that is, administered by the same group of people. You likely will place these objects into the same OU later.

- Scope of administration. You can create a list of administrative scopes in which the same administrative rights are required for a particular group, such as branch offices, computer types (laptops, desktops), server types (Microsoft Internet Information Services (IIS), Microsoft SQL Server®, or domain controllers), and group types (project groups, departmental groups, distribution groups per location, or groups that grant rights and permissions).

Planning Administrative Processes

Most organizations have many administrative entities. One team might be responsible for user and group management, while another team might be responsible for the Active Directory forest and domain infrastructure. Different departments might also manage virtualization, backup, storage, network, and client computers. The application administrator might own the application servers, together with groups to manage rights on the application servers. Certain groups, such as SharePoint site owners, distribution list owners, or project managers, could self-manage access to their resources.

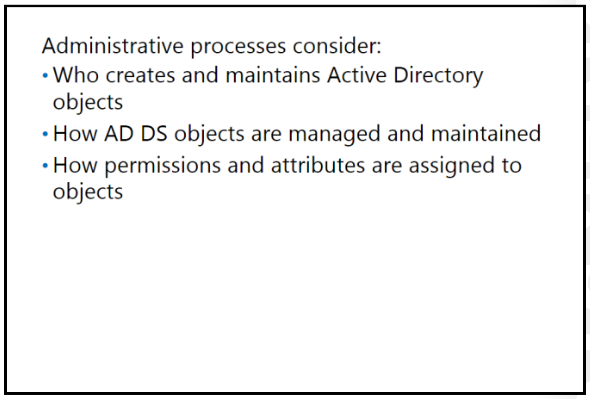

Administrative processes consider:

- Who creates and maintains Active Directory objects
- How AD DS objects are managed and maintained
- How permissions and attributes are assigned to objects

When you design your Active Directory administrative tasks model, you should clearly understand who is, and who should be responsible for administrating the resources. You can establish administrative processes by using the information described in the previous topics and applying it to a practical design for your Active Directory administrative task model, OU structure, and group strategy.

Administrative processes include determining the following aspects of Active Directory domain administration:

- Who will create and maintain Active Directory objects?
- How will you manage and maintain Active Directory objects?
- How will you assign permissions and attributes to these objects?

Best Practices for Administrative Processes

When you gather information about administrative processes, compare the following best practices to your current processes, and keep the best practices in mind as you implement or modify your administrative tasks processes and model:

- Use personalized but separate accounts for administrative purposes. If a user has administrative rights in addition to their user rights, create a second account that is clearly for administrative purposes. Instruct the user to use their administrative account when fulfilling administrative tasks, but to use their regular account for regular day-to-day tasks such as reading email or accessing the Internet.
- Grant administrative permissions only to administrative accounts.
- Avoid assigning permissions to individuals. Instead, group together any accounts that have the same administrative scope and rights, and then assign permissions to this larger group. It is better to create a single-member group and assign permissions to it than to assign permissions to a user object. This simplifies managing permissions when someone leaves the company or moves to a new role, and another person replaces that individual.

- In your OU structure, put administrative groups and administrative accounts in separate places within the structure. You want to avoid a situation in which individuals other than domain administrators are managing permissions in AD DS. By putting administrative groups or accounts in the same part of the OU structure where less-privileged administrative users might obtain delegated rights, you risk having unauthorized personnel elevating their own rights.
- Place objects for regular user accounts, groups, computers, and servers together in OUs based on object type, if they are managed by the same administrators.
- When you assign permissions to objects, always grant the minimum required permissions. Users and groups should have only the permissions that they require on a specific object or OU, and no more. The same principle applies particularly to technical accounts. Administrators often work with application vendors who insist on high-level rights, such as domain administrators. However, many applications can work with the least required privileges. The application of minimized permissions is often called the *principle of least privilege*.
- Always assign permissions at the highest level within the Active Directory OU structure, so that you can benefit from permission inheritance in the structure hierarchy. However, keep in mind the principle of least privilege.

Considerations for Branch Office Delegation

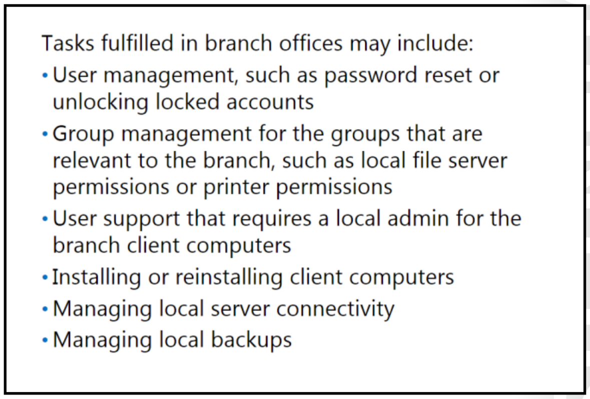

Tasks fulfilled in branch offices may include:

- User management, such as password reset or unlocking locked accounts
- Group management for the groups that are relevant to the branch, such as local file server permissions or printer permissions
- User support that requires a local admin for the branch client computers
- Installing or reinstalling client computers
- Managing local server connectivity
- Managing local backups

To design your AD DS administrative tasks delegation model, you must understand which tasks the branch office administrative staff are performing currently.

You should control certain areas of user management carefully, such as unlocking users who locked their own accounts, or modifying information about the rooms or buildings in which those users work. Examples of control practices include the following:

- Control certain parts of group management, such as groups that assign permissions to local file servers, or printer object permissions.
- Have local administrative permissions on the client computers to support users who have problems with their computers.
- Install or reinstall client computers. This requires permissions to create and delete computer objects. This also requires local administrative rights to the client computers, which you can also manage via Group Policy.
- Manage networking and other connectivity issues with local servers. By grouping the servers into a specific branch office OU, you can delegate administration of the specific branch office servers to the branch office administrators.
- Manage backups if you need to perform decentralized backups. By grouping servers into OUs and delegating rights in a GPO, you can centrally manage branch office administrative rights.

Lesson 2
Designing an OU Structure

You can partition and organize objects in your Active Directory domain by creating a hierarchical structure of OUs. Your OU structure is the key element for establishing a delegation model for your administrative tasks. You can use your design to apply GPOs to certain users or computers. Your design also helps you structure the different types of objects and their organizational roles. If you design your OU structure properly, you can streamline administrative processes and save the administrative staff time when they manage their Active Directory domain.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the advantages and disadvantages of different OU strategies.
- Describe the options for administrative delegation.
- Describe the different types of activities that you can delegate.
- Describe Group Policy inheritance and its impact on the OU design.
- Describe permission inheritance and scenarios that may require changing default inheritance.
- Explain how to protect OUs from accidental deletion.
- Explain how to implement OUs.

Strategies for Designing OUs

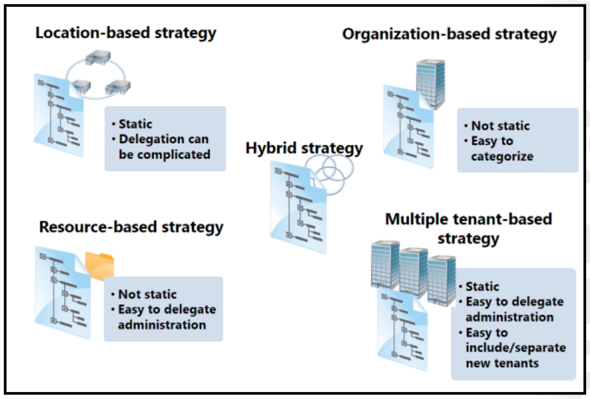

Location-based strategy
• Static
• Delegation can be complicated
Organization-based strategy
• Not static
• Easy to categorize
Hybrid strategy
Resource-based strategy
• Not static
• Easy to delegate administration
Multiple tenant-based strategy
• Static
• Easy to delegate administration
• Easy to include/separate new tenants

As the main components in your Active Directory domain, OUs enable you to and organize all the objects in the domain into a hierarchical structure.

OUs can help you to with the following:

- Delegate rights to administrative groups, enabling them to administer certain objects or object attributes that are located in or below that OU.
- Apply GPOs to users and computer objects that are located in or below that OU.
- Create a hierarchy that enables you to administer the objects in the domain quickly.

You can use several strategies to design your OU structures. The following OU design strategies are the most common:

- Location-based strategy. This strategy uses locations for each top-level OU in the root of the domain. These location-based OUs are the main organizational element of the OU structure. For example, A. Datum Corporation might use a location-based strategy to create OUs for each of its physical locations, which are in London, Toronto, and Sydney, and then create additional OUs for their smaller distribution centers. Each AD DS resource (such as users, groups, and computers), is located in the OU that corresponds to the location where the resource resides.

Organizations commonly use the Location-based strategy when each location operates relatively independently, or when many tasks are delegated to decentralized administrators. For example, the administrative staff in London can perform the main administrative duties, while administrators in Sydney or Toronto who have some delegated rights can access their users, groups, and computers quickly. Therefore, it is more advantageous for the local staff to fulfill some administrative duties across different types of objects in the same branch.

In addition, you do not have to move objects frequently between the top-level OUs, unless the objects move to another physical location. Objects moving to another physical location likely require moving home folders or Microsoft Exchange Server mailboxes. The location-based strategy works well for organizations that anticipate expanding into new locations, because you can add new locations easily to the OU structure.

- Resource-based strategy. This strategy is built around the functions of resources that are in the OU structure. Typically, you separate resources by function (or objects by type), and you create OUs to represent these functions. For example, some common top-level OUs are Servers, Workstations, Groups, and Users. You typically use the resource-based strategy in smaller organizations or in organizations that are maintained centrally by the same administrative staff, and where administrative delegation is based on the object type rather than on the location or department. Examples of these administrative groups include User Helpdesk, Client Support, Virtualization Administration, and Application-Specific Support. In large organizations, those top-level OUs typically are more defined in the next subordinate level. For example, the Servers OU might contain child OUs named after their applications, such as Exchange Server or SQL Server.
- Organization-based strategy. This strategy reflects the structure of an organization's business logic. Top-level OUs represent departments within the organization, such as Sales, Research, or Finance. This strategy works well if resources move frequently, or if they are not affiliated with a physical location, and if employees rarely move between departments. You should consider this strategy when administrative tasks are delegated on a per-department basis rather than per-location. For example, an organization with traveling sales teams and other units that are not location-bound would benefit from an organization-based strategy. However, this strategy is not a good choice for organizations that frequently realign their business model or that encourage employees to shift between roles.
- Multi-tenancy-based strategy. This strategy is suitable for companies that provide the Active Directory infrastructure as a service (IaaS) to other companies, such as a group of affiliated companies that share the same domain, a hosted environment, an outsourced environment, or even a private cloud or public cloud provider. This strategy is appropriate when one company is responsible for maintaining the Active Directory infrastructure, while another company is delegated management of certain Active Directory objects. This strategy is also applicable if a company relies entirely on a hosting company for Active Directory administration.

  For example, A. Datum might maintain AD DS for Trey Research and Contoso, Ltd. Trey Research might want to administer their own Users, Groups, and Computers OUs independently, but might not want to manage Active Directory replication and Domain Name System (DNS). Conversely, Contoso relies fully on the IT staff from A. Datum for all of these tasks. In this scenario, A. Datum would create a top-level OU for ITServices that contains all administrative accounts and groups, and top-level OUs named Adatum, Contoso, and TreyResearch, for each of the managed companies. Under these OUs, regular user, group, workstations, and perhaps even server accounts are represented as deeper levels. Therefore, the Trey Research IT staff would be delegated the responsibility to maintain their own accounts as designed.
  In the multi-tenancy-based strategy, it is possible to allow different tenants to work together or to create privacy settings such that each organization only sees its own resources. In this strategy, it can also be a more straightforward process to include or separate new companies.

- Hybrid strategy. This strategy uses a combination of OUs based on location, organization, and resources. The multiple tenancy-based strategy is also a hybrid strategy. Other hybrid strategies might assign the location at the top level, and separate object types on the next level. Hybrid strategies can be mixed, depending on the organizational requirements. For example, the Servers OU could contain OUs for file servers, IIS, SQL Server, Exchange Server, and other application servers. The Users OU might contain location or departmental OUs, the Workstations OU might distinguish between desktops and laptops, and the Groups OU could incorporate departmental, location, project, or application-specific groups.

Regardless of which strategy you use to design your OU structure, always remember that the main strategic purpose is to enable you to implement an Active Directory administrative tasks model, while a secondary priority might be simplifying the deployment of GPOs. We also recommend separating administrative accounts and groups from regular user accounts and groups that delegated administrators might administer.

Question: What is the OU structure that you use at your workplace, and why is it designed that way?

Question: What current issues are you facing with your OU model?

How Administrative Permissions Work

To implement the Active Directory administrative tasks model, you merge the OU design with the OU permissions. This enables delegated administrators to fulfill administrative tasks. To create the administrative task model and design the OU structure to support this, you must understand how Active Directory administrative delegation works, and the options that exist for delegating administrative control.

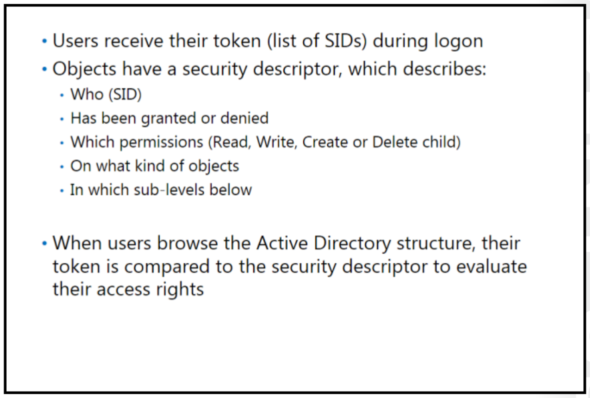

- Users receive their token (list of SIDs) during logon
- Objects have a security descriptor, which describes:
  - Who (SID)
  - Has been granted or denied
  - Which permissions (Read, Write, Create or Delete child)
  - On what kind of objects
  - In which sub-levels below
- When users browse the Active Directory structure, their token is compared to the security descriptor to evaluate their access rights

How Do Users Get Permissions?

When users sign in to an Active Directory domain, they receive a token, which is a list of security identifiers (SIDs) for their individual account, historical accounts (if they have been migrated), and every group they belong to (even recursively). If any group was migrated from another domain, it is likely that the user also received the historical SIDs of those groups in the former domain.

In Windows® operating systems, many objects (such as files, folders, registry keys, processes, and Active Directory objects), contain a security descriptor. Based on the SID, the security descriptor defines which rights are granted or denied, and to whom.

When a user browses files and folders or registry keys, or navigates through the Active Directory domain structure, the list of SIDs in the token is compared with the list of SIDs that are in the security descriptor. If there are any matching SIDs, the system validates the type of access, and allows or prohibits the current operation.

Active Directory OU Permissions

In Active Directory, the permission model is more complex than in most other Windows operating system services. Security settings on the Active Directory domain are inherited hierarchically in the OU structure of that domain. At any point in the structure, you can configure additional security settings that could be inherited throughout the hierarchy. Inheritance depends on the scope of inheritance that is defined in the security setting, and whether inheritance is blocked at a lower level. New objects obtain default security settings (which are defined in the schema class), and inherit security settings from their parents.

For example, in the OU schema definition, Account Operators are granted full rights to create and delete objects for computer accounts, user accounts, group objects, and **inetOrgPerson** objects. Therefore, if you remove the default Account Operators group from the security permissions of an OU, and then you create a child OU, the child OU retains the explicit security settings of the Account Operators.

Active Directory Object Security Descriptors

A security descriptor of an Active Directory object contains the following parts:

- The owner of the object. The owner can reset security settings even if he or she accidentally configures them to have no permissions on the object.
- The primary group of the owner (SID). The default primary group is Domain Users.
- A control field that specifies whether the discretionary access control list (DACL) or system access control list (SACL) are present. This field also contains information regarding whether there is inheritance blocking.
- An optional DACL. The optional DACL contains permissions for granting or denying access.
- An optional SACL. The optional SACL contains the auditing permissions when you enable Success or Failure auditing.

The DACL and SACL are containers that contain one or more access control entries (ACEs). An ACE stores the following information:

- Who (which security principal, such as a user or group) is allowed or denied access
- To Read, Write, Create, or Delete
- Which objects, or object attributes
- At what sublevels (on the OU level only, on objects only in the OU, or objects in any sub-OU)

To verify or adjust the security settings, you use the Security tab in the OU Properties dialog box, and in particular, the Advanced Security dialog box in the properties of an OU. The Security Delegation Wizard assists with some common tasks, but you cannot use it for reviewing security settings.

Designing OUs for Delegating Administrative Control

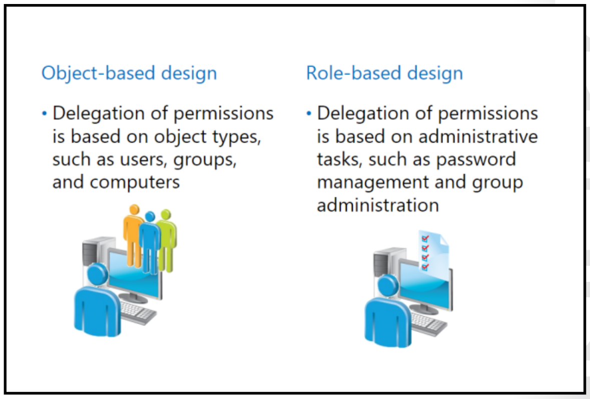

Object-based design

- Delegation of permissions is based on object types, such as users, groups, and computers

Role-based design

- Delegation of permissions is based on administrative tasks, such as password management and group administration

Designing your Active Directory OU and administrative tasks model is a key element to enabling secure and effective administration of your domain. Administrator delegation must be done for other IT administrative groups, and for applications, such as mobile messaging applications, which also require rights in AD DS.

Delegation Control Methods

When you delegate control of partial structures in your Active Directory OU, you must consider two factors: to whom you are granting permissions, and to where. In AD DS, you can grant specific rights on resources. You can allow the creation or deletion of only certain object types, or you can select the individuals who have rights on a particular attribute of a specific object type, such as group account descriptions, or their members.

Except in rare cases (such as service accounts), you should always grant administrative control to groups rather than users. Even if the group contains only one user, this individual might leave the company, and determining where that individual had permissions to would be more difficult than modifying the appropriate group memberships.

Two methods for delegating administrative control over Active Directory domain resources are as follows:

- Object-type delegation: In this delegation model, you can delegate various levels of control to groups based on the objects that the groups control. An example of an object-type delegation would be if you delegated control to the Toronto Admins group for objects within the Toronto OU. In this case, the Toronto Admins become responsible for the majority of administrative tasks within the Toronto OU.

  You typically use object-type delegation if you have only a few administrators, or if minor delegation is required. This type of delegation also works well if many administrators require the same level of control, typically over most of the domain structure.

  We do not recommend object-type delegation in an environment where different users require various levels of control over different objects. This is because determining what level of control is granted to which users for a specific object can be difficult.

- Role-based delegation. This delegation model involves creating several specific groups to which you delegate administrative control. These groups usually relate to a specific resource (or resources), and you can name groups for the level of control that you assign to them. Unlike object-based delegation, role-based delegation involves granting permissions to modify only some of the attributes of an object. For example, you could create the role-based group Change Finance User Password, and then assign permissions to that group to change passwords for any users in the Finance OU.

  To ensure that your role-based delegation is effective, all functions or roles within the Active Directory domain structure should have an associated group. This level of precision can help you to determine which level of control you have assigned to an individual user, because you simply examine the role-based groups to which the user belongs.

  Role-based delegation can take longer to implement than object-type delegation. However, if you design the OU and group structure properly, role-based delegation saves administrative effort, especially in larger organizations.

Note: You also can combine aspects of each method to create a hybrid method.

Designing Administrative Control Delegation

You should consider the following points when you design your OU structure to delegate administrative control:

- Delegate control as high in the OU structure as possible to make use of inheritance. Use a hierarchy of OUs when appropriate.

- Delegate only the administrative privileges that are necessary. This might require you to create more OUs to separate objects further. For example, if all workstation computers in your organization are in an OU called Workstations, and if you want to delegate control over joining computers to the domain only for Workstations in a single branch, you need to redesign your OU structure to separate Workstations by branch (likely underneath the workstations OU).

If you require many administrators to have different levels of control in the domain, consider using role-based delegation and a larger number of specific OUs.

Designing OUs for Applying GPOs

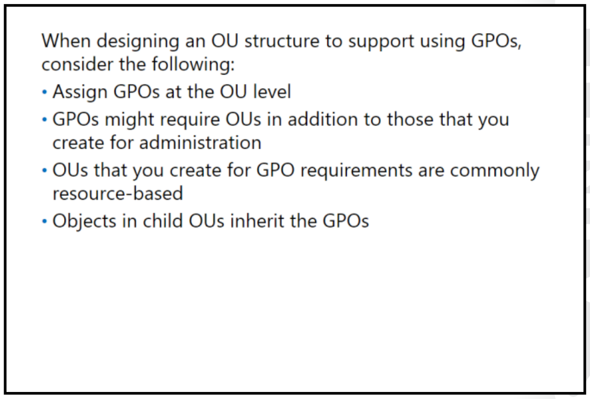

When designing an OU structure to support using GPOs, consider the following:

- Assign GPOs at the OU level
- GPOs might require OUs in addition to those that you create for administration
- OUs that you create for GPO requirements are commonly resource-based
- Objects in child OUs inherit the GPOs

Another important consideration in designing your OU structure is the application of GPOs. In the Active Directory domain structure, you can link GPOs to the following objects:

- The Active Directory domain object, where the Default Domain Policy is linked by default. Policies linked at this level apply to all objects in the domain.
- The Active Directory site object. Policies linked to this object apply to all objects in the site.
- The Active Directory OU. Policies can be linked to any level of an OU, and by default are inherited by OUs at lower levels.

In addition, you can configure security settings (called *security filtering*), on the GPO to allow only certain users or computers to apply that GPO. You can also filter by using Windows Management Instrumentation (WMI) filters, for example to apply a GPO only to computers of a certain hardware vendor. You must design your OU structure to account for the use and application of GPOs.

In most cases, you design the OU structure at two or more levels. You use the first and highest-level OUs to delegate administration and to assign permissions to Active Directory domain objects. You typically create these OUs according to location, department, or resource.

If you use the location-based or department-based models, you commonly use a second level of OUs to separate the resources into logical units that align with the IT structure. You typically apply GPOs to these OUs. For example, Contoso, Ltd might have OUs for each of its locations: New York, Toronto, and Tokyo. Within these OUs, you create a second level that separates resources by type. You then name these OUs for resources, such as Servers, Workstations, and Printers. For many organizations, these two levels provide the necessary framework to delegate control and to apply GPOs. In scenarios that are more complex, you might need to implement further levels to enable a granular delegation model or GPO implementations.

If you require even more granularity, you can break down resource-based OUs even further. For example, the Servers OU might have child OUs named Exchange Servers, SQL Servers, Domain Controllers, and File Servers. You then can apply GPOs that are specific to these different server types to the categorized OUs.

GPO applications, by default, use inheritance in the OU structure. This means that if you apply a GPO to the Servers OU, AD DS by default applies the GPO to the Exchange Servers, SQL Servers, and File Servers OUs.

Considerations for Designing OU Hierarchies

When planning Active Directory functionality, you should consider the following high-level aspects of OU design:

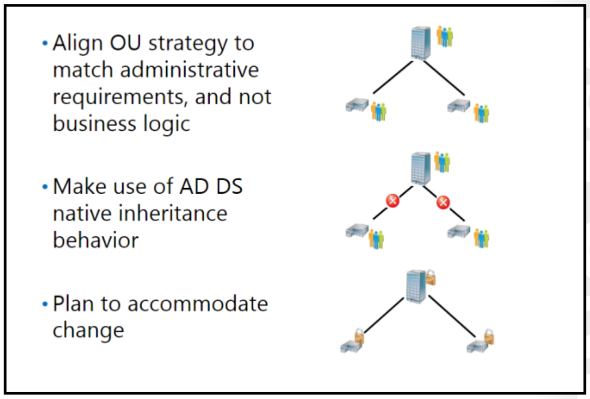

- Align OU strategy to match administrative requirements, and not business logic
- Make use of AD DS native inheritance behavior
- Plan to accommodate change

- The OU structure should align primarily to administrative purposes, such as to your Active Directory tasks delegation model and your GPOs. Avoid mimicking your organizational chart, unless it benefits the administrative model. This is because organization charts change more frequently than you will want to change your OU structure. If you have requests for department-specific GPOs, it might be preferable to reflect those departments in a lower level of the OU structure

  Users and Managers usually do not see the OU structure, so there is no benefit in having corporate hierarchies listed in the structure. If you want users to enable browsing the organizational structure, ensure that you fill in the manager attribute of the user objects. This enables users to use current and previous versions of Microsoft Office Outlook® in combination with an Exchange Server messaging infrastructure, Microsoft Office 365™, or Microsoft SharePoint Server to navigate through the organizational structure. The OU should be for administrative purposes only.

- Inheritance is an important part of OU functionality, both for delegation of control and GPO application. You should design the OU structure to include objects that require the same administrative control or Group Policy settings within the same OU structure. This way, you can assign the delegation of control or the GPO setting only once at a higher level in the OU structure, rather than individually at each child level. Remember that you can block inheritance for certain child OUs if you do not want AD DS to apply the settings from a higher OU level to certain child OUs. When designing the OU structure, remember to consider inheritance blocking, and how you might design a structure that rarely requires it.

- Design your OU structure to accommodate change. After you implement an OU structure, it can be difficult to change. This is especially true if you also change design strategies, such as changing from an organization-based to a resource-based OU structure. Ensure that your OU structure leaves room for organizational growth and a reasonable level of structural change.

Protecting OUs from Accidental Deletion

Accidental deletion is one of the major causes of Active Directory recoveries. Therefore, Windows Server® 2008 introduced the Protect OUs from Accidental Deletion feature. When designing your OU structure, or when migrating your AD DS domain from earlier versions of Windows Server to Windows Server 2012, we recommend that you use this feature to protect all OUs from accidental deletion.

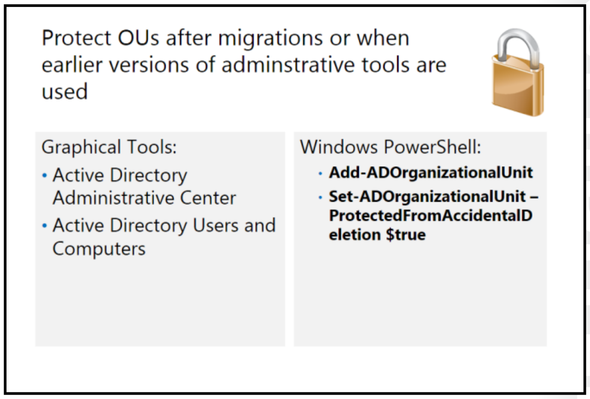

Protect OUs after migrations or when earlier versions of adminstrative tools are used

Graphical Tools:

- Active Directory Administrative Center
- Active Directory Users and Computers

Windows PowerShell:

- **Add-ADOrganizationalUnit**
- **Set-ADOrganizationalUnit –ProtectedFromAccidentalDeletion $true**

OUs that you protect from accidental deletion share the following benefits:

- OUs cannot be deleted accidentally. If an administrator wants to delete an OU purposely, he or she must remove the protection prior to deleting the OU.
- OUs cannot be moved accidentally.
- OUs that are newly created using Active Directory Administrative Center, or Active Directory Users and Computer in Windows Server 2008 or newer, are protected automatically against accidental deletion.

However, you should also consider the following:

- OUs that were created prior to migrating the domain are unlikely to be protected. Therefore, as a best practice, you should protect them after migration. Running the Windows Server 2012 Best Practice Analyzer can also tell you if you have unprotected OUs.
- OUs that were created using a custom script or older versions of administrative tools are also unlikely to be protected from accidental deletion.
- Protection from accidental deletion does not work if you delete a higher-level, unprotected OU, even if OUs further down in the hierarchy are protected. Therefore, we recommend protecting all OUs from accidental deletion.

You can configure OU protection in the Active Directory Administrative Center on the properties of the OU. You can also use Active Directory Users and Computers, after enabling View\Advanced View, on the Object Tab in the OU Properties dialog box.

Protecting OUs from Accidental Deletion Using Windows PowerShell

You can also use the Windows PowerShell® command-line interface to search for OUs that are not protected from accidental deletion. Use the following Windows PowerShell command to identify the unprotected OUs:

```
Get-ADOrganizationalUnit –filter * -properties ProtectedFromAccidentalDeletion | where
{$_.ProtectedFromAccidentalDeletion –eq $false}
```

You can also use commands in Windows PowerShell to identify all OUs that are not protected from accidental deletion, and then protect them.

```
Get-ADOrganizationalUnit –filter * -properties ProtectedFromAccidentalDeletion | where
{$_.ProtectedFromAccidentalDeletion –eq $false} | Set-ADOrganizationalUnit –
protectedFromAccidentalDeletion $true
```

Question: Are there things about the OU structure in your organization that you would like to change?

Demonstration: Implementing OUs

In this demonstration you will see how to:

- Create an OU.
- Verify that the OU is protected against accidental deletion.
- Examine the default security settings of the OU.
- Delete a protected OU.

Demonstration Steps

Create an OU

- On LON-DC1, open Active Directory Administrative Center and create a new root-level OU with the following parameters:
  - Name: **Contoso-IT**
  - Description: **OU to contain Accounts / Groups for administrative purposes**

Verify that the OU is protected against accidental deletion

1. From Server Manager, open Active Directory Users and Computers.
2. Enable the **Advanced Features** view.
3. View the **Contoso-IT** OU properties.
4. On the **Objects** tab, verify that the OU is protected from accidental deletion.

Examine the default security settings of the OU

1. Switch to Active Directory Administrative Center.
2. Open the **Contoso-IT** OU properties.
3. In **Advanced Security Settings**, review the default security settings.

Delete a protected OU

1. In Active Directory Administrative Center, open the **Contoso-IT Properties** OU dialog box.
2. Clear the **Protected from accidental deletion** check box.
3. Delete the **Contoso-IT** OU.

Lesson 3
Designing and Implementing an AD DS Group Strategy

Active Directory groups are the fundamental objects that you use to control access to Active Directory resources and other Windows security-enabled applications in your Active Directory domain and forest. As you design the administration of an Active Directory domain, you must consider group strategy. To implement an effective Active Directory group strategy, you should understand the available group types and scopes, and how they interact. In addition, you should observe naming and location strategies so that you can make your group simple to administer and use. Finally, you should use the various best practices, such as group nesting.

Lesson Objectives

After completing this lesson, you will be able to:

- Summarize the different types of Active Directory groups in Windows Server 2012.
- Explain how to develop an Active Directory group naming strategy.
- Describe strategies for using Active Directory groups to access resources.
- Describe considerations for planning Active Directory group administration.
- Describe guidelines for designing an Active Directory group strategy.
- Explain how to create and manage groups.

Active Directory Groups in Windows Server 2012

An *Active Directory security group* is an Active Directory domain object. It can contain users and computers, and it enables administrators to grant permissions at the group level instead of on each individual object. You can characterize groups by type and scope.

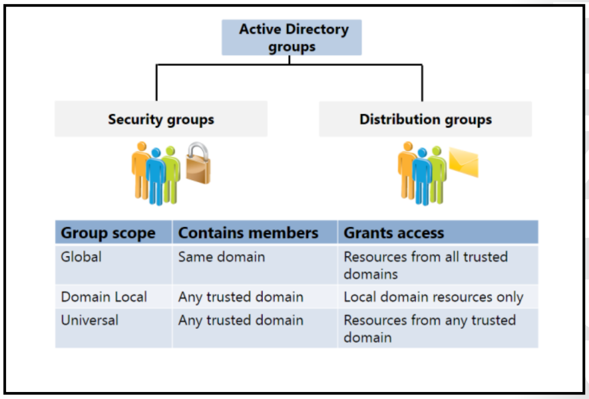

Group scope	Contains members	Grants access
Global	Same domain	Resources from all trusted domains
Domain Local	Any trusted domain	Local domain resources only
Universal	Any trusted domain	Resources from any trusted domain

Group Types

Windows Server 2012 features two types of groups:

- Security groups. You can use security groups to assign user rights and permissions to a group of users and computers.
- Distribution group. You can use distribution groups with certain email applications such as Exchange Server, so that you can send email messages to a group of users.

Note: You can convert groups between security groups and distribution groups. Although you can email-enable a security group, you cannot grant security permissions to a distribution group. If you grant permissions to a security group and then convert it to a distribution group, the permissions remain but do not take effect.

Group Scope

The group scope defines whether the group can have members of other domains, or if the group can be granted access to resources in other domains. The group scope also defines how you assign permissions to group members.

Groups can have the following scopes:

- Global groups. You can use groups with a global scope to grant rights in trusted domains. However, Global groups can only contain members of the same domain.
- Domain Local groups. You can use groups with a domain local scope only to grant rights in the local domain. However, Domain Local groups can contain members of other, trusted domains.
- Universal groups. Universal groups can be granted permissions in any trusted domain, and can contain members of any trusted domain. Universal groups are designed to consolidate groups from various domains that require access to the same resources.

  Question: What issues are you currently facing in your organization regarding to your group strategy?

Developing an Active Directory Group Naming Strategy

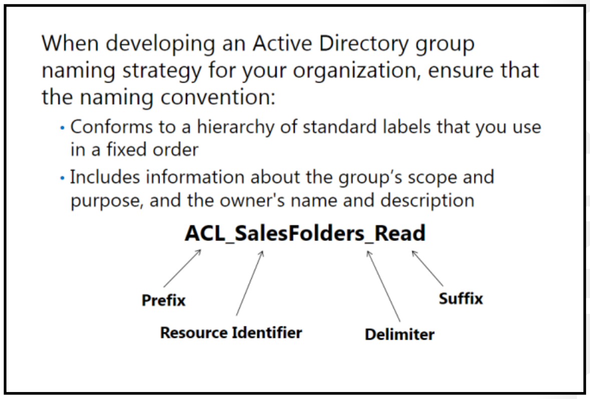

When developing an Active Directory group naming strategy for your organization, ensure that the naming convention:

- Conforms to a hierarchy of standard labels that you use in a fixed order
- Includes information about the group's scope and purpose, and the owner's name and description

ACL_SalesFolders_Read

Prefix

Resource Identifier

Delimiter

Suffix

When developing your Active Directory group naming strategy, use a universal naming convention to ensure that people within your organization can identify groups easily. Otherwise, your Active Directory structure can become unorganized and misunderstood within the organization.

The names that you select should help you manage the groups and the enterprise. A best practice is to use a naming convention that identifies the type of group and its purpose, such as the following examples:

- Role groups. These identify people who are grouped according to business logic. Use a simple, unique name, such as Sales, Marketing, or Finance.
- Resource groups. You typically use resource groups to assign specific permissions to specific resources within the domain. You should identify the functionality of these groups by name.

  For example, ACL_SalesFolders_Read has the following elements:

  - *Prefix*. The prefix identifies the management purpose of the group. In the example, ACL identifies that the group appears in the Access Control List of shared resources.
  - *Resource identifier*. This uniquely identifies the resource that the group is managing and to which it has access. In the example, the resource identifier is SalesFolders.
  - *Suffix*. The suffix further defines the access rights that the group membership is granted. For resource access management groups, the suffix defines the level of access that group members have. In this example, the suffix is Read, indicating that the group has Read permissions.
  - *Delimiter*. This should be a consistent marker, such as an underscore (_), which separates the prefix, identifier, and suffix. Do not use the delimiter elsewhere in the name. Use it only as a delimiter between those parameters in your naming strategy. Note that you do not use the delimiter between the words Sales and Folders. Group names can include spaces, but you must then enclose those group names in quotes if you refer to them in commands or scripts.

You can facilitate auditing and reporting by creating scripts. You can also adjust OU structures or other business requirements by creating scripts that use the delimiter to deconstruct group names.

Remember that nontechnical users often use role groups that define user roles. For example, you might email-enable the Sales group so that people in your organization can use the group as an email distribution list. Therefore, keep your naming convention for role groups simple and straightforward. In other words, your naming convention for role groups should not overly use prefixes, suffixes, or delimiters, and instead it should emphasize user-friendly, descriptive role group names. In larger organizations it might be more appropriate to use simple prefixes or suffixes, such as in the following examples:

- US_Sales or UK_Marketing, to distinguish between countries or locations.
- Sales_Hardware, Sales_Software, and Sales_All, to distinguish between sales roles.

Strategies for Using Groups to Access Resources

Groups are the principal way to control access to resources within your Active Directory domain and forest. When implementing groups within AD DS, you should take into account several aspects of group behavior.

Group Nesting

In almost all cases, you should use groups to control access to resources instead of giving permissions to individual user objects. Placing groups within groups (also called *group nesting*), is an important part of designing and using groups to control access to resources. If you nest groups, you can manage multiple objects and groups simultaneously, and you can provide a more modular and flexible group structure.

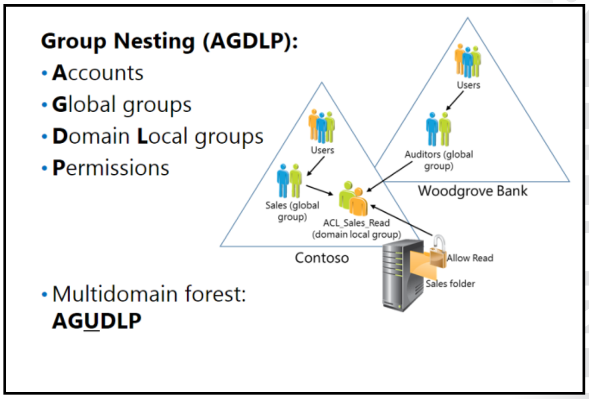

Group Nesting (AGDLP):

- **A**ccounts
- **G**lobal groups
- **D**omain **L**ocal groups
- **P**ermissions

- Multidomain forest: **AGUDLP**

There are different methods of nesting, depending on whether the group scope is Global, Domain Local, or Universal. The following table lists what each group can include as members.

Group scope	Members
Global	• Accounts from the same domain as the global group • Global groups from the same domain as the global group
Domain Local	• Accounts from any domain • Global groups from any domain • Universal groups from any domain • Domain Local groups, but only from the same domain
Universal	• Accounts from any domain within the forest • Global groups from any domain within the forest • Universal groups from any domain within the forest

Group Nesting Best Practices

Group nesting best practices are commonly defined by the acronym AGDLP. AGDLP defines the order in which group nesting should occur:

- **A**ccounts (user and computer accounts) are members of:
- **G**lobal groups, which represent business roles, for example Sales. Global groups are members of:
- **D**omain **L**ocal groups, which represent management rules, such as determining who has Read permissions to a specific collection of folders. Domain Local groups are granted:
- **P**ermission to access resources. In the case of a shared folder, access is granted by adding the Domain Local group to the access control list (ACL) of the folder, with a permission that provides the appropriate level of access.

AGDLP Example

This best practice for implementing group nesting translates well even in multiple domain scenarios, such as in the following illustration.

Consider the following figure, which illustrates an AGDLP scenario:

This figure represents a group implementation that reflects the technical view of group management best practices (AGDLP), and the business view of role-based, rule-based management.

Consider the following scenario:

The sales force at Contoso, Ltd has just completed its fiscal year. Sales documents from the previous year are in a folder named Sales. The sales force requires Read access to the Sales folder. A team of auditors from Woodgrove Bank, a potential investor, requires Read access to the Sales folder to perform an audit.

To implement security for this scenario, you must complete the following steps:

1. Assign users who have common job responsibilities or other business characteristics to role groups, which you then implement as global security groups. Do this separately in each domain. Add Contoso sales people to the Sales role group, and Woodgrove Bank auditors to the Auditors role group.
2. Create a group with Read permission to manage access to the Sales folders. You should implement this in the domain that contains the resource that you are managing, for example, the file server on which the sales folder resides—in this case, the Sales folder resides in the Contoso domain. In addition, create the rule group for resource access management as a domain local group, ACL_SalesFolders_Read.
3. Add the role groups to the rule group for resource access management so that you can represent the management rule. These role groups can come from any domain in the forest, or from a trusted domain, such as Woodgrove Bank. Global groups from trusted external domains, or from any domain in the same forest, can be members of a domain local group.
4. Assign the permission that implements the required level of access. In this case, grant the Allow Read permission to the domain local group.

Universal Groups and AGUDLP

A multiple domain forest also has universal groups, which fit between global and domain local groups. AD DS replicates universal groups throughout the forest, and these groups can contain members from anywhere within the forest. You often use universal groups to consolidate global groups from multiple domains. In the group-nesting scenario, you can place that universal group into domain local groups in multiple domains. You can remember the nesting of universal groups by adding the letter U to the AGDLP acronym, which now becomes AG**U**DLP, where the U stands for universal.

Considerations for Planning Group Administration

Options for group placement in AD DS include the following:

- Place group objects in the same OU that contains the group accounts
- Place group objects in the same OU where a resource exists
- Place all groups centrally in the same location in AD DS
- Place groups in separate OUs
- Allow group self-management
- Hybrid scenarios

Although the main function of groups in AD DS is to group users to simplify assigning permissions for accessing resources, you also must consider how you administer the group objects themselves. This includes how you create and delete groups, and how you manage group membership and permissions. In AD DS, designing a group management strategy is primarily about structuring group objects. When planning group administration, you can use the following strategies:

- Place group objects in the same OU containing the user objects. With this strategy, you can delegate both user and group management under one OU. You could use this strategy when one administrator per department administers departmental user accounts and their groups.
- Place group objects in the same OU as where the resource exists. With this strategy, you can manage resources and the access to those resources under one delegation group. However, this method is not very common or practical. For example, it is usually unnecessary to administer computer objects for servers, other than reinstallation or changing an administrative description in AD DS. If you manage groups for an application on the server, it is more likely that the same administrative group should have administrative rights on the server locally, which is different from having rights on the computer object of the server.
- Place all group objects centrally in the same location in AD DS. This strategy works well for small Active Directory environments. It places all of your groups in one container, while keeping them separate from other AD DS resources. This strategy can be useful if you only want to delegate control of groups to one or more users.
- Place group objects in separate OUs. This is useful if you have a large number of groups, or if you delegate the management of several groups to different administration teams.
- Self-management of groups. This approach allows group members to self-manage, which means that they can add other members to their group, or remove themselves. This is useful for community distribution groups or community projects within an organization that do not rely on the organizational structure or have security implications. In Windows Server 2012, the Managed By tab in Active Directory Users and Computers, and the Managed By section in the group properties in Active Directory Administrative Center both include a check box that the group manager can select to update the group members.
- Hybrid scenarios. This scenario is most common in large organizations. Depending on the purpose of the group, you can choose from several different hybrid scenario strategies. Groups that reflect the organizational structure are administered by a central administrative entity. Groups that grant access to applications or SharePoint sites are administered by the application owner or the owner of the SharePoint site. Groups that collaborate in communities are self-managed.

When deciding on group placement, you should consider who manages and administers the groups. You can assign a group manager for each group, or you can leave this task to a domain administrator or a designated delegated group that provides all group management. The first approach provides more decentralized administration, while the second approach allows only specific privileged users to manage groups.

Guidelines for Designing an Active Directory Group Strategy

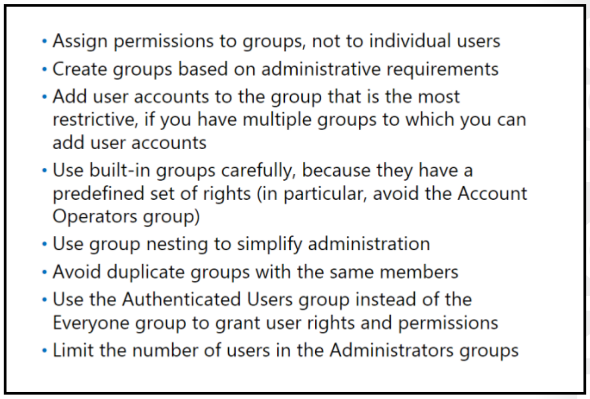

- Assign permissions to groups, not to individual users
- Create groups based on administrative requirements
- Add user accounts to the group that is the most restrictive, if you have multiple groups to which you can add user accounts
- Use built-in groups carefully, because they have a predefined set of rights (in particular, avoid the Account Operators group)
- Use group nesting to simplify administration
- Avoid duplicate groups with the same members
- Use the Authenticated Users group instead of the Everyone group to grant user rights and permissions
- Limit the number of users in the Administrators groups

Consider the following guidelines when designing an Active Directory group strategy:

- Grant permissions to groups only, not to individual users, even if there is only one user who requires access. If you assign permissions directly to a user account and the account is deleted, it leaves orphaned entries in the ACLs for the objects to which the permissions apply.
- Create groups based on administrative requirements. If you create a group based on a job function and then another person takes over that job, you must change only the group membership; you do not have to change all permissions that are granted to the individual user account.
- If you have multiple groups to which you can add user accounts, add user accounts to the group that has the least permissions.
- Use caution with default groups. In most cases, these groups have predefined rights, which might permit greater access than is required. This applies particularly to the Account Operators group, which has rights on every OU to create and delete objects, and full control over user, group, computer, and inetOrgPerson objects. When you implement an OU structure, you could have separate OUs for each of these object types. You would not want to enable anyone to create computer accounts where user accounts are supposed to reside, and vice versa.
- Use group nesting wherever possible to simplify administration. Nesting groups is the process of adding groups to other groups. By nesting groups, you simplify some administrative tasks such as granting rights to resources.
- Avoid duplicating groups. For example, if you have departmental groups and you want one group to control access to the department's folder share and another group to control access to a distribution list, consider using the same role group and then email-enable it. If you have two groups that contain the same members, it is likely that over time those groups will be administered inconsistently, and will contain different members. Additionally, you increase the token size, which can cause issues as the infrastructure grows.
- Use the Authenticated Users group instead of the Everyone group to grant user rights and permissions to all users. Authenticated Users limits access to users who are signed in to the Active Directory domain, but the Everyone group also includes anonymous users.
- Limit the number of users in the Administrators, Domain Admins, and Enterprise Admins groups. The permissions for these groups are far-reaching, and one user seldom requires the entire set of permissions. For a more secure environment, you can create separate administrative groups that have the necessary permissions only.

Demonstration: Creating and Managing Groups

Adam Barr is a marketing employee who resides in London and works for A. Datum Corporation. In his spare time, Adam enjoys different types of sports. As such, he has decided to start a Sports In London group at A. Datum to gather interested colleagues together. Due to the low security relevance, administrators at A. Datum have agreed that the group members should manage their own community distribution list.

In this demonstration, you will see how to:

- Create an OU.
- Create a group, and then configure management of the group.
- Add a user to the group.
- Verify that the community group can manage itself.

Demonstration Steps

Create an OU

- On LON-DC1, open Active Directory Administrative Center and create a new root-level OU using the following settings:
    - Name: **SelfService**
    - Description: **OU for Groups that are managed by themselves**

Create a group, and then configure management of the group

1. In Active Directory Administrative Center, create a group in the SelfService OU, using the following settings:

- Name: **SportsInLondon**
- E-mail: **SportsInLondon@adatum.com**
- Description: **SelfManaged DL to contain the members of the Sports in London community**

2. Configure the Group to manage itself.

Add a user to the group

1. In Active Directory Administrative Center, in the Marketing OU, in the details pane, click **Adam Barr**.
2. Make Adam a member of the **SportsInLondon** Group.

Verify that the community group can manage itself

1. Sign off LON-DC1.
2. Sign in to LON-CL1 as **Adatum\Adam** with the password **Pa$$w0rd**.
3. In Active Directory Administrative Center, add **Pat Coleman** to the group **SportsInLondon**.
4. Sign off LON-CL1.

Lab: Designing and Implementing an Active Directory OU Infrastructure and Delegation Model

Scenario

In the past, A. Datum Corporation has used a highly centralized approach to managing its IT infrastructure. However, because the company has expanded to other countries, this centralized approach is no longer efficient. As a result, IT management wants the Active Directory design team to recommend how to change the Active Directory administration structure to meet new requirements.

Objectives

After completing this lab, you will be able to:

- Design an OU infrastructure.
- Implement the OU design.
- Design and implement an AD DS permissions model.

Lab Setup

Estimated Time: 120 minutes

Virtual machine	20413C-LON-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 **20413C-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in using the following credentials:

- User name: **Administrator**
- Password: **Pa$$w0rd**
- Domain: **Adatum**

Exercise 1: Designing an Organizational Unit Infrastructure

Scenario

A. Datum Corporation has expanded its business through acquisitions. Their IT management is considering providing AD DS as a service to ensure that critical Active Directory replication and name resolution works across companies. However, IT management also is planning to allow local administration to maintain administration of the Active Directory objects.

In this exercise, you need to design an OU structure that meets these new business and organizational requirements.

Supporting Documentation

Brad Sutton

From: Bill Malone [Bill@adatum.com]
Sent: 25 September 09:05
To: Brad@adatum.com
Subject: Active Directory Redesign
Attachments: Corporate Domains and Sites.vsd

Hello Brad,

AD DS is a very important service to us. We want to avoid outages anywhere in our corporate environment from preventable errors, especially now that we have acquired Trey Research and Contoso, Ltd. We must ensure that we can work together as efficiently and effectively as possible.

While you were in Cambridge last week, our management team discussed whether it would be possible to provide Active Directory Domain Services as a central service out of A. Datum, and in the midterm perhaps even merging Trey Research and Contoso into our domain. We can discuss these mid-term requirements later.

For now, I would like you and Charlotte to propose an OU and administrative structure that meets the requirements that I have outlined in the attached proposal template.

For the first steps, I'd like you to update the proposal documentation. Then, I would like you to create the OU structure, but maintain the old one until we have tested the scenario and obtained management approval for the migration.

Regards,

Bill

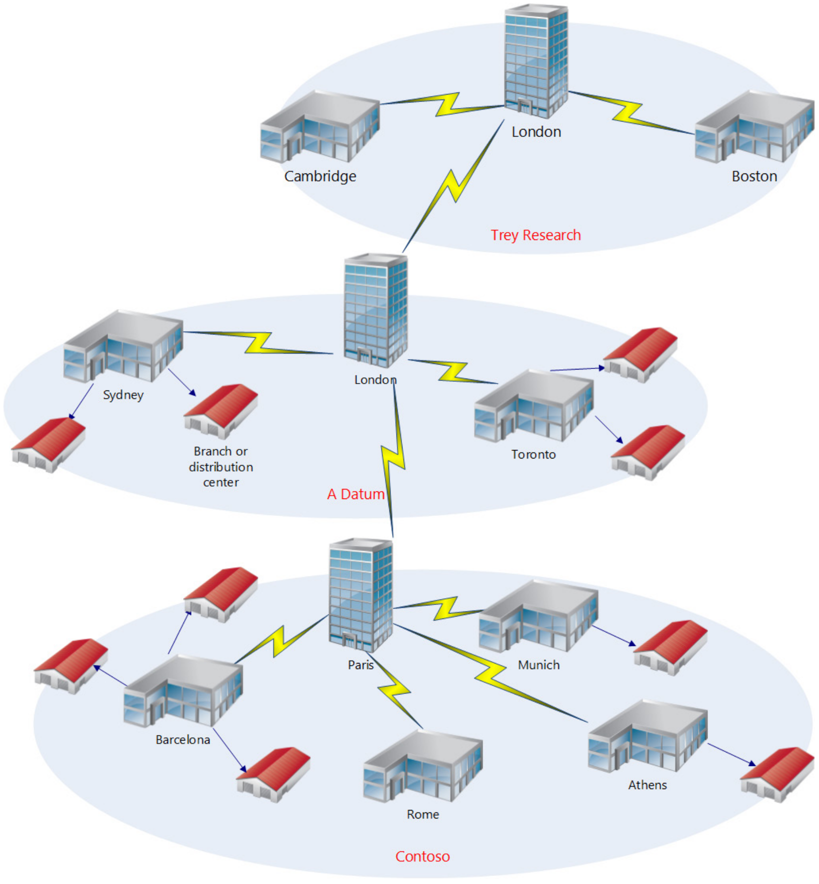

A. Datum OU Redesign Proposal

Document Reference Number: BS00925/1

Document Author	Brad Sutton
Date	25th Sep

Requirements Overview

Design an OU structure that meets the following requirements:

- For administrative purposes, only use personalized administrative accounts.
- The Administrators group in A. Datum must be the only group of users who can make domain-wide changes such as configuring Active Directory sites, creating top-level OUs, managing domain controllers, and creating and managing administrative user accounts and groups.
- Facilitate future integration of Contoso and Trey Research into the A. Datum OU structure (excluding the research.treyresearch.net domain). The central A. Datum team should run the Active Directory replication and infrastructure services. However, it is up to Trey Research and Contoso to determine what levels of day-to-day administration of users, groups, and computers they would like to retain.
- A single group of administrators in the London office must be the only people who can create and delete all regular user and group accounts in the Adatum.com domain.
- Each regional hub office must have a local administrators group, which is responsible for troubleshooting local computer issues and supporting their users in their respective regions. The local administrators must have full access to all servers in their offices (other than domain controllers), and they must be able to support user issues on those client computers.
- Every department has a local Administrators group that can reset passwords for their users.
- The A. Datum team that is responsible for managing enterprise-wide application servers must manage access to the applications and resources that these servers host. Plan for the following three roles for these servers: SQL, WEB, and APP.
- Users from the Contoso.com domain and users from all other offices must be able to access files on a file server that is in Paris.
- The processes for adding multiple user accounts simultaneously, for changing the attributes for a large number of user accounts and computers, and for managing groups must be automated.

Summary of Information

- A. Datum:
  - Responsible for the Active Directory infrastructure.
  - Manages administrative accounts and associated groups.
  - Contains three major locations: London, Toronto, and Sydney. (Smaller locations will not be included in the draft design.)
  - Provides some enterprise-wide servers.
  - Has a dedicated team to create or delete all users.
  - Has departmental administrators who are responsible for user support, such as password resets.
  - Has regional hub administrators who are responsible for client computer and local server management.

A. Datum OU Redesign Proposal

- Trey Research:
  - Will merge into the Adatum.com domain at a later point, and the OU structure should allow for this.
  - Will not integrate the research.treyresearch.net domain.
- Contoso:
  - Will merge into the Adatum.com domain at a later point, and the OU structure should allow for this.
  - Has a file server that provides services to all users in the enterprise.

Proposals

1. Which OU design model do you think would work in this situation?

2. How does centralized administration in London affect the overall OU design?

3. How can you place Active Directory objects within the OU structure?

4. Which OUs do you suggest for the top-level OUs?

5. Which OUs should you create so that you can delegate administration?

6. Which OUs should you create so that you can manage the local servers?

7. Which OUs should you create for managing the enterprise application servers?

8. Create a drawing of your draft design.

The main tasks for this exercise are as follows:

1. Read the supporting documentation

2. Update the proposal document with your planned course of action

3. Examine the suggested proposals in the Lab Answer Key

4. Discuss your proposed solution with the class, as guided by your instructor

▶ Task 1: Read the supporting documentation

- Read the documentation provided.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposal section of the A. Datum OU Redesign Proposal document.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones in the Lab Answer Key.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you should have created an OU design that reflects the A. Datum Corporation's administrative task model.

Exercise 2: Implementing the OU Design

Scenario

You will now implement the OU design while keeping the old OU structure in place as a backup. This is a common scenario in many companies that are changing their OU structure. To implement the OU design, you will complete the following steps:

1. Implement the new structure.

2. Link GPOs into the new structure.

3. Move objects into the new structure.

4. Verify that everything works as expected.

In the next task, we will migrate back to the default scenario.

The main tasks for this exercise are as follows:

1. Create the new OU structure

2. Migrate the user and group accounts to the new OUs

▶ Task 1: Create the new OU structure

1. Switch to LON-DC1 and, if necessary, sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.

Note: For the following steps, you can also use the Windows PowerShell cmdlet **New-ADOrganizationalUnit**. The syntax is:

```
New-ADOrganizationalUnit –name OU-Name –path “parentDN”
```

2. Open Active Directory Administrative Center, and create the following OUs in the root of the domain:

- Name: **Central-IT**
- Description: **Admin accounts and Groups for Delegation. Administered from DOMAIN ADMINS ONLY**
- Name: **Enterprise**
- Description: **Enterprise-wide managed objects**
- Name: **Adatum**
- Description: **Regular objects for A. Datum**
- Name: **Contoso**
- Description: **Regular objects for Contoso**
- Name: **TreyResearch**
- Description: **Regular objects for Trey Research**

3. Run the **E:\Labfiles\Create-SubOUs.ps1** Windows PowerShell script to create the sub-OUs.

Use the provided Windows PowerShell script to create the Sub-OUs.

Create-SubOUs.ps1

```
$subous = @(`
("Admin-Accounts","ou=Central-IT,dc=adatum,dc=com","Admin-accounts only"),`
("Groups","ou=Central-IT,dc=adatum,dc=com","Groups for administrative tasks
delegation"),`
("Servers","ou=Enterprise,dc=adatum,dc=com","Enterprise-wide managed Servers"),`
("SQL","ou=Servers,ou=Enterprise,dc=adatum,dc=com",""),`
("WEB","ou=Servers,ou=Enterprise,dc=adatum,dc=com",""),`
("APP","ou=Servers,ou=Enterprise,dc=adatum,dc=com",""),`
("Users","ou=Adatum,dc=adatum,dc=com","Adatum regular user accounts"),`
("Groups","ou=Adatum,dc=adatum,dc=com","Adatum regular group accounts"),`
("Clients","ou=Adatum,dc=adatum,dc=com","Adatum clients"),`
("Servers","ou=Adatum,dc=adatum,dc=com","Adatum Servers"),`
("Marketing","ou=Users,ou=Adatum,dc=adatum,dc=com",""),`
("Sales","ou=Users,ou=Adatum,dc=adatum,dc=com",""),`
("Development","ou=Users,ou=Adatum,dc=adatum,dc=com",""),`
("IT","ou=Users,ou=Adatum,dc=adatum,dc=com",""),`
("Research","ou=Users,ou=Adatum,dc=adatum,dc=com",""),`
("London","ou=Clients,ou=Adatum,dc=adatum,dc=com",""),`
("Sydney","ou=Clients,ou=Adatum,dc=adatum,dc=com",""),`
("Toronto","ou=Clients,ou=Adatum,dc=adatum,dc=com",""),`
("London","ou=Servers,ou=Adatum,dc=adatum,dc=com",""),`
("Sydney","ou=Servers,ou=Adatum,dc=adatum,dc=com",""),`
("Toronto","ou=Servers,ou=Adatum,dc=adatum,dc=com",""))

$subous | %{New-ADOrganizationalUnit -Name $_[0] -Path $_[1] -Description $_[2]}
```

4. Verify that the OUs have been created.

▶ Task 2: Migrate the user and group accounts to the new OUs

1. On LON-DC1, switch to Windows PowerShell Integrated Scripting Environment (ISE).

2. For the Marketing, Sales, IT, Development, and Research departments, run the **E:\Labfiles\Move-AdatumUserGroups.ps1** Windows PowerShell script to move user and group accounts to their new OUs.

Use the provided Windows PowerShell script to move users and groups into their new OUs.

Move-AdatumUserGroups.ps1

```
$depts = @("Marketing","Sales","IT","Development","Research")
ForEach ($dept in $depts) {
$userDN = "OU=" + $dept + ",ou=users,ou=adatum,dc=adatum,dc=com"
$users = Get-ADUser -Properties Department -Filter {department -eq $dept}
    ForEach ($user in $users) {
        Move-ADObject $user -TargetPath $userDN
    }
$group = Get-ADGroup -Filter {Name -eq $dept}
Move-ADObject $group -TargetPath "ou=groups,ou=adatum,dc=adatum,dc=com"
}
```

3. Switch back to Active Directory Administrative Center, and verify that the user and group objects have been moved.

Results: After completing this exercise, you should have successfully implemented a part of the OU design.

Exercise 3: Designing and Implementing an Active Directory Permissions Model

Scenario

After designing the OU structure to support the future Active Directory administrative tasks model for A. Datum, you are tasked with determining how to implement groups and perform delegation in the new OU structure.

The following emails are sorted from newest to oldest. You should read the emails starting at the bottom of the email chain.

Supporting Documentation

Brad Sutton

From:	Bill Malone [Bill@adatum.com]
Sent:	27 September 11:27
To:	Brad@adatum.com
Subject:	Re: Active Directory Redesign - Done

Hello Brad,

That was fast, thank you very much. The OU design looks good. It is just as we wanted, and we can modify it in the future to provide Contoso and Trey Research with Active Directory-as-a-Service from our central team.

Would you be so kind as to prepare a proposal documentation? Would you then implement examples in the new OU structure to demonstrate how the administrative tasks delegation model would work? I'm really curious to see how that would be done.

If you can spare a few extra minutes, yesterday the management team discussed whether we should allow self-management of the "community groups and distribution lists" that our central team creates. It does not have security requirements, and the community members should be able to self-manage. It's a low priority, but it would be great since members of the community want to use Office Outlook to add additional members to their communities. I would like to discuss this option with you.

Thanks, and great job so far,

Bill

----- Original Message -----
From: Brad Sutton [Brad@adatum.com]
Sent: 27 Sep 08:45
To: Bill@adatum.com
Subject: Active Directory Redesign – Done

Attachment: A. Datum OU Redesign Proposal.docx

Good morning Bill,

Charlotte and I finished the proposal. We think that this is the best approach for a new, flexible OU structure that maintains the current infrastructure while enabling expansion when the time comes to migrate Contoso and Trey Research to our domain.

Let me know what you think.

Brad

A. Datum OU Administrative Tasks Delegation Model

Document Reference Number: BS00927/1

Document Author Date	Brad Sutton 27th Sept

Requirements Overview

This proposal documentation is an extension to BS00915/1, which describes the OU structure, and how to implement the Administrative Tasks Delegation Model that is supported by the OU structure.

While the previous proposal focused on OU design, this proposal describes how administrative delegation will be used.

The requirements for both tasks are as follows:

- For all administrative purposes, only use personalized Admin accounts.
- The Administrators group in A. Datum must be the only group of users who can make domain-wide changes, such as configuring Active Directory sites, creating top-level OUs, managing domain controllers, and creating and managing administrative user accounts and groups.
- The model must facilitate the future integration of Contoso and Trey Research (excluding the research.treyresearch.net domain) into the OU structure. AD DS replication and infrastructure services should be run by the central A. Datum team. However, it is up to Trey Research and Contoso to decide what levels of day-to-day administration of users, groups, and computers they would like to manage.
- A single group of administrators in the London office must be the only people who can create and delete all regular user and group accounts in the Adatum.com domain.
- Each regional hub office will have a local administrators group, which is responsible for troubleshooting local computer issues and supporting their users in their respective regions. The local administrators must have full access to all servers in their offices (other than domain controllers), and they must be able to support user issues on those client computers.
- Every department has a local Administrators group, which can reset their users' passwords.
- A. Datum has a dedicated team that is responsible for managing their enterprise-wide application servers. This team must manage access to the applications and resources that these servers host. Plan for the following three roles for these servers: SQL, WEB, and APP.
- Users from the Contoso.com domain and all other offices must access files on a file server that

A. Datum OU Administrative Tasks Delegation Model

is in Paris.

- As the number of users at A. Datum increases, the administrators want to automate adding multiple user accounts simultaneously, changing the attributes for a large number of user accounts and computers, and managing groups.

Summary of Information

- A. Datum:
  - Should be responsible for the Active Directory infrastructure.
  - Is managing administrative accounts and associated groups.
  - Includes three major locations: London, Toronto, and Sydney (smaller locations will not be included in the draft design).
  - Provides some enterprise-wide servers.
  - Has a dedicated team that creates and deletes all users.
  - Has Departmental Admins for user support, such as password resets.
  - Has Regional Hub Admins for client management, and for local server management.
- Trey Research:
  - Might move into the Adatum.com domain at some point in the future, and the OU structure should allow for this.
  - Must not integrate the Research.treyresearch.net domain.
- Contoso:
  - Might move into the Adatum.com domain at some point in the future, and the OU structure should allow for this.
  - Has a file server that provides services to all users in the enterprise.

Proposals

1. Which groups must you create to facilitate the high-level requirements?

2. What delegations must you configure?

The main tasks for this exercise are as follows:

1. Read the supporting documentation

2. Update the proposal document with your planned course of action

3. Examine the suggested proposals in the Lab Answer Key

4. Discuss your proposed solution with the class, as guided by your instructor

5. Create the required groups

6. Delegate permissions to management groups

7. To prepare for the next module

▶ Task 1: Read the supporting documentation

- Read the documentation provided.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the A. Datum OU Administrative Tasks Delegation Model document.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones shown previously.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

▶ Task 5: Create the required groups

1. On LON-DC1, switch to Active Directory Administrative Center, navigate to the **Central-IT\Groups** OU, and then create the following Role Groups:

- Name: **London-UserGroupprovisioning**
  - Type: **Global/Security**
- Name: **Enterprise-ServerOps**
  - Type: **Global/Security**
- Name: **Marketing-Admins**
  - Type: **Global/Security**
- Name: **London-Admins**
  - Type: **Global/Security**

2. In the Central-IT\Groups OU, create the following resource groups:

- Name: **acl_adatum-users_ccdc**
  - Type: **Domain Local/Security**
  - Member: **London-UserGroupProvisioning**
- Name: **acl_adatum-groups_ccdc**
  - Type: **Domain Local/Security**
  - Member: **London-UserGroupProvisioning**
- Name: **acl_enterprise-serveroperator**
  - Type: **Domain Local/Security**
  - Member: **Enterprise-ServerOps**
- Name: **acl_Users-Marketing_resetpwd**
  - Type: **Domain Local/Security**
  - Member: **Marketing-Admins**

- Name: **acl_clients-London_computer_ccdc**
  - Type: **Domain Local/Security**
  - Member: **London-Admins**
- Name: **acl_servers-London_computer_ccdc**
  - Type: **Domain Local/Security**
  - Member: **London-Admins**

▶ Task 6: Delegate permissions to management groups

Note: You can optionally complete this task by using the command-line tool Dsacls.exe. For example, to grant rights to create or delete group objects, the syntax is as follows:

```
dsacls ou=… /G adatum\acl_...:CCDC;group
```

For help with the tool Dsacls.exe, at a command prompt type **dsacls /?**.

1. Switch to Active Directory Administrative Center.
2. For the following OUs, grant the permissions to the groups in the **Advanced Security** dialog box of the OU properties:

- OU: **Adatum\Users**
  Group: **acl_adatum-users_ccdc**
  Permission: **Create User objects, Delete User objects**

Note: By using **dsacls**, you can perform the same task with the following command:

```
dsacls ou=users,ou=adatum,dc=adatum,dc=com /G adatum\acl_adatum-users_ccdc:CCDC;group
```

For help with the tool dsacls, at a command prompt type **dsacls /?**.

- OU: **Adatum\Groups**
  Group: **acl_adatum-groups_ccdc**
  Permissions: **Create Group objects, Delete Group objects**
- OU**: Adatum\Users\Marketing**
  Group: **acl_users-Marketing_resetPwd**
  Applies to: **Descendant User objects**
  Permissions: **Reset password**
- OU: **Adatum\Clients\London**
  Group: **acl_clients-London_computer_ccdc**
  Permissions: **Create Computer objects, Delete Computer objects**
- OU: **Adatum\Servers\London**
  Group: **acl_servers-London_computer_ccdc**
  Permissions: **Create Computer objects, Delete Computer Objects**

▶ 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 Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.

Results: After completing this exercise, you should have successfully designed and implemented an administrative permissions model.

Question: What was your suggested OU design? What were the reasons behind your design decisions?

Question: While the lab had you use Windows PowerShell to move user objects based on a certain attribute, can you think of other ways to do this?

Question: Bill suggested self-management for certain groups. How would you implement this? What are the benefits and what are the risks associated with this recommendation?

Module Review and Takeaways

Best Practices:

- Use the AG(U)DLP model when designing your group strategy. By doing this, accounts are grouped in global groups for the business roles. If required, you can consolidate these groups across domains in a universal group. The role groups are then assigned via Domain Local groups that grant access to the specific resource.
- Design your Active Directory administrative tasks model with least privileges in mind. As a best practice, make a list of tasks in your organization, and then grant each task to a specific team. If a team wants the permissions, they become responsible for those tasks.
- Use scripts to implement your design. Review the Windows PowerShell cmdlets and the Dsacls.exe tool for setting permissions.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
When using the Security Delegation Wizard to set permissions, once the wizard reopens, the permissions do not display.	
How to determine what attributes must be delegated.	
How to change the security settings of an attribute that does not display in the Advanced Security dialog box.	

Review Questions

Question: Why is it a good idea to implement the least privileges required when delegating administrative tasks?

Question: Why should you use administrative accounts and store them in a different location than regular user accounts?

Question: What must you consider when you want to migrate your OU structure to a new model?

Module 8

Designing and Implementing a Group Policy Object Strategy

Contents:

Module Overview

The Group Policy environment in an Active Directory® Domain Services (AD DS) infrastructure is the Active Directory–integrated technology that you can use to configure and manage clients and servers that are joined to your domain. An effective Group Policy design provides you with a more standardized and manageable environment in which you can perform many administrative tasks.

This module introduces the key concepts for designing Group Policy as they relate to planning, implementing, and management in AD DS. It also describes the best practices for designing Group Policy Objects (GPOs) and GPO inheritance.

Objectives

After completing this module, you will be able to:

- Collect and analyze the information required to facilitate a GPO design.
- Create a GPO design and implement it.
- Create a GPO processing design.
- Plan GPO management.

Lesson 1
Collecting the Information Required for a GPO Design

When considering design changes or planning tasks for your organization's AD DS environment, it is important that you configure the AD DS resources properly and secure them.

Before designing a Group Policy environment, you must have a clear idea of what it is that your organization needs from the environment. You also will need to understand your AD DS structure, and how you can leverage that structure using your Group Policy implementation.

Lesson Objectives

After completing this lesson, you will be able to:

- Collect organizational information.
- Collect information regarding security requirements.
- Collect information regarding desktop management requirements.
- Collect information regarding administrative processes.

  Question: How is Group Policy used in your organization? What are the issues you face or have faced in your organization regarding Group Policy? What settings or tasks would you would prefer to do with GPOs, but are unable to?

Collecting Organizational Information

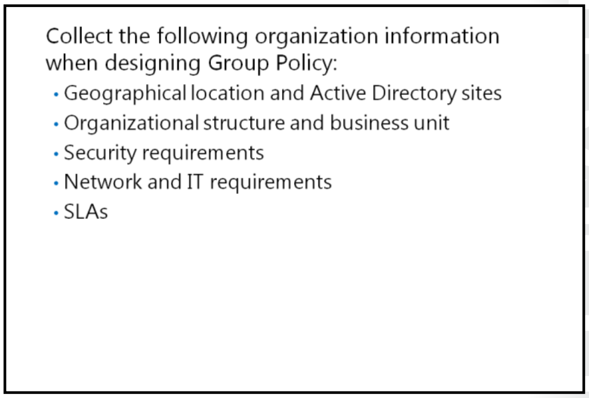

Collect the following organization information when designing Group Policy:

- Geographical location and Active Directory sites
- Organizational structure and business unit
- Security requirements
- Network and IT requirements
- SLAs

Within the AD DS structure, Group Policy provides you with the ability to configure user and computer environments. You also can use Group Policy to enforce the technical aspects of organizational policies and legal requirements. Collecting information relevant to your organization is the first step in planning an effective Group Policy design.

When you are starting your design process for Group Policy, you must collect the following organizational information:

- Company locations and AD DS sites by location. Your organization's geographical layout and your network's topology are important considerations. The needs of each of your locations and your AD DS sites configurations influence the GPOs that you design and create, and influence how you link and apply those GPOs within the AD DS structure.
- Organizational structure and business unit. Similar to the geographical structure, organizational structure also has a significant impact on your GPO design. In many cases, your organizational unit (OU) structure will reflect your organization's administrative model, and you will need to apply the GPOs that you create to accommodate the requirements for each business unit. For example, if each department in your organization has a different set of desktop settings, you should design your OU structure to implement those settings.

- Security requirements. Security is one of the first and most important areas that you configure by using Group Policy. You typically use Group Policy in the Active Directory environment to configure password policies, account lockout settings, and general AD DS security measures such as authentication methods. Additionally, you can configure the Windows® Firewall settings in Group Policy. Depending on the security requirements of your server and application infrastructure, you also can design the application communication infrastructure through Windows Firewall by using GPOs.
- Network and IT requirements. Many administrative practices and technical aspects of your environment will require GPOs to function properly, or will benefit from the administrative ease that Group Policy can provide. You should consider IT requirements such as application-specific Group Policy settings. Additionally, you should collect information about how you will manage the Group Policy infrastructure. This information will include delegated administration and GPO ownership.
- Service level agreements (SLAs). Group Policy is a key method of SLA enforcement within the Active Directory environment. In many cases, SLAs are established for your organization, your departments, or individual components and applications. These SLAs might require, or benefit from one or more Group Policy settings. As you design your Group Policy infrastructure, you need to consider which settings should be combined with which GPOs, who is administering the GPOs, and at which level they are applied to ensure that SLAs are met.

Collecting Information on Security Requirements

When creating a Group Policy design, you need to collect information pertaining to your enterprise's security requirements. When collecting this information, you need to:

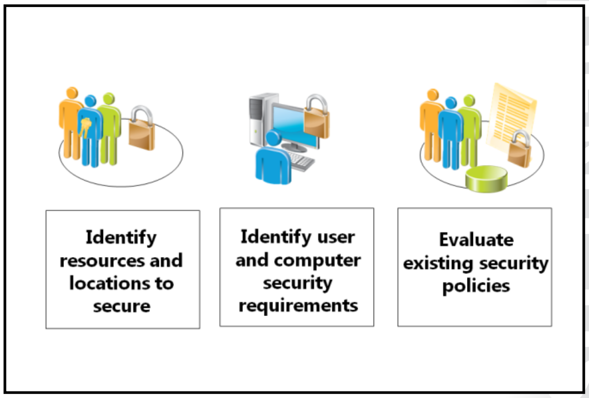
Identify resources and locations to secure
Identify user and computer security requirements
Evaluate existing security policies

- Identify the resources that you need to secure. In your Active Directory environment, you need to account for and include the resources that you need to secure in your Group Policy security plan. Resources include any user objects (standard users, administrative users, and service accounts), and computer objects (client computers, laptops, tablets, kiosks, servers, and classroom or lab computers), within AD DS that are joined to the domain, and to which you may need to apply Group Policy settings.
- Identify the locations that you need to secure. OU-based geographical locations, and locations defined in Active Directory sites based on topographical locations may possess specific security requirements. Consider the business activity that occurs within each of these locations and whether the locations require specific security settings such as drive encryption, specific Windows Firewall settings, or limited access to applications or administrative control.
- Identify the various user and computer security requirements. User and computer security requirements typically are easy to identify, but they also require the most effort for planning and implementation. Security requirements usually affect the IT environment directly. You apply them to the domain, site, or OU that contains the applicable users and computers. These requirements can include settings for local logon rights, and modifications to administrative access for certain users or computers.
- Consider the difference between public-facing servers and intranet servers. Public-facing servers require more specific and stringent hardening policies to protect them against hackers. Internal servers that are connected to trusted networks typically do not need strict hardening policies.

- Evaluate existing corporate security policies. Examine general and written policies, and those that you already implement using Group Policy. Written, corporate policies may contain security requirements that you need to implement in your environment by using Group Policy. Often, policies that involve limited access rights and privacy laws require the implementation of both Group Policy and other security layers such as file and folder permissions and auditing.

Collecting Information on Desktop Management Requirements

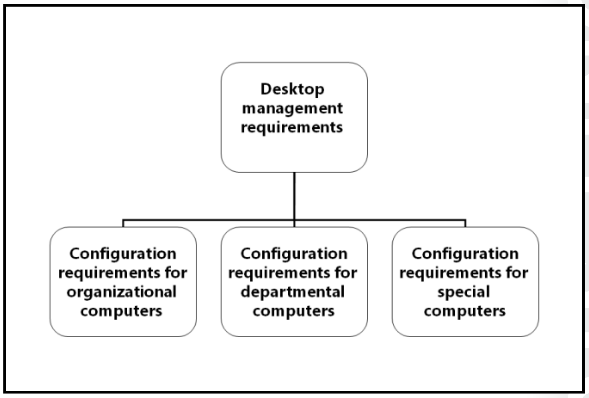
Desktop management requirements
Configuration requirements for organizational computers
Configuration requirements for departmental computers
Configuration requirements for special computers

Group Policy settings define user desktop settings that you need to manage. When designing GPOs for desktop management, you need to collect the following information:

- Organizational requirements for computer configuration. You may need to set and apply certain policy settings to all of your organization's computers. Your design must determine the policy settings that are applicable to the entire organization. You then must link those policy settings to a location in the AD DS structure, and ensure that the GPOs are applied to all computers.

  GPOs that you apply to a domain affect all of its users and computers. You can filter GPOs to apply only to a specific group of users or computers, or you can use Windows Management Instrumentation (WMI) filters. However, as a best practice, avoid filtering unless it is the only way to achieve the goal.

- Departmental requirements for computer configuration. You can create OUs to organize and manage objects for each department, business unit, and geographical location. You then can apply GPOs to OUs or sites that contain multiple computers with similar requirements.

Note: In a multi-domain environment, you must be careful when assigning a GPO to a site. The GPO files are contained in SYSVOL, and these files will always be created in the domain in which you are creating the GPO. Sites do not necessarily map to domains. Therefore, you might assign the GPO to a site that does not have a domain controller for the domain in which the GPO was created. When you are creating site policies, if possible, try to create them in domains that span all sites.

- Special requirements for computer configuration. You can meet special computer configuration requirements by applying Group Policy to dedicated OUs containing computer objects that you need to manage. For example, you may need to apply specific settings to kiosk computers in a public place, or to laptops that belong to executive management.

Collecting Information on Administrative Processes

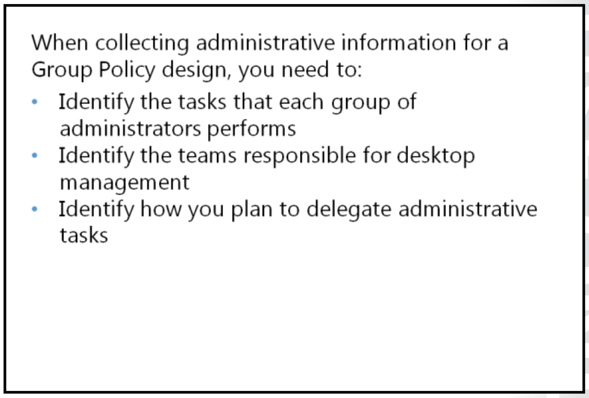

When collecting administrative information for a Group Policy design, you need to:

- Identify the tasks that each group of administrators performs
- Identify the teams responsible for desktop management
- Identify how you plan to delegate administrative tasks

Applying a Group Policy design depends on the model that you are using for administrative control and delegation. When collecting administrative information for a Group Policy design, consider the following:

- Identify the administrators and their administrative tasks. The users that will be administering your environment and the tasks that they perform will determine how you apply Group Policy settings to restrict or enable administrative control over your Active Directory environment. You should consider administrators and their requirements when you design Group Policy settings that control local logon to servers, access to administrative tools, and other administrative tasks.
- Identify the teams that are responsible for desktop management. Users outside of the core administrative group may also require special consideration when designing Group Policy in your Active Directory environment. Consider these users when you design Group Policy settings that remove management capabilities, such as Control Panel access or registry access. Additionally, when you apply GPOs, you should account for the scope of the GPO. For example, GPOs that allow access to Control Panel for Toronto users should not apply to resources in other locations.
- Identify how you plan to delegate administrative tasks. You should account for delegating administrative tasks with limited rights to additional users, such as branch manager, application server operators, or helpdesk personnel. For these users, you may need to configure different GPO settings to override the GPO settings that apply higher in the GPO structure.

Lesson 2
Designing and Implementing GPOs

Each GPO within your environment represents an important element of configuration management. Designing GPOs includes assessing your environment, identifying the configuration that your resources require, and identifying the Group Policy tools necessary to implement those configuration changes.

This lesson describes the information that you need to design GPOs that effectively meet your organization’s needs.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the options in Group Policy settings.
- Design administrative templates.
- Manage GPO storage.
- Design Group Policy Preferences.
- Implement best practices for GPO design.
- Describe the technologies you would complement with GPOs and their use.
- Implement GPOs.

Overview of Group Policy Settings

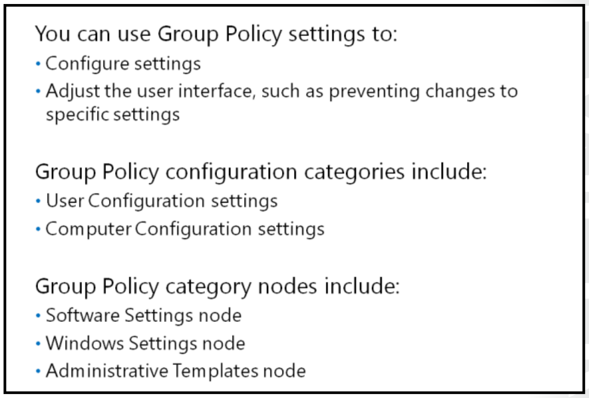

You can use Group Policy settings to:

- Configure settings
- Adjust the user interface, such as preventing changes to specific settings

Group Policy configuration categories include:

- User Configuration settings
- Computer Configuration settings

Group Policy category nodes include:

- Software Settings node
- Windows Settings node
- Administrative Templates node

In the AD DS environment, every GPO contains settings that allow administrators to configure the operating system behavior of computers, or to modify user settings in a centralized way to preconfigure or enforce certain configuration options.

How Group Policy Works

Group Policy in AD DS works with the group policy engine and Group Policy client-side extension dynamic-link libraries (DLLs) that reside on Windows operating systems. The group policy engine uses the Group Policy client-side extensions to perform required changes and modifications to the operating system environment, primarily by making changes to the Windows operating system registry. After these changes are made, the group policy engine manages them, and depending on the configuration, you can prevent users from modifying them. Most settings generally have two mechanisms that you should consider:

- Settings configurations. You can use GPOs to configure certain settings, such as the proxy server (which is used by Internet Explorer®), the home page, or the screen saver configurations.
- Locking administrative interfaces. Although the GPO configuration settings are written when the GPO is applied, they are not enforced. By default, the user can change many of these settings, which is appropriate for settings such as such as the screen saver or the home page. However, you should consider which settings are preconfigured, which settings the user is allowed to change, and which

settings should be enforced. (Examples of settings include preventing the user from changing the proxy server, or preventing the user from disabling BitLocker® Drive Encryption.) Therefore, you should remove the user's ability to change these specific settings. You can do this by using GPOs.

By using both mechanisms, the administrator can ensure that configuration changes made by using Group Policy settings are persistent, enforced, and mandatory.

GPO Settings Structure

Within each GPO, Group Policy settings are classified into two categories:

- User Configuration settings. The settings that you configure within the User Configuration node apply to user objects such as OUs, domains, and sites. GPOs link to these objects in the Active Directory structure. The settings that you configure in the User Configuration node do not apply until a user to which the GPO applies logs on to the computer.
- Computer Configuration settings. The Computer Configuration settings apply to the Windows operating system environment, based on the computer account object. These settings apply when the computer starts, but prior to user logon. This behavior makes the Computer Configuration settings a better candidate for general security-related policies.

Within each configuration category, there are three different nodes for Group Policy settings. They are:

- Software Settings node. The Software Settings node contains settings that relate to software installations.
- Windows Settings node. The Windows Settings node contains settings that modify items such as scripts and folder redirection. This node also contains a broad set of security settings that are critical to implementing your domain's security by using Group Policy.
- Administrative Templates node. The Administrative Templates node contains the majority of the individual settings within Group Policy. The default settings in this node provide control over almost all aspects of the Windows operating system environment. You can add additional nodes to the Administrative Templates node by importing .adm or .admx/.adml (Administrative template) files that contain definitions for other Group Policy settings.

Note: Administrative templates are discussed in greater detail later in this module.

The settings in these nodes can be either single-value (Enabled, Disabled, or Not Configured), or multi-value, such as Internet Explorer proxy settings. This means that nodes can contain several configuration details within a single Group Policy setting.

In addition to these policy settings, Group Policy Preferences settings also exist within each GPO object that you create in the Windows Server® 2012 environment. Group Policy Preferences settings were introduced in Windows Server 2008, and apply only to Windows Vista® and newer Windows client operating systems, or to Windows Server 2008 and newer Windows Server operating systems. Group Policy Preferences are discussed in greater detail later in this module.

Considerations for Designing Administrative Templates

Group Policy administrative templates contain the majority of configurable settings within a GPO. A default GPO contains the Administrative Templates node, which in turn contains the Computer Configuration node and the User Configuration node. In Windows Server 2012 R2, each of these nodes contains the following nodes:

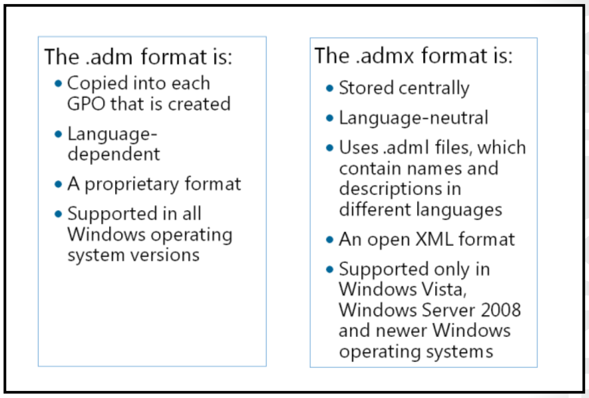

The .adm format is:

- Copied into each GPO that is created
- Language-dependent
- A proprietary format
- Supported in all Windows operating system versions

The .admx format is:

- Stored centrally
- Language-neutral
- Uses .adml files, which contain names and descriptions in different languages
- An open XML format
- Supported only in Windows Vista, Windows Server 2008 and newer Windows operating systems

- Computer Configuration:
  - Control Panel
  - Network
  - Printers
  - Server
  - Start Menu and Taskbar
  - System
  - Windows Components
- User Configuration:
  - Control Panel
  - Desktop
  - Network
  - Shared Folders
  - Start Menu and Taskbar
  - System
  - Windows Components

There also is an **All Settings** node, which contains the combined settings for all other nodes.

Extending the Administrative Templates Node

Unlike the other settings nodes in a GPO, you can customize the Administrative Templates node. You can import administrative templates for Microsoft® applications or non-Microsoft applications, or import custom-developed administrative templates that extend Group Policy's configuration capabilities. For example, Microsoft provides an administrative template for Microsoft Office 2013 and previous Office version applications, which enable administrators to use Group Policy settings to configure the behavior of those applications. You should consider which applications in your environment you should configure by using a Group Policy administrative template.

.adm and .admx/.adml Files

Windows Server 2003 and earlier Windows Server operating systems store the Group Policy administrative template files in a proprietary format with the .adm extension. When you import these .adm files, they become part of the Group Policy structure in AD DS and part of every GPO that you create. The .adm templates are stored in every GPO, and are replicated within SYSVOL. This can make the sum of GPOs rather large.

Windows Server 2008 and newer Windows Server editions store the Group Policy templates as .admx files. These .admx files are XML files that are much easier to develop and manage than .adm files, and are language-neutral. In addition to .admx files, .adml files contain language translations, which enable administrators of different languages to view the names and descriptions of GPO settings in their preferred language. You can apply GPOs based on .admx and .adml files to clients and servers by using computers that run Windows Vista, Windows Server 2008, or newer Windows operating systems.

You can import .adm, .admx, and .adml administrative templates into a GPO in Windows Server 2008 and newer Windows Server operating systems. However, unlike .adm-based administrative templates, you import the .admx administrative templates by copying the associated .admx files into the %SYSTEMROOT%\Policy Definitions folder, (or to the Central Store, which is discussed in the next topic). You must store the .adml files in a subfolder that states the language, for example, en-US for US-English, or de-DE for German.

After you copy the template files into the folder structure, the administrative template settings that are defined in the .admx file become available for configuration in GPOs.

Managing GPO Storage

The Group Policy Management Console (GPMC) stores GPOs under an administrative node named *Group Policy Objects*. In AD DS, the Group Policy Objects node is the container cn=Policies, cn=System, dc=*domain*. Every GPO is stored underneath that node as a special container of the Active Directory class **groupPolicyContainer**, and contains some general settings, such as whether the GPO applies to user or computer objects. The settings within the GPO are stored in the file system and are replicated by using the SYSVOL folder structure.

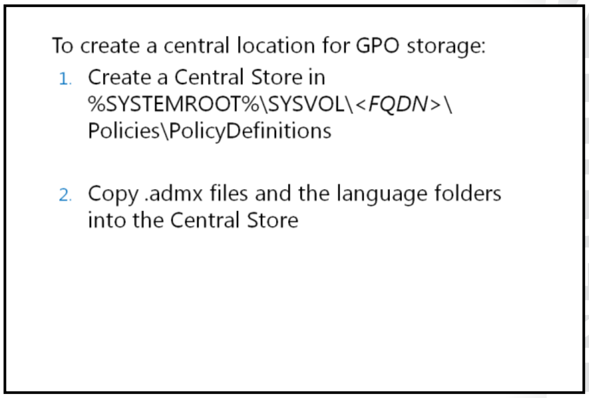

To create a central location for GPO storage:

1. Create a Central Store in %SYSTEMROOT%\SYSVOL\<*FQDN*>\Policies\PolicyDefinitions

2. Copy .admx files and the language folders into the Central Store

Storage Options for Group Policy Templates

Additional Group Policy components are stored outside of the AD DS database and the SYSVOL folder. In Windows Server 2008 and newer Windows Server operating systems, the %SYSTEMROOT%\PolicyDefinitions folder stores the .admx administrative template files with the .adml language files, by default. This storage schema means that each domain controller may potentially have a different set of administrative templates available to the Group Policy infrastructure. If you frequently use custom or non-Microsoft .admx/.adml files, storing these files locally on each domain controller can lead to an inconsistent Group Policy editing environment or behavior. To mitigate the potential issues, you should manually update the local PolicyDefinitions folder on each domain controller or write a script that automates this update process.

The Central Store

To utilize replication that occurs between AD DS domain controllers, you can create a Central Store for .admx /.adml files. The Central Store is a file location that the Group Policy tools check when creating and editing GPOs. Because the .admx/.adml files are in the SYSVOL folder, they replicate to all domain controllers in the domain.

To create a Central Store for .admx/.adml files, create the following folder location on any domain controller, with *FQDN* as the fully qualified domain name: %SYSTEMROOT%\SYSVOL\<*FQDN*>\Policies\PolicyDefinitions.

For example, to create a Central Store for the treyresearch.net domain, create the following folder on a domain controller in the treyresearch.net domain: %SYSTEMROOT%\SYSVOL\treyresearch.net\Policies\PolicyDefinitions.

You then should store all .admx-based administrative template files in this location, and then copy the language folder (for example, en-US) with the .adml files to the same location.

Considerations for Group Policy Storage

When you are designing Group Policy storage, consider the following:

- Use the Central Store if your organization uses custom or additional .admx and .adml files.
- All computers that are being used for administering GPO settings will first attempt to locate the Central Store for administrative templates.
- Domain-joined clients or domain-joined servers do not need to contact the Central Store if they only apply GPOs.

If a Central Store does not exist, clients used for Group Policy management will use the files in their local %SYSTEMROOT%\PolicyDefinitions folder. This can be problematic if the client computer does not contain the same administrative templates as a domain controller or any other administrative client that was used to edit the GPO.

Considerations for Designing Group Policy Preferences

Group Policy Preferences include a range of configurable GPO settings. Configurable GPO settings include mapping a network drive, modifying registry keys, and creating and sharing local folders and files. A GPO stores these settings in its Preferences node, which is included in the root of both the Computer Configuration and the User Configuration nodes. Previously, these kinds of settings required you to implement them separately from Group Policy, often by using a logon script.

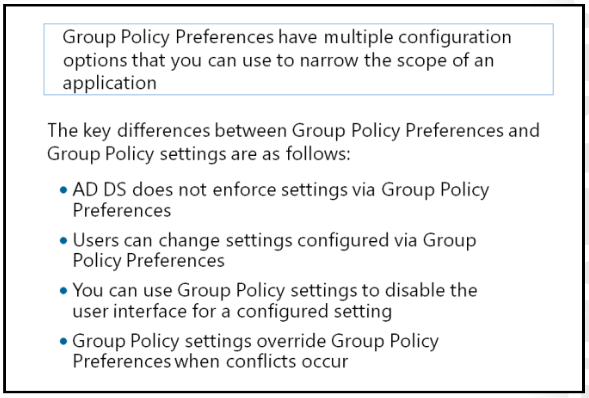

Group Policy Preferences have multiple configuration options that you can use to narrow the scope of an application

The key differences between Group Policy Preferences and Group Policy settings are as follows:

- AD DS does not enforce settings via Group Policy Preferences
- Users can change settings configured via Group Policy Preferences
- You can use Group Policy settings to disable the user interface for a configured setting
- Group Policy settings override Group Policy Preferences when conflicts occur

The following operating systems support Group Policy Preferences natively:

- Windows Server 2008 and newer Windows Server operating systems
- Windows Vista Service Pack 2 (SP2) and newer Windows client operating systems

For Windows Server 2003 and Windows Vista Service Pack 1 (SP1) and older Windows operating systems, you can download and install Group Policy client-side extensions to provide support for preferences on those systems.

Group Policy Preferences are similar to Group Policy settings in that they apply configurations to the user account or to the computer. However, how you configure and apply Group Policy Preferences differs. The differences between Group Policy Preferences and Group Policy setting configuration options are as follows:

- Unlike Group Policy settings, AD DS does not enforce Group Policy Preference settings.
- You can use Group Policy settings to disable the user interface for settings that the policy manages, but Group Policy Preferences in general do not have that option.

- AD DS applies Group Policy settings at regular intervals, but AD DS applies Group Policy Preferences only once, or at specific intervals.
- The end user can change preference settings that are applied through Group Policy, but you can configure Group Policy settings to prevent users from changing these preference settings.
- In some cases, you can configure the same setting though a Group Policy setting and a Group Policy Preference. If you configure and apply conflicting settings to the same object, the value of the Group Policy setting takes precedence over the Group Policy Preference.
- Group Policy Preferences overwrite original settings, while Group Policy settings generally do not.
- Some Group Policy Preferences do not work if they are not run in the right security context. By default, the SYSTEM account is used for processing Group Policy Preferences. However, you can enable the Run in logged-on user's security context option in the GPO to switch processing from the SYSTEM account to the signed in user account. This can be helpful when mapping drives or in other situations that require the permissions of the user account.

Configuring Group Policy Preferences

Group Policy Preferences have several configuration options that you can use to manage how AD DS applies the preferences, including the following:

- Stop processing items in this extension if an error occurs. If an error occurs while processing a Group Policy Preference, no other preferences in this GPO will process.
- Run in logged-on user's security context (user policy option). Group Policy Preferences can run as the System account for the logged-on user. This option forces the logged-on user context, and is therefore only available on Group Policy Preferences in the User Configuration node.
- Remove this item when it is no longer applied. Unlike most policy settings, AD DS does not remove Group Policy Preferences when the GPO that delivered it is removed. This option changes that behavior and has AD DS remove the Group Policy Preference when the GPO is removed.
- Apply once and do not reapply. Normally, AD DS refreshes Group Policy Preferences at the same interval as Group Policy settings. This option changes that behavior to apply the setting only once on logon or startup.
- Item-level targeting. This option enables you to specify criteria that determines exactly which users or computers will receive a Group Policy Preference. Criteria includes:
    - Computer name
    - IP address range
    - Operating system
    - Security group
    - User

Best Practices for GPO Design

Although there are several key factors that you should consider when you design GPOs, one of the most important factors to consider is manageability.

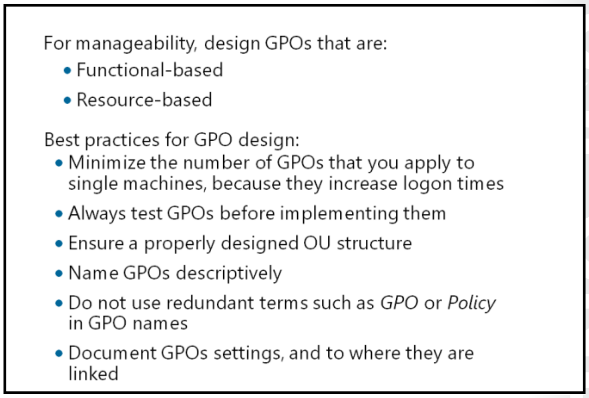

For manageability, design GPOs that are:

- Functional-based
- Resource-based

Best practices for GPO design:

- Minimize the number of GPOs that you apply to single machines, because they increase logon times
- Always test GPOs before implementing them
- Ensure a properly designed OU structure
- Name GPOs descriptively
- Do not use redundant terms such as *GPO* or *Policy* in GPO names
- Document GPOs settings, and to where they are linked

Managing GPOs

The number of GPOs in your Group Policy design and the complexity of each GPO's settings will affect the manageability of the Group Policy infrastructure. Typically, you base GPO creation and naming on either function or resources. Depending on your organization's needs, your environment may contain either type of GPOs, or both.

GPOs Based on Functions

GPOs that contain fewer and more specific settings enable you to create a more specific name for each GPO based on its functionality. For example, you may create GPOs named Limit Control Panel Access, Map Sales Drives, and Disable Administrative Tools. This type of GPO naming method makes it easier to determine the configuration that your Group Policy is imposing. When Group Policy applies these configurations to a group of computers in your environment, you can link the corresponding Group Policy to the applicable OUs. You also can reuse GPOs, and link them to multiple groups of computers that require the same type of functionality. However, this method also requires that you create and administer many GPOs and many GPO links.

GPOs Based on Resources

When naming GPOs based on resources, you name the GPO after the resources to which you are applying it. For example, you may create a single GPO named New York Laptops, and place all Group Policy settings that need to apply to the New York location's laptop computers within this GPO. The general design and naming scheme of this method describes more of what the GPO applies to rather than what it does. This method typically results in fewer GPOs overall, unless you have a large number of resource groups. This method also makes it easier for administrators to link and apply GPOs within the structure, because the GPO is already named for the resource it is designed to manage.

Best Practice for Managing GPOs

GPOs significantly affect logon times and boot up times of users and computers. Therefore, a smaller number of GPOs is preferred. We recommend that you join all settings that apply to a larger group of clients or servers at the highest level, such as corporate clients, location-based clients, or clients based on the type of device (desktop, laptop, or tablet). You can then create smaller policies for specific settings that only apply to a subset of computers or users.

Other GPO Best Practices

Consider the following best practices for GPO design and implementation:

- Always test the effects of a GPO in a test environment before applying it to your production's Active Directory domain structure. If you ensure that you use test computers and test users, you can create a Test-Branch in your OU structure with which you can test the policies before applying them to the broad range of users.
- Create a properly designed OU structure to ensure that naming and assigning GPOs is a more straightforward task. Wherever possible, restructure your OUs according to the best practices discussed previously in this module.

- Name your GPOs with a reasonable amount of specificity. Administrators should be able to determine both the general function of a GPO or the resources to which a GPO applies—and sometimes both—from the GPO's name.
- Do not use the terms *Policy* or *GPO* in your GPO names. These terms are redundant, and use characters that you could otherwise use for descriptive naming.
- Document GPO settings, and where the GPOs are linked. This helps when troubleshooting issues. You can use the Group Policy Management Console (GPMC) and its associated scripts to export a report of Group Policy settings, which you can then store on a file share or a Microsoft SharePoint® site for your administrative staff.

Alternatives to Using GPOs

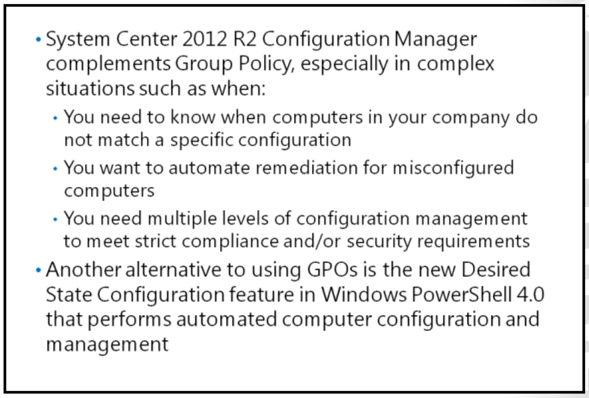

- System Center 2012 R2 Configuration Manager complements Group Policy, especially in complex situations such as when:
  - You need to know when computers in your company do not match a specific configuration
  - You want to automate remediation for misconfigured computers
  - You need multiple levels of configuration management to meet strict compliance and/or security requirements
- Another alternative to using GPOs is the new Desired State Configuration feature in Windows PowerShell 4.0 that performs automated computer configuration and management

GPOs are in widespread use across the vast majority of companies. However, administrators often find the need to supplement GPOs with other system management software and services. One of the most popular products to supplement or complement GPOs is System Center 2012 R2 Configuration Manager. Configuration Manager is a fully featured configuration management application. It is one of the components of the System Center suite and it integrates closely with some of the other System Center components such as System Center 2012 R2 Operations Manager. You can combine GPOs and configuration management software to provide expanded capabilities or more efficient operational procedures in scenarios, such as the following:

- You need to know when computers in your company are no longer matching a specific configuration. For example, you might want to know when the membership of the local Administrators group does not match the desired membership configured in a restricted group's GPO. Although a GPO can configure aspects of a computer (such as restricted groups in this scenario), it can only do so when the computer is in scope of the GPO and when the computer is processing Group Policy correctly. Administrators can use the Desired Configuration Management feature in Configuration Manager to report on or enforce a specific configuration.
- You want to automate the remediation of misconfigured computers. In this scenario, Operations Manager is monitoring computers and discovers that one of the critical Windows services is in a stopped state. Configuration Manager can remediate the service automatically by starting the service. Although an administrator can configure a GPO to manage services and ensure that these services are set to start automatically, a GPO only functions when computers are in the scope of the GPO and when computers process GPOs correctly.
- You need to have multiple levels of configuration management to meet strict compliance and/or security requirements. In such a scenario, by combining Group Policy and system management software, you can provide multiple layers of configuration management. Similar to a layered security strategy (which is considered a best practice), layering management software can often provide the most consistent configuration for a corporation.

Another alternative to using GPOs is the new Desired State Configuration feature in Windows PowerShell® 4.0. This feature offers new functionality to control environments configurations. It can perform automated configuration and management of computers, including:

- Ensuring that a Windows role or feature is installed on a group of computers.

- Ensuring that specific Windows services are set to start automatically and restart if stopped.
- Ensuring that the membership of the local Administrators group contains the Domain Admins group.

Additional Reading: To learn more about the Windows PowerShell Desired State Configuration feature, visit http://go.microsoft.com/fwlink/?LinkID=391884.

Demonstration: Implementing GPOs

Consider the following scenario for implementing GPOs. At A. Datum Corporation, the Chief Information Officer has recently announced a security policy that passwords must not be older than 45 days. After implementing those policies by using fine-grained password policies, administrators discovered that the Directory Services Restore Mode Administrator password is not easy to change. This password needs to be reset on every domain controller by using the command-line tool Ntdsutil. Furthermore, resetting this password is very hard to script because the password must be entered twice, and in this scenario, using an input file does not work.

However, one of the administrators has discovered that you can synchronize a user password. Therefore, A. Datum has decided to create a policy and a scheduled task on the domain controllers to ensure that all domain controllers change the Directory Services Restore Mode Administrator password when the administrators reset the password on a specific service account.

In this demonstration, you will see how to:

- Create a Directory Services Restore Mode service user.
- Create a GPO.
- Create a scheduled task by using Group Policy Preferences.
- Link the policy to the Domain Controllers OU.

Demonstration Steps

Create a Directory Services Restore Mode service user

1. Switch to LON-DC1.
2. Open the Active Directory Administrative Center, and create a new user in the Users container with the following settings:
   - Name: **srv_dsrm**
   - Password: **Pa$$w0rd**
   - Password options: **Password never expires**
3. Disable the user **srv_dsrm**.

Create a GPO

- Create a Group Policy with the name **DSRM_Pwd**.

Create a scheduled task by using Group Policy Preferences

- Edit the Group Policy **DSRM_Pwd** under **Computer Configuration\Preferences**, and create a scheduled task with the following settings:
  - Action: **Create**
  - Name: **Sync DSRM Password**

- Runs as: **System**
- Runs whether the user is logged on or not: **Enabled**
- Do not store password. The task will only have access to local resources: **Enabled**
- Trigger: Under **On a schedule**, specify the following settings:
  - Settings: **Daily**
  - Recur every: **1 days**
  - Repeat task every: **1 hour**
- Action: **Start a program**
- Program: **C:\windows\system32\ntdsutil.exe**
- Arguments: **"set dsrm password" "sync from domain account srv_dsrm" quit quit**

Link the policy to the Domain Controllers OU

- Link the GPO **DSRM_Pwd** to the **Domain Controllers** OU.

Lesson 3
Designing GPO Processing

How you apply GPOs within your Active Directory domain environment determines which settings apply to users and computers. A poorly designed GPO implementation can prevent settings from affecting the appropriate users and computers, and can make troubleshooting the Group Policy environment difficult. Effective design of Group Policy processing involves knowing how Group Policy processes settings, and understanding the tools and methods available to make the processing of those Group Policy settings as effective as possible.

Lesson Objectives

After completing this lesson, you will be able to:

- Design Group Policy inheritance.
- Design Group Policy filtering.
- Describe slow link detection.
- Design Group Policy processing.
- Configure GPO inheritance and filtering.

Considerations for Designing Group Policy Inheritance

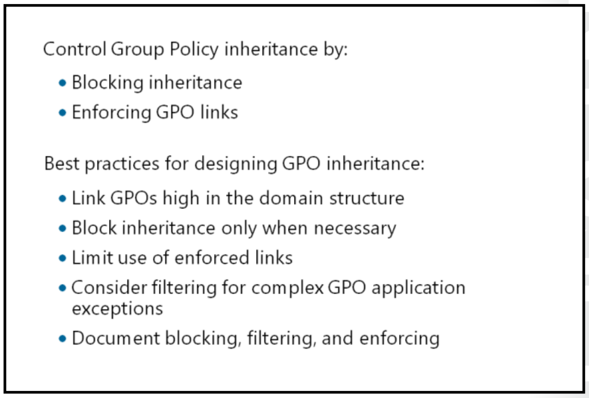

Control Group Policy inheritance by:

- Blocking inheritance
- Enforcing GPO links

Best practices for designing GPO inheritance:

- Link GPOs high in the domain structure
- Block inheritance only when necessary
- Limit use of enforced links
- Consider filtering for complex GPO application exceptions
- Document blocking, filtering, and enforcing

You can create and link multiple GPOs under any OU within AD DS. You should define the settings that you want to apply to a broad set of an organization's users or computers by using GPOs that link to parent containers. Child containers inherit the settings of parent containers through Group Policy inheritance. Group Policy inheritance combines the GPO settings that apply to parent containers with the settings that link to the child containers.

You can use the following methods to control inheritance within the Group Policy environment:

- Block inheritance. You can block policy inheritance for a domain or an OU. Blocking inheritance prevents child objects from inheriting settings from any GPOs that are linked to parent containers.
- Enforcement. You can set a GPO link to take precedence over any child object's settings. With this enforcement, if the GPO contains conflicting settings the parent GPO link always has precedence over the child GPO link. Enforced GPOs override block inheritance.

You should use block inheritance and GPO enforcement only if you are unable to use other options. This is because block inheritance and GPO enforcement are likely to impact future Group Policy troubleshooting.

Best Practices for Designing GPO Inheritance

When designing a Group Policy infrastructure for your domain, consider the following GPO inheritance best practices:

- Link GPOs as high as possible in your Active Directory domain structure. This allows GPOs—particularly those with general, domain-wide settings—to take advantage of inheritance. Linking to GPOs higher in the structure will reduce the number of GPO links that you have to create or manage when you add additional OUs at lower levels.

- Block inheritance only when necessary. Blocking inheritance can be a useful option for providing certain resources in your AD DS structure with a set of GPOs that differ from other GPO resources. However, overusing block inheritance can lead to a confusing and inconsistent GPO environment. Administrators who apply GPOs at a higher level may not be aware of areas that you have blocked in your inheritance structure.

- Enforce links only for critical and security-related GPOs. Enforcement overrides potential exception methods that apply further down in the structure, such as blocking inheritance and blocking GPOs that apply to child OUs. Therefore, you should use enforcement for GPOs that contain settings that you want to apply to all clients.

- Consider using options such as security groups filtering and WMI to configure specific exceptions to GPOs that you want to apply at high levels in the OU structure. Filtering often provides a greater level of control over GPO application, and enables you to apply GPOs to objects in different OUs.

- Document your GPO environment. Documenting your GPO environment should include detailing the levels where GPOs are blocked, and documenting specific settings such as enforcing and filtering GPOs. This will enable other administrators to understand the GPO design.

Considerations for Designing Group Policy Filtering

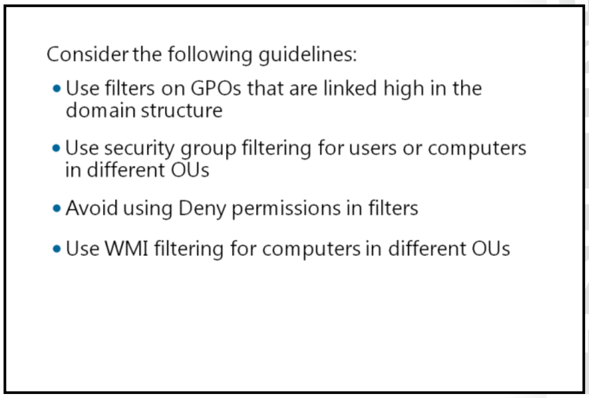

Consider the following guidelines:

- Use filters on GPOs that are linked high in the domain structure
- Use security group filtering for users or computers in different OUs
- Avoid using Deny permissions in filters
- Use WMI filtering for computers in different OUs

You can use Group Policy filtering if you need to deny or allow specific users or computers to apply GPOs. Group Policy filtering can be useful in folder redirection and software installation when you want to apply a GPO to a large set of domain resources, while also needing to further refine the domain resources to whom the policy applies. When the existing OU structure does not provide adequate separation or categorization, you also can use Group Policy filtering to control a GPO's application scope.

Consider the following guidelines when designing Group Policy filtering:

- Use filters on GPOs that are linked high in the domain structure. Generally, security group and WMI filters are most effective when you use them within a GPO that applies either to the entire domain, or to a large portion of the domain. These GPOs affect most (or sometimes all) users and computers. By filtering at this level, you are less likely to inadvertently omit users or computers that meet the group or WMI query search criteria, but are not within the OU tree to which the GPO applies.

- Use security group filtering to select a specific group of users or computers that are in different OUs. A user or computer must have Read and Apply Group Policy permissions for the GPO if you want it to apply. By default, GPOs allow both of these permissions to the Authenticated Users group. By using security groups, you can refine the groups of computers and users to which to apply the GPO.

- Avoid using Deny permissions for managing Group Policy. When you deny access, the denied permissions take precedence over other allowed permissions.

- Use WMI filters to select a specific group of computers. To determine the scope of GPOs based on the target computer's attributes, you can use WMI filters. A WMI filter consists of one or more WMI Query Language queries that run against the target computer's WMI repository. This is especially beneficial if you want to apply policies based on hardware—for example, when targeting a specific laptop model, or computers with enough space on their drive C. When you cannot use WMI filters, you should use security group filtering.

Understanding Slow Link Detection

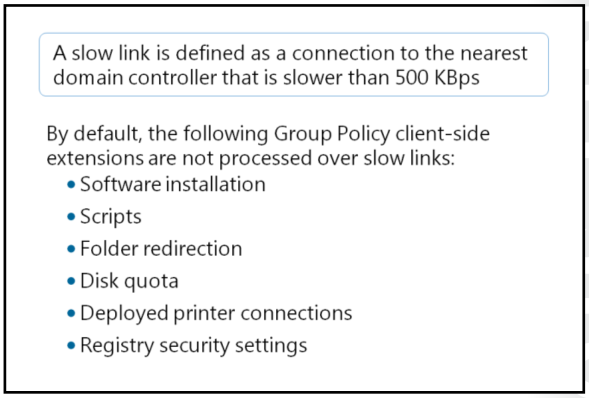

A slow link is defined as a connection to the nearest domain controller that is slower than 500 KBps

By default, the following Group Policy client-side extensions are not processed over slow links:

- Software installation
- Scripts
- Folder redirection
- Disk quota
- Deployed printer connections
- Registry security settings

In Windows Server 2012, Group Policy determines the link speed by using the network location awareness service. This service samples the current Transmission Control Protocol (TCP) traffic between the client and the domain controller.

You can use a Group Policy setting to define a slow link's parameter when you are applying GPOs. If you do not configure this setting, the default value defines a rate slower than 500 kilobytes per second (KBps) as a slow link.

You can partially control the Group Policy extensions that Group Policy can process over a slow link. By default, when you process Group Policy extensions over a slow link, Group Policy does not process all components. Slow links are available for the following settings:

- Computer Configuration. Use this to specify settings for Group Policy slow link detection for computers. To define this setting, access the Configure Group Policy slow link detection policy setting, which is in the Group Policy node underneath Computer Configuration\Administrative Templates\System. The unit of measurement for the connection speed is KBps.

- User Configuration. Use this setting to configure the slow link policy for users. To define this setting, access the Configure Group Policy slow link detection policy setting in the Group Policy node under User Configuration\ Administrative Templates\System.

- User Profiles. Use this setting to configure a slow network connection timeout for user profiles. To define this setting, access the Wait for remote user profile policy setting in the User Profiles node under Computer Configuration/Administrative Templates/System. If you enable the Do not detect slow network connections policy setting, the slow network connection timeout for the user profiles policy setting is invalid.

If AD DS detects a slow link, the client will not process Group Policy client-side extensions that might require significant network bandwidth. By default, AD DS does not process the following Group Policy client-side extension over a slow link:

- Software installation
- Scripts
- Folder redirection
- Disk quotas
- Deployed printer connections
- Registry security settings

Additionally, the Group Policy client-side extension for security and administrative template settings are always applied, regardless of whether a slow link is detected. You cannot modify this behavior.

When using roaming profiles, enabling the Delete cached copies of roaming profiles policy setting deletes the local copy of the roaming profile. Therefore, when the computer detects a slow connection, there will be no local copy of the roaming profile available to load.

GPO Caching

Another feature in Window 8.1 and Windows Server 2012 that benefits environments with slow links is GPO caching. GPO caching allows computers to retrieve the latest version of a specific GPO, and then write it locally to a datastore (located at C:\Windows\System32\GroupPolicy\Datastore). Subsequently, instead of downloading the GPO over the network the Windows operating systems can use the local copy of the GPO. Note that GPO caching works in very specific situations and relies on specific Group Policy settings. This means that GPO caching is not useful across all GPOs on all networks.

Considerations for Designing Group Policy Processing

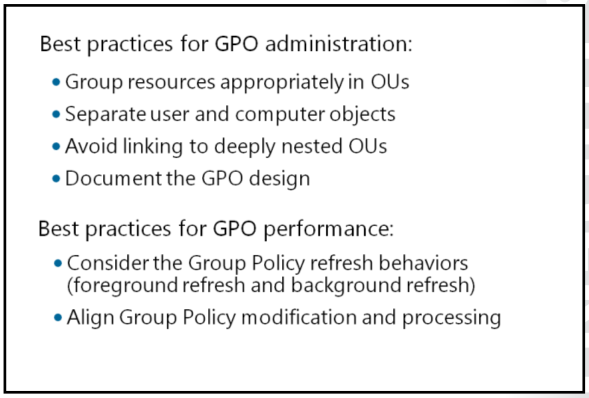

Best practices for GPO administration:

- Group resources appropriately in OUs
- Separate user and computer objects
- Avoid linking to deeply nested OUs
- Document the GPO design

Best practices for GPO performance:

- Consider the Group Policy refresh behaviors (foreground refresh and background refresh)
- Align Group Policy modification and processing

When designing the overall processing for your Group Policy environment, you should consider several factors. These factors generally fall into two categories: administration, and performance.

Administration

Generally, the number of GPOs in your Active Directory environment and how you create and name those GPOs affects how you administer their application and processing. When you are implementing GPOs, to ease the administrative effort that GPO processing requires, you should consider the following best practices:

- Group computers and users that have the same requirements, into the same OU.
- Separate the computer objects from the user objects by placing them in separate OUs.
- Avoid linking GPOs to deeply nested OU structures, because multiple GPOs at each level are difficult to manage and can increase the time it takes users to sign in.
- Document the GPO design. Be sure to include the Active Directory OU structure with the GPOs defined and linked to each container in your GPO documentation. You also should document each GPO’s settings. You can use the GPMC to generate HTML reports for all GPO settings.

Performance

Another important consideration when designing GPO processing is how client machines process GPOs. When designing your GPO structure, you should consider the following performance factors for your GPO environment:

- Group Policy refresh
- GPO modification and processing

Group Policy Refresh

GPOs are processed on client machines during two different refresh phases: foreground refresh, and background refresh:

- Foreground refresh. The foreground refresh occurs during computer startup for computer configuration settings, and at user logon for user configuration settings. During this refresh phase, Group Policy processes all Group Policy client-side extensions.
- Background refresh. The background refresh occurs after the initial foreground refresh, and at repeated intervals according the Group Policy refresh interval setting. (This setting is 90 minutes by default on domain members, and 5 minutes by default on domain controllers.) Some Group Policy client-side extensions, such as software installation and folder redirection, do not process during a background refresh.

During these refresh phases, the Group Policy client-side extensions on client machines are responsible for enacting operating system changes according to applied GPOs, and any changes that might have been made to the GPOs.

GPO Modification and Processing

When you modify the settings of a single GPO, often only the settings for the associated Group Policy client-side extensions are processed on client machines. However, if a policy has other policies that fall within its scope, the Group Policy client-side extensions for those policies are processed regardless of whether they have been modified. This behavior ensures that GPOs linked further down (or up) in the Active Directory structure do not affect or override the newly configured setting.

Because of this behavior, as a best practice you should avoid having multiple GPOs with a large number of settings. You also should avoid using performance-intensive Group Policy client-side extensions within the same scope, especially if you will be frequently changing any of the GPOs.

> **Question:** How would you want to redesign your Group Policy infrastructure based on the information from the last three lessons? What issues do you expect to encounter when implementing these changes?

Demonstration: Configuring GPO Inheritance and Filtering

Consider the following scenario for implementing GPOs. At A. Datum Corporation, the Chief Information Officer has recently announced the following policies:

- A security policy that prohibits the use of the Authenticated Users group
- A development policy that prohibits the use of corporate GPOs for quality assurance (QA) user and computer objects
- A right-sizing policy to ensure that GPOs apply only to intended user and computer objects

In this demonstration, you will see how to:

- Configure GPO inheritance.
- Configure a security filter.
- Configure a WMI filter.

Demonstration Steps

Configure GPO inheritance

1. Sign in to LON-DC1 and then open the Group Policy Management Console.
2. Create a new GPO named **Corp Settings**.
3. Under the **Development** OU, create an OU named **QA**.

4. Link the **Corp Settings** GPO to the **Development** OU.
5. Configure the **QA** OU to block inheritance.

Configure a security filter

- Modify the security filtering of the **Corp Settings** GPO by adding the **Development** group and removing the **Authenticated Users** group.

Configure a WMI filter

1. Create a new WMI filter named **Windows 7 clients or newer**. Use the following query for the filter:

```
SELECT Version, ProductType FROM Win32_OperatingSystem WHERE Version >= '6.1' AND
ProductType = '1'
```

2. Modify the **Corp Settings** GPO to use the **Windows 7 clients or newer** WMI filter.

Lesson 4
Planning Group Policy Management

Managing a Group Policy environment can be a challenging task, especially when several administrators are working together to administer a large AD DS environment. Ensuring that you protect your Group Policy environment from unwanted changes, and ensuring that the environment is resilient if a disaster occurs are important aspects of the Group Policy administration process. As such, you must implement, back up, and document your GPOs properly to ensure that you can recover your Group Policy environment in case of a failure.

Lesson Objectives

After completing this lesson, you will be able to:

- Design Group Policy backup and recovery.
- Migrate GPOs.
- Design Group Policy administration.
- Manage GPOs.

Considerations for Designing Backup and Recovery for Group Policy

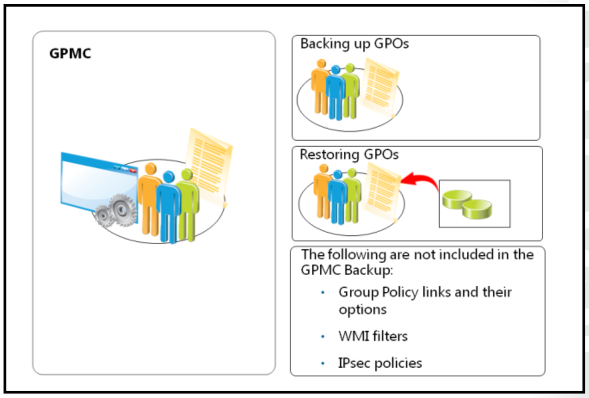

GPMC

Backing up GPOs

Restoring GPOs

The following are not included in the GPMC Backup:

- Group Policy links and their options
- WMI filters
- IPsec policies

You can use GPMC to back up and recover GPOs. Backing up a GPO saves all information that the GPO stores, to the file system. You can then use the GPO backup to restore the GPO to the backed-up state, or to restore the settings in the backup to another GPO.

GPMC includes the following features for Group Policy management:

- Backup. Backing up a GPO copies the data in the GPO into the file system. The backup function also serves as the export capability for GPOs.
- Restore. Restoring a GPO re-creates the GPO from the backup data. You can use the restore operation to recover a deleted GPO to its last backed-up version, and to roll back an active GPO to a known previous state.

In your recovery strategy, you need to consider the data that is stored outside of the GPO, because you cannot back up this data by using GPMC. The restore feature does not restore objects that are not part of the GPO, including links to a site, domain, OU, WMI filters, and Internet Protocol security (IPsec) policies.

Additionally, your design document should include the steps necessary to recover or re-create data should a disaster occur.

Windows Server 2012 provides a set of Windows PowerShell cmdlets that you can use to retrieve additional settings, which are helpful when restoring GPOs. A few examples of these cmdlets are listed in the following table.

Cmdlet	Description
Get-Command * -Module GroupPolicy	Lists all cmdlets that are available for Group Policy administration.
Get-GPOReport –ReportType HTML –All | Out-File <outputfile.html>	Creates an HTML file of all GPOs, which contains the following information: • Where is the GPO linked to? • Which WMI filters are used? • What are the GPO's security filters? • What are the GPO's settings?
Get-GPPermissions –Name <GPO> -All	Lists security filtering and other permissions for the specified GPO.
Invoke-GPUpdate –force	Reevaluates and reapplies all GPOs to the local system and currently logged-on user.

Considerations for Migrating GPOs

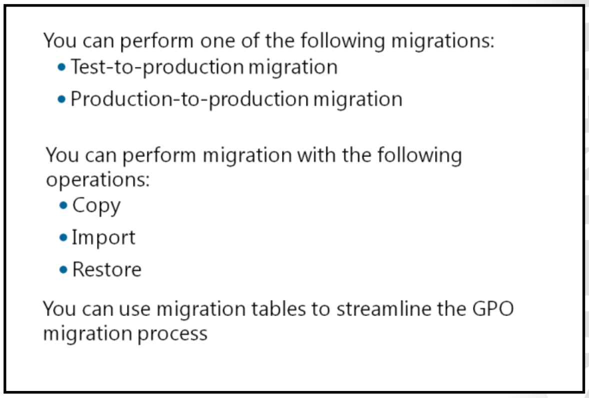

You can perform one of the following migrations:

- Test-to-production migration
- Production-to-production migration

You can perform migration with the following operations:

- Copy
- Import
- Restore

You can use migration tables to streamline the GPO migration process

Part of implementing a Group Policy design within a new domain environment involves migrating the applicable settings from the old domain environment. Group Policy migration involves moving or copying one or more GPOs from their original source domain to the new destination domain. In general, GPO migrations consist of two types:

- Test-to-production migration. These migrations usually involve a test domain structure that you establish for pre-implementation testing. After you complete testing and the test environment is ready for production, you can migrate GPOs from the test domain to the new production domain.
- Production-to-production migration. This migration scenario involves migrating GPOs from a former production domain to a new production domain.

To migrate GPOs, you can choose from the following three operations:

- Copy. A copy operation copies an existing, active GPO to the desired destination domain. This process always creates a new GPO. You can run the copy process from the GPMC in the original domain.
- Import. An import operation starts with a GPO backup to the file system, and then transfers the settings in that backup to an active GPO in the domain. Unlike the copy operation, the import does not create a new GPO. Instead, you must utilize a GPO that already exists.
- Restore. The restore operation also starts with a backup of one or more GPOs in the file system, but unlike the import process, this operation restores the backed-up GPOs to the newly created GPOs in the destination domain.

Migration Tables

Migration tables enable you to map your old environment's users and computer objects (which GPO settings may specify) to their counterparts in the new Active Directory environment. For example, the users in a security group called Adatum\Sales might now be in the new environment's Contoso\Sales group. If you are using this group for security group filtering in any of the imported GPOs, you will want them to use the new group instead.

You can edit migration tables by using the Migration Table Editor in the GPMC. To access the Migration Table Editor, right-click the Domains node or Group Policy Objects node in the GPMC, and then click Open Migration Table Editor.

Typically, you use the Migration Table Editor to create migration tables. You can also use the GPMC to populate relevant entries in your migration table from a set of GPO backups or from an existing GPO. To use the population feature, in the Migration Table Editor, click Tools, and then click either Populate from GPO or Populate from Backup. After selecting a GPO, the security principals and Universal Naming Convention (UNC) paths referenced in the selected GPO or backups are then extracted and entered into the migration table. After you create the initial set of entries in the table, you can update the Destination Name field to the appropriate values.

Considerations for Designing Group Policy Administration

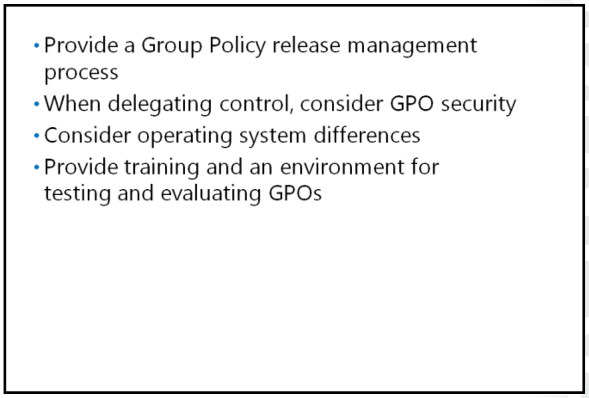

- Provide a Group Policy release management process
- When delegating control, consider GPO security
- Consider operating system differences
- Provide training and an environment for testing and evaluating GPOs

The Group Policy design calls for delegation of certain Group Policy administrative tasks. When assessing your organization's needs, you need to determine the degree to which you should centralize or distribute administrative control of Group Policy. To make these decisions, you need to consider the following factors:

- The release management process for Group Policy
- GPO security after you delegate control
- Differences between operating systems
- The GPO testing and evaluation environment

You should provide users to whom you delegate Group Policy administrative tasks with an environment for testing and evaluating GPOs. By doing this, the users can familiarize themselves with the Group Policy tools and tasks before using them in production.

When delegating GPO administration to other users, consider the following guidelines:

- Manage Windows Server 2012 policy settings by running the GPMC and the Group Policy Object Editor only from Windows Server 2012 or Windows 8 computers. To provide a consistent experience and to avoid making any accidental changes, we recommended that you use the most up-to-date Windows operating systems in your environment to manage GPOs. For example, if you use older Windows operating system versions to edit GPOs, certain settings that only apply to newer versions might not display.
- Consider the default permissions that allow certain predefined groups to create and manage GPOs, and the default permissions that determine who can perform these tasks.
- Use the GPMC to manage permissions at the task level. Grant additional permissions to the GPO's users.

- Provide training prior to delegating permissions in the production environment.

In general, you can delegate the following tasks for GPOs:

- Create and delete GPOs.
- Edit settings for specific GPOs.
- Link objects to existing OUs, and configure their options.

Advanced Group Policy Management

While designing for Group Policy administration, you can decide whether to use the Advanced Group Policy Management (AGPM) tool and its features. AGPM is a component of the Microsoft Desktop Optimization Pack (MDOP). MDOP is available for purchase for Microsoft Software Assurance customers (customers that maintain Software Assurance). AGPM provides extended functionality for Group Policy administrators, including:

- More precise administrative control of GPOs. AGPM provides the ability to create, review, and edit GPOs with zero impact. It also provides for a check-in and checkout system, which mitigates the risk of two different administrators modifying the same GPO simultaneously. In addition, workflow capabilities are built into AGPM, which allows higher-level administrators to approve GPO changes before deploying them to the production environment.
- Extended searching of GPOs. AGPM provides complex searching capabilities. You can find all GPOs that were changed by a specific administrator, or all GPOs that were changed on a specific date (such as the day before a production problem was reported). You can even search for GPOs that changed in a defined time period (for example, last week).
- Enhanced GPO importing and exporting. You can use AGPM to export a GPO from one forest and import it into another forest.
- Capability of rolling back a GPO to an earlier version. This feature enables you to roll back a GPO to a previous version if a problem arises with a new version. This rollback feature is made possible by the GPO archive of AGPM. The GPO archive is responsible for archiving older GPO versions as new versions are introduced.

Additional Reading: For additional AGPM features, visit
http://go.microsoft.com/fwlink/?LinkID=391885.

Demonstration: Managing GPOs

This demonstration shows you how to:

- Create a backup of all GPOs.
- Document GPO settings.

Demonstration Steps

Create a backup of all GPOs

1. Create a folder on drive C named **GPO-Backups**.
2. Open the GPMC, and back up all GPOs into the C:\GPO-Backups folder.

Note: You can also use the following Windows PowerShell cmdlet to back up the GPOs:

```
Backup-GPO –All –Path c:\GPO-Backups
```

3. View the backed-up GPOs.

Document GPO settings

1. In the GPMC, select a random GPO (for example, the **DSRM_Pwd** GPO from the previous demonstration).
2. Right-click, and save the GPO report.
3. In File Explorer, navigate to and open the **HTML-Report**.
4. View the Links and Security Filtering sections, view Delegation, and view the actual settings specified in the GPO.

Note: You can also use the following Windows PowerShell cmdlet to document GPO settings:

```
Get-GPOReport –Name GPO-Name –ReportType HTML –Path c:\GPOReports\GPOReport1.html
```

Lab: Designing and Implementing a Group Policy Object Strategy

Scenario

After completing the OU design, the next step in the AD DS design project for A. Datum Corporation is to complete a GPO design to manage user desktops and server security. The new GPO design will be a foundation that enables administrators to centrally manage desktop settings and user configurations. The GPO design should also enable A. Datum to meet its security requirements and to configure compliance settings.

Objectives

After completing this lab, you will be able to:

- Design a GPO strategy.
- Implement the GPO design.

Lab Setup

Estimated Time: 60 minutes

Virtual machines	20413C-LON-DC1 20413C-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, on the **Start** screen, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Microsoft Hyper-V® Manager, click **20413C-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on by using the following credentials:
   - User name: **Administrator**
   - Password: **Pa$$w0rd**
   - Domain: **Adatum**
5. Repeat steps 2 through 4 for 20413C-LON-CL1.

Exercise 1: Designing a Group Policy Object (GPO) Strategy

Scenario

Administrators at A. Datum Corporation are still considering whether to implement AD DS as a service model. In the meantime, you are tasked with planning a GPO model for the current infrastructure to manage user desktops and server security. You need to finalize the administrative tasks delegation model and determine the administrators who will have rights on clients.

A. Datum also wants to configure Windows Update settings, and restrict administrative tools for regular user accounts. Additionally, security requirements mandate that the company has a compliance warning related to misuse of corporate computers.

As the administrator of A. Datum, you are tasked with translating the business requirements into GPO settings. You must then design and implement the GPOs at the appropriate levels of the OU design.

In this exercise, you will design the GPO strategy for A. Datum that meets the business and organizational requirements.

Supporting Documentation

Brad Sutton

From:	Bill Malone [Bill@adatum.com]
Sent:	2 October 11:43
To:	Brad@adatum.com
Subject:	GPO Design

Hello Brad,

As we've discussed in our meeting yesterday, we need to strengthen the security of servers and configure the users' desktops according to the first initial design.

I've included the notes of our meeting in the attached proposal document. Please read through the document. Also, it would be great if you could send me the updated proposal document later this week.

Thank you very much,

Bill

A. Datum GPO Strategy Proposal

Document Reference Number: BS00918/1

Document Author	Brad Sutton
Date	2nd Oct

Requirements Overview

Design a GPO strategy that meets the following requirements:

- All of the organization's computers should have a core group of GPO settings that must be applied. These settings should include:
  - Configuring the local administrator accounts.
  - Configuring update settings.
  - Restricting certain options, such as access to the registry editor.

  These settings should not apply to administrator desktops.
- Each office should have a core group of settings that apply to their workstations. As of now, you need to implement the following:
  - Display a security warning prior to computer logon stating that only A. Datum employees can use the computers. This setting needs to be applied to each location, and to display automatically in other languages for foreign locations.
- All users must have a default set of mapped drives assigned to them. You should base the mapped drive on the department membership.
- The central IT administrators in London must be able to manage all GPOs and settings in the organization. Administrators in each office should be able to manage only GPOs that apply to that office.

A. Datum GPO Strategy Proposal

Summary of Information

The supporting OU structure includes the following:

- Users are currently grouped by department in a top-level OU.
- Clients are in the top-level Clients OU, which is separated by location on the next level.

Proposals

1. Which of the requirements will necessitate creating one or more GPOs?

2. Can you fulfill any of the requirements without creating GPOs?

3. Are there any exceptions to the default GPO application that you must consider?

4. List the GPOs that you must create to fulfill the lab scenario’s requirements. Provide the following information in the table provided:
   - Name of the GPO
   - The requirements that the GPO fulfills
   - The configuration settings (user policies, computer policies, user preferences, or computer preferences) that the GPO will contain
   - The container (domain, OU, site) to which the GPO will be linked

Name	Requirements fulfilled	Configuration settings	Applies to

A. Datum GPO Strategy Proposal
5. List other configuration tasks that you must perform within the Group Policy Management Console to fulfill the scenario requirements.

The main tasks for this exercise are as follows:

1. Read the supporting documentation.

2. Update the proposal document with your planned course of action.

3. Examine the suggested proposals in the Lab Answer Key.

4. Discuss your proposed solution with the class, as guided by your instructor.

▶ Task 1: Read the supporting documentation

- Read the documentation provided.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the A. Datum GPO Strategy Proposal document.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones in the Lab Answer Key.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you should have created a GPO design that meets the GPO requirements of A. Datum Corporation.

Exercise 2: Implementing the GPO Design

Scenario

After designing the GPO strategy for A. Datum, you will now implement the GPO design.

The main tasks for this exercise are as follows:

1. Prepare the environment.

2. Create the required GPOs and link them to the required domain containers.

3. Configure filtering.

4. Test the design.

▶ Task 1: Prepare the environment

1. Switch to LON-DC1, and if necessary, sign is as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. Run the following Windows PowerShell script to create the required OUs:

```
New-ADOrganizationalUnit –name Clients –path “dc=adatum,dc=com”
New-ADOrganizationalUnit –name London
–path “ou=clients,dc=adatum,dc=com”
Get-ADObject –Filter {name –eq ‘LON-CL1’} | Move-ADObject –TargetPath
“ou=London,ou=Clients,dc=adatum,dc=com”
```

3. Run the following Windows PowerShell script to create the required shared folders:

```
New-Item c:\shares –ItemType Directory
New-Item c:\shares\Marketing –ItemType Directory
New-SmbShare –Name Marketing –Path c:\shares\Marketing –FullAccess ADatum\Marketing
```

▶ Task 2: Create the required GPOs and link them to the required domain containers

1. On LON-DC1, in Server Manager, start the Group Policy Management Console.
2. Create the **All_Clients** policy. Configure the following settings:
   a. Navigate to **Computer Configuration\Policies\Windows Settings\Security Settings\Restricted Groups**.
   b. Create the new **Restricted Groups** setting for the local **Administrators**, and add the **IT** group.
   c. Navigate to **Computer Configuration\Policies\Administrative Templates\Windows Components\Windows Update**.
   d. Open the **Configure Automatic Updates** Group Policy setting, **Enable** the setting, and in the **Configure automatic updating** drop-down list box, click **4 – Auto download and schedule the install**.
3. Link the policy **All_Clients** to the **Clients** OU.
4. Create the **All_Users_but_Admins** policy. Configure the following setting:
   a. Navigate to **User Configuration\Policies\Administrative Templates\System**.
   b. Select the **Prevent access to registry editing tools** policy, and then enable the Group Policy setting.
5. Link the policy to the **adatum** domain.
6. Create the **London_Clients** policy. Configure the following setting:
   a. Navigate to **Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\Security Options**.
   b. Open the **Interactive Logon: Message text for users attempting to log on** policy setting. Select the **Define this policy setting in the template** check box, and then type the message **Only A. Datum Employees are allowed to log on to this computer**.
   c. Open the **Interactive Logon: Message title for users attempting to log on** policy setting. Select the **Define this policy setting** check box, and then type the message title **Property of A. Datum**.
7. Link the policy to **ou=London,ou=Clients**.

8. Create the policy **Marketing_Share**, and then configure it as follows:
   a. Navigate to User Configuration\Preferences\Windows Settings\Drive Maps.
   b. Location: \\LON-DC1\Marketing
   c. Label: Marketing-Materials
   d. Drive Letter: M
9. Link the policy to the **Marketing** OU.

▶ Task 3: Configure filtering

1. On LON-DC1, in the Group Policy Management Console, under the **Group Policy Objects** node, click the **All_Users_but_Admins** Group Policy.
2. On the **Delegation** tab, under **Advanced**, add the **IT** group, and deny the **Apply group policy** permissions to the group.

▶ Task 4: Test the design

1. Switch to LON-CL1, and sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. At a command prompt, run **gpupdate /force** to ensure that all GPOs are updated, and then restart LON-CL1.
3. On LON-CL1, sign in as **Adatum\Adam** with the password **Pa$$w0rd**.

Note: Adam Barr is a member of the Marketing Group.

4. Verify that the following settings apply to the client and to Adam:
   - Prior to sign in, the user receives a compliance warning.
   - Windows Update is configured to download and install updates.
   - Registry editing tools are prohibited.
   - The Marketing Share is mapped to drive letter M.
5. Sign out from LON-CL1.
6. Sign back in to LON-CL1 as **Adatum\Brad** with the password **Pa$$w0rd**.

Note: Brad Sutton is a member of the IT Group.

7. Verify that the following settings apply to the client and to Brad:
   - Prior to sign in, the user is receiving a compliance warning.
   - Windows Update is configured to download and install updates.
   - Registry editing tools are allowed.
   - The Marketing Share is not mapped.
   - Brad is a member of the local Administrators group.
8. Sign out from LON-CL1.

Results: After completing this exercise, you should have successfully implemented the GPO design that you created.

▶ Task: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20413C-LON-CL1**.

Module Review and Takeaways

Review Question(s)

Question: What are the options for applying a GPO to specific users or computers?

Question: What do you need to consider when applying a GPO to a site?

Best Practices

- Enable the Central Store for Group Policy Administrative Templates if you have multiple administrators who are editing GPOs, and if you are editing GPOs from different computers.
- Avoid using site-linked GPOs.
- Carefully plan your Group Policy backup and recovery strategy.
- Plan for Group Policy testing before you apply GPOs to production users and computers.
- Limit the number of GPOs that apply to users and computers. Use high-level GPOs for common settings, and try to limit individual settings in individual GPOs. A high number of GPOs increases startup and logon times.
- Take time on a regular basis to document or update your GPOs, their settings, and where they are linked in the OU structure.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
A recently changed policy does not yet apply.	
Security filtering does not work as expected.	

Module 9

Designing and Implementing an AD DS Physical Topology

Contents:

Module Overview

You should design the site topology for your network after you design the logical structure of your organization's Active Directory® infrastructure. The site topology is a logical representation of the physical network. The Windows Server® 2012 operating system uses site information for many purposes, including routing replication, client affinity, SYSVOL replication, Distributed File System (DFS) namespaces, and service locations.

When you design your infrastructure for Active Directory Domain Services (AD DS), you also should create a detailed design for domain controller placement, deployment, and high availability. Domain controllers are core components for every Active Directory environment. When designing domain controllers, you should be aware of the available deployment options and understand how to ensure that domain controllers are always available for user authentication and authorization.

In this module, you will learn how to design a distributed AD DS forest that supports domain controllers in your network that are separated by expensive, slow, or unreliable links. You will also learn about designing domain controller deployment, placement, and high availability.

Objectives

After completing this module, you will be able to:

- Design and implement Active Directory sites.
- Design and configure Active Directory replication.
- Design domain controller placement.
- Plan for virtualization of the domain controller role.
- Design domain controller deployments for high availability.

Lesson 1
Designing and Implementing Active Directory Sites

When designing an Active Directory deployment, you need to ensure that user authentication is efficient, and that you optimize replication between domain controllers. AD DS provides efficiency by using the concept of sites to map the Active Directory design to the physical network.

To design Active Directory sites, you first need to collect information about the existing network and locations. Then, you need to determine whether to implement a single site or multiple sites. As you design the Active Directory sites, consider automatic site coverage so that you can determine how the domain controllers in one site can provide authentication and related services for another site that does not have a domain controller. In addition, you should also consider existing applications that are site-aware, such as Microsoft® Exchange Server 2007 or newer, and include these applications in your site design.

In this lesson, you will learn about designing Active Directory sites and about considerations that you should account for when designing these sites.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe Active Directory sites and their benefits.
- Describe the options for designing Active Directory sites.
- Explain how to collect information for an Active Directory site design.
- Explain how automatic site coverage works.
- Describe the considerations for designing Active Directory sites.
- Create site objects.

Benefits of Deploying Active Directory Sites

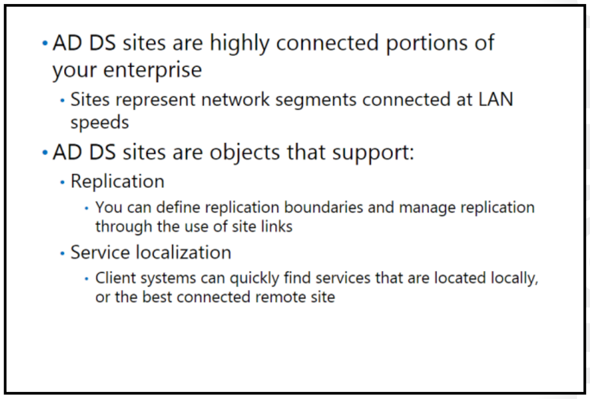

- AD DS sites are highly connected portions of your enterprise
  - Sites represent network segments connected at LAN speeds
- AD DS sites are objects that support:
  - Replication
    - You can define replication boundaries and manage replication through the use of site links
  - Service localization
    - Client systems can quickly find services that are located locally, or the best connected remote site

When administrators describe their network infrastructure, they often mention how many sites comprise their enterprise. They often equate a site with a physical location, such as an office or a city. These sites usually connect via network links. Together, the physical locations and the links between locations make up the physical network infrastructure.

AD DS logically represents the physical network infrastructure with objects called sites and site links. Although the words are similar, Active Directory sites and site links are not the same as the network sites and network connectivity. Network sites and network connections are physical components, and they usually represent real locations and links between them. However, Active Directory sites are logical components that ideally should be mapped closely to the physical network components.

To understand the difference, you must understand the properties and roles of sites in AD DS. *Active Directory sites* are objects in the directory, specifically in the Configuration container (CN=Configuration, DC=forest root domain), that help you manage replication traffic and facilitate service localization.

Replication Traffic

Replication is the transfer of changes between domain controllers. When you add a user account or change a user's password, a domain controller commits that change to the directory. The domain controller then replicates the change to all other domain controllers in the domain. The goal for replication is to bring all domain controllers to convergence, for example, to ensure that all the domain controllers have the same data and are consistent with each other. In larger organizations, not every domain controller may communicate directly with every other domain controller. In this case, a change replicates to replication partners until it replicates to all domain controllers. Some changes, such as changes to the configuration or schema container, replicate to all domain controllers in the forest.

AD DS is based on the assumption that your enterprise has two types of networks: highly connected, and less highly connected. Conceptually, a change made to AD DS should replicate immediately to the other domain controllers within a highly connected network. However, you might not want the change to replicate immediately over a slower, more expensive, or less reliable link. Instead, you might want to manage replication over less highly connected segments of your enterprise, so that you can optimize performance, reduce costs, or manage bandwidth.

An AD DS site represents a highly connected portion of your enterprise. The domain controllers within the site replicate changes almost instantly. AD DS site links represent the network connections between organization locations, and you can configure the properties on the AD DS site links to optimize replication across slower wide area networks (WANs).

Service Localization

AD DS is a distributed service. If you have at least two domain controllers, they provide the same authentication and directory access services. If you have more than one logical site, and if you place a domain controller in each site, client computers authenticate against the domain controller in their site. This is an example of service localization.

Active Directory sites help you localize services, including those services that domain controllers provide. During logon, Domain Name System (DNS) and AD DS automatically direct client computers to a domain controller in their site. If a domain controller is not available in their site, DNS and AD DS direct the client computers to a domain controller in the closest site that can authenticate the client computer.

In addition, some services require that you have a domain controller in a site where you deploy that service. For example, if you implement Exchange Server 2007 or newer, you should deploy a domain controller and global catalog in that same site. Examples of other site-aware applications are DFS and Microsoft System Center 2012 Configuration Manager.

Options for Designing AD DS Sites

AD DS in Windows Server 2012 has two site-design models: a single-site model, and a multiple-site model.

In general, choose the single-site model if your organization has one physical location and the domain controllers connect with high-speed links. Additionally, you can use the single-site model if you have multiple locations that are connected with high-speed links, such as fiber optics. If you use a single-site model, you are actually using a default configuration of Active Directory sites, because AD DS comes with one site that is created by default.

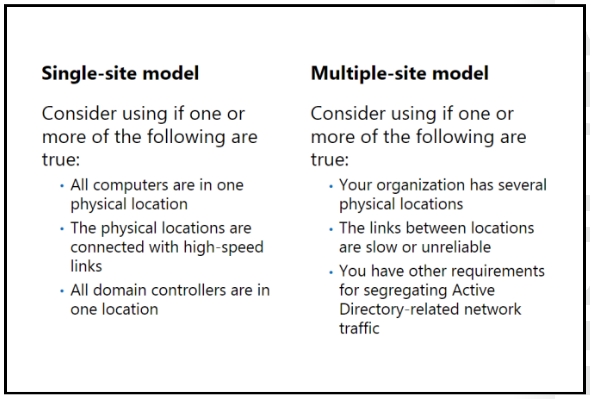

Single-site model

Consider using if one or more of the following are true:

- All computers are in one physical location
- The physical locations are connected with high-speed links
- All domain controllers are in one location

Multiple-site model

Consider using if one or more of the following are true:

- Your organization has several physical locations
- The links between locations are slow or unreliable
- You have other requirements for segregating Active Directory-related network traffic

Choose the multiple-site model if your organization's locations connect with WAN links and you want to manage the network traffic used by AD DS across those WAN links. However, do not base your site design only on the number of physical locations and links between them, although these two factors are important. Active Directory sites might not always map one-to-one to your network sites. Consider the following scenarios:

- You have offices in two locations. The two offices connect via a WAN link with low latency and highly available bandwidth. To simplify management, you could place a domain controller in each location but deploy only a single AD DS site.
- You have an enterprise on a large, highly connected campus. For replication purposes, the enterprise can be considered a single site. However, you want to encourage client computers to use services that are distributed to their location, so you configure multiple sites to support service localization. For example, you might deploy a separate site just for Exchange Server, or you might deploy a site so that you can segment client traffic to specific domain controllers.

Therefore, an AD DS site can include more than one network site, or it can be a subset of a single network site. The key is to remember that sites help with both replication management and service localization.

Several characteristics of your enterprise help you determine which type of sites you require. Because Active Directory sites manage Active Directory replication and service localization, it is not useful to create a site for a network location, unless the site hosts a domain controller or other Active Directory–aware service, such as DFS.

To ensure that client computers authenticate against a domain controller in their location, associate the IP subnet to which the client computers belong, with the Active Directory site. When you create an Active Directory site design, consider the cost of deploying a domain controller in a smaller branch office. It might not be cost-efficient to deploy a domain controller where the number of users is very low. Instead, those users could authenticate to a domain controller in a central location. You also must consider which option creates more traffic: logging on to a domain controller in another location, or replicating traffic between the domain controllers. Additionally, consider the work disruption if you do not deploy a domain controller in a certain location and the network connection fails.

Collecting Information for an AD DS Site Design

Because site design is linked closely to network infrastructure and physical computer deployments, the first step in creating an Active Directory site design is to document the existing network infrastructure. To begin, create a location map of your organization's physical network infrastructure. Typically, you get the information about the WAN topology from your organization's networking group. On the location map, identify the geographic locations where you have deployed computers, and identify the WAN connections that are deployed between all the locations.

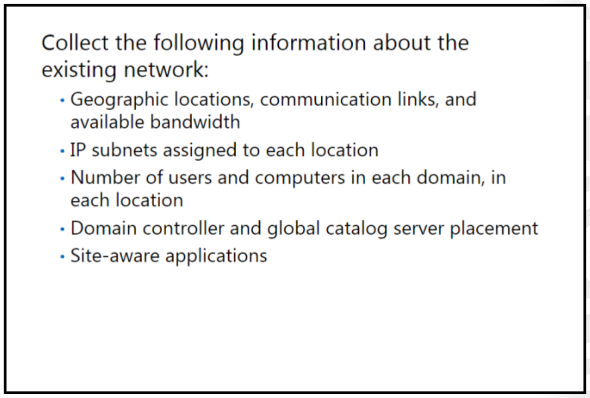

Collect the following information about the existing network:

- Geographic locations, communication links, and available bandwidth
- IP subnets assigned to each location
- Number of users and computers in each domain, in each location
- Domain controller and global catalog server placement
- Site-aware applications

As part of the location map, document the type of communication link, the link speed, and the available bandwidth between each location. You can obtain information about the available bandwidth from your network group, or you can analyze traffic on each link by using a protocol analyzer, such as Microsoft Message Analyzer. It is important that you get an approximate average usage for links between locations. For example, if you have a 1 megabit per second (Mbps) link between locations, which is sufficient for replication, but average usage of that link during peak periods is 80 percent of total capacity, you might have to schedule replication traffic for periods of lower link usage.

Next, record the IP subnets within each location. If you do not already know the subnet mask and IP address within each location, ask your organization’s networking group. Each site in AD DS must correspond to an IP subnet.

For each location, detail the number of users in each domain, and the number of workstations and servers deployed in the location. Document the types of network services that are deployed in the location.

Next, document the domain controllers' and the global catalog servers' placement in the office locations. If you are planning to deploy both a domain controller and a global catalog server, record this fact. If you are not planning to deploy a domain controller, identify the closest location where you plan to deploy one.

Consider applications that you deploy currently in sites, or which you plan to deploy, in case the applications are site-aware. An application is site-aware if it can recognize a site object and can use the resources that are deployed in the same site.

The existing network infrastructure influences your design. Use the following checklist when you collect information for the site design:

- The physical and logical network topology
- The geographic locations and the WAN links that connect them
- The type of WAN links
- The available bandwidth for each WAN link
- All existing site-aware applications
- The IP subnets that are assigned to each location
- The domain names of each domain in each location (if you deploy multiple domains)
- The number of users and computers in each domain, in each location

How Does Automatic Site Coverage Work?

If you deploy sites that do not have domain controllers, it is very important for you to understand automatic site coverage. If you do this, you must plan site links and site link costs so that you can optimize automatic coverage.

Not every site needs a domain controller. In cases where a site does not require a domain controller, you can use the automatic site-coverage functionality in Windows Server 2012. You can use this functionality to ensure that every site has a domain controller designated, even if the domain controller is physically located elsewhere.

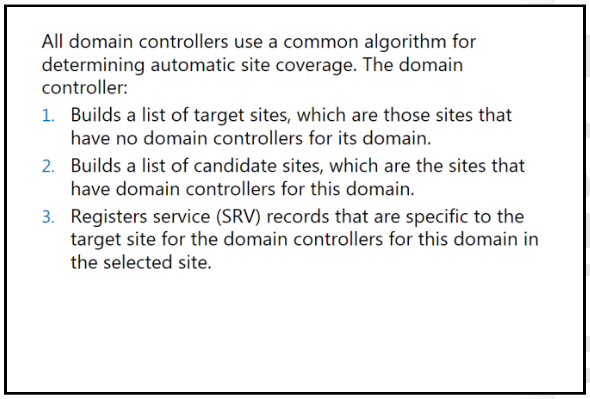

All domain controllers use a common algorithm for determining automatic site coverage. The domain controller:

1. Builds a list of target sites, which are those sites that have no domain controllers for its domain.
2. Builds a list of candidate sites, which are the sites that have domain controllers for this domain.
3. Registers service (SRV) records that are specific to the target site for the domain controllers for this domain in the selected site.

When you add a domain controller to the domain, the domain controller advertises its services by creating service (SRV) resource records (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. The domain controller advertises its ability to provide authentication and directory access by registering SRV records for the Kerberos version 5 (V5) authentication protocol and the Lightweight Directory Access Protocol

(LDAP). These SRV records are added to several folders within the DNS zones for the forest. The first folder, named _tcp, is within the domain zone and contains the SRV records for all domain controllers in the domain. The second folder is specific to the site in which the domain controller is located, with the path _sites*sitename*_tcp.

Automatic site coverage enables each domain controller to check all sites in the forest and then to check the replication cost matrix to determine the lowest-cost connection for providing services to a site. Site coverage is determined according to site link costs, and domain controllers register themselves in sites accordingly.

A domain controller registers a site-related SRV record in DNS in any site that does not have a domain controller for that domain, and for which the site has the lowest-cost connection. All domain controllers use a common algorithm for determining automatic site coverage.

That algorithm works as follows:

1. The domain controller builds a list of target sites, which are those sites that have no domain controllers for their domain (the domain of the current domain controller).
2. The domain controller builds a list of candidate sites, which are those sites that have domain controllers for their domain. For every target site:
   a. The domain controller builds a list of candidate sites to which this domain belongs. If the domain does not belong to any candidate site, do nothing.
   b. From the candidate sites, the domain controller builds a list of sites that have the lowest site link cost to the target site. If there are no sites, do nothing.
   c. If there is more than one site with the same lowest cost, the domain controller breaks the tie and reduces the list to one candidate site by selecting the site that has the largest number of domain controllers.
   d. If there is still more than one potential target site, the domain controller again breaks the tie by selecting the site that is first alphabetically.
3. The domain controller registers SRV records that are specific to the target site for the domain controllers for this domain in the selected site.

Considerations for Designing AD DS Sites

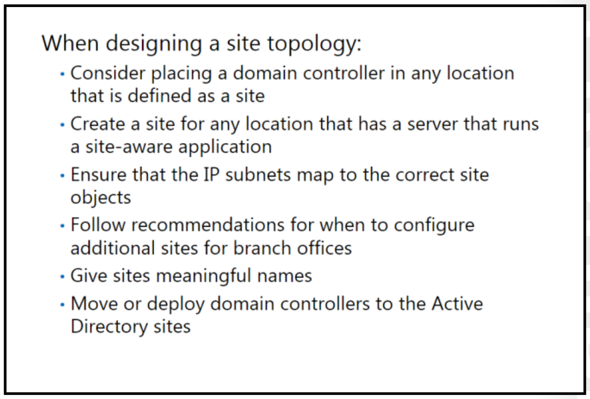

When designing a site topology:

- Consider placing a domain controller in any location that is defined as a site
- Create a site for any location that has a server that runs a site-aware application
- Ensure that the IP subnets map to the correct site objects
- Follow recommendations for when to configure additional sites for branch offices
- Give sites meaningful names
- Move or deploy domain controllers to the Active Directory sites

When designing Active Directory sites, you should follow these guidelines:

- Before starting to create the AD DS site design, always collect all necessary information about the physical network, the WAN links, and the number of users, computers, and domain controllers.
- Try to deploy at least one domain controller in a location where you establish an Active Directory site. Otherwise, you reduce the functionality of the site.
- Configure at least one domain controller per site to be the global catalog server. This speeds up the Active Directory logon and search processes for users.

- Consider site-aware applications.
- Always associate the appropriate IP subnet to each Active Directory site. This is crucial so that computers are aware to which site they belong.
- Create additional sites if:
    - A part of the network is separated by a slow WAN link.
    - A part of the network has enough users to warrant hosting domain controllers or other services in that location.
    - Directory query traffic warrants a local domain controller.
    - You want to control service localization.
    - You want to control replication between domain controllers.
- Give the sites meaningful names. In most cases, the site name should reflect the location for the site. You should rename the default site named Default-First-Site-Name to a more meaningful name.
- When you create Active Directory sites, move the existing domain controllers from Default-First-Site-Name to the appropriate sites. If you are deploying new domain controllers, you can associate the domain controller with the Active Directory site during the deployment.

Demonstration: Creating Site Objects

This demonstration shows how to:

- Create a new Active Directory site.
- Create a new Active Directory subnet.

Demonstration Steps

Create a new Active Directory site

1. From Server Manager, open **Active Directory Sites and Services**.
2. Right-click the **Sites** node, and then click **New Site**. Specify the name for the new site as **Paris**, and then associate the new site with the default site link.
3. Create additional sites, as needed.

Create a new Active Directory subnet

1. In the navigation pane, right-click **Subnets**, and then click **New Subnet**.
2. Provide the prefix, **10.10.0.0/16** and then associate the IP prefix to the **Paris** site object.

Lesson 2
Designing Active Directory Replication

After completing the site design, the next step is to create an Active Directory replication design. An effective replication design ensures that the Active Directory data replicates efficiently, even across slow and expensive WAN connections. Frequently, designing the optimal Active Directory replication configuration requires a balance between minimizing the use of WAN bandwidth, and ensuring that any Active Directory changes replicate as quickly as possible.

This lesson prepares you to create a replication design in a multiple-site Active Directory deployment.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the elements of intrasite replication.
- Describe the Knowledge Consistency Checker (KCC) and the Inter-Site Topology Checker in AD DS.
- Describe the options for designing intersite replication topologies.
- Choose the right replication protocol.
- Plan replication for global catalog and read-only domain controllers (RODCs).
- Plan SYSVOL replication.
- Design site links.
- Design bridgehead servers.
- Design site link bridging.
- Describe how to configure Active Directory replication.

Active Directory Replication Elements

Active Directory replication is one of the key processes that operate constantly in any Active Directory environment. Active Directory consists of multiple partitions, each of which replicates to all the domain controllers in the domain. Active Directory replication keeps each partition consistent across the domain controllers where they reside. Not all domain controllers have exactly the same information in their replicas at any one moment, because AD DS makes changes constantly to the directory. However, Active Directory replication ensures that all changes to a partition transfer to all replicas of that partition. Active Directory replication balances accuracy (or *integrity*) and consistency (called *convergence*) with performance (maintaining replication traffic at a reasonable level). Active Directory replication is designed to optimize the speed of replication.

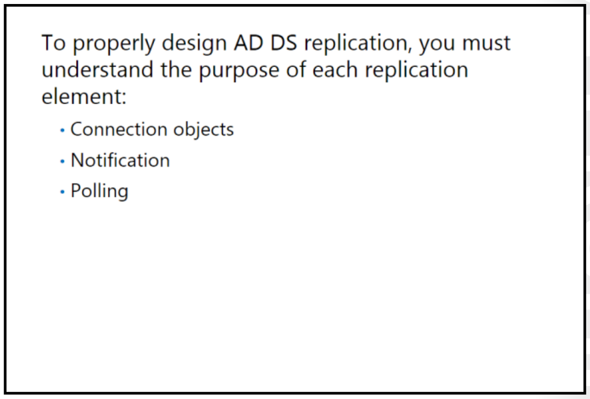

To properly design AD DS replication, you must understand the purpose of each replication element:

- Connection objects
- Notification
- Polling

Replication within a site and between sites is different when it comes to management. However, several components are common to both intersite and intrasite replication. To design Active Directory replication properly, you must understand the purpose of each replication component.

Connection Objects

If you want to manually enable one domain controller to replicate changes from another domain controller, you must create a connection object for this replication. A *connection object* represents the logical connection between two domain controllers that Active Directory replication is using.

Connection objects display in administrative tools in the Active Directory Sites and Services snap-in, in the NTDS Settings container of the server object for a domain controller.

Connection objects are one-way, representing inbound-only replication, because AD DS always uses pull replication. For example, if DC02 pulls changes from DC01, then DC02 is referred to as a downstream replication partner of DC01, and DC01 is the upstream partner. By default, connection objects exist in the NTDS Settings of both DC01 and DC02.

Note: You can use connection objects to force replication between two domain controllers by right-clicking the connection object, and then clicking Replicate Now. Remember, however, that replication is inbound-only. Therefore, to replicate both domain controllers, you must replicate the inbound connection object of each domain controller.

AD DS generates most of the connection objects automatically. On each domain controller, the KCC component of AD DS helps generate and optimize the replication automatically between domain controllers within a site. The KCC evaluates the domain controllers in a site, and then it creates connection objects to build the two-way, three-hop topology. If you add or remove a domain controller from the site, 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.

To specify replication paths that you want to persist, you can create connection objects manually. The KCC does not delete any connection objects that you create manually.

Within a site, very few scenarios require that you create a connection object. One such scenario is standby operations masters. You can select domain controllers to be standby operations masters, which AD DS uses when you transfer or seize the operations master role. A standby operations master should be a direct replication partner with the current operations master. Therefore, if a domain controller named DC01 is the relative ID (RID) master, and if DC02 is the system that will take over the RID master role if DC01 goes offline, you should create a connection object in DC02 so that it replicates directly from DC01.

Notification

Notification is the process by which an upstream partner informs its downstream partners that a change is available. For example, if the DC01 domain controller makes a change to a partition, it queues the change for replication to its partners. By default, DC01 waits 15 seconds to notify its first replication partner, DC02, of the change. By default, DC01 waits three seconds before it notifies additional partners. These delays, called the *initial notification delay* and the *subsequent notification delay*, stagger network traffic that intrasite replication causes.

Upon receiving the notification, the downstream partner, DC02, requests the changes from DC01, and the directory replication agent transfers the changed attribute from DC01 to DC02. In this example, DC01 made the initial change to AD DS, and therefore it is the originating domain controller. When DC02 receives the change from DC01, DC02 changes its directory. Then DC02 queues the change for replication to its own downstream partners.

If DC03 is a downstream replication partner of DC02, after 15 seconds DC02 notifies DC03 that it has a change. DC03 replicates the change to its directory, and then it notifies its downstream partners. The change has made two hops, from DC01 to DC02, and from DC02 to DC03. The replication topology ensures that all domain controllers in the site receive the change within three hops. At approximately 15 seconds per hop, the change fully replicates within the site within one minute.

Polling

DC01 might not make any changes to its replicas for quite a long time, particularly during off hours. As a result, DC02, its downstream replication partner, does not receive notifications from DC01. Alternatively, DC01 might be offline, which also would prevent it from sending notifications to DC02. Therefore, DC02 needs to know that its upstream partner is online and simply does not have any changes. This is achieved through *polling*.

In polling, the downstream replication partner queries the upstream replication partner as to whether it is queuing any changes for replication. By default, the polling interval for intrasite replication is once per hour. Although not recommended, you can configure a different polling frequency from the properties of a connection object by clicking Change Schedule.

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 site's replication topology rebuilds to accommodate the change.

What Are the KCC and the Intersite Topology Generator?

AD DS has two important mechanisms to manage both intrasite and intersite replication automatically. These are the KCC, and the intersite topology generator. Both services work automatically.

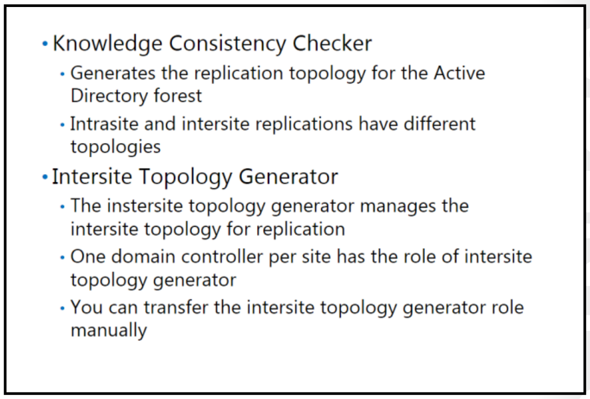

- Knowledge Consistency Checker
  - Generates the replication topology for the Active Directory forest
  - Intrasite and intersite replications have different topologies
- Intersite Topology Generator
  - The instersite topology generator manages the intersite topology for replication
  - One domain controller per site has the role of intersite topology generator
  - You can transfer the intersite topology generator role manually

KCC

Within a site, the connections between writable domain controllers are arranged in a bidirectional ring, with additional shortcut connections to reduce latency in large sites. However, the intersite topology is based on the spanning tree algorithm, which means that one intersite connection exists between any two sites for each directory partition, and the topology generally does not contain any shortcut connections.

Within a site, the KCC service in Windows Server automatically generates a topology for replication among the domain controllers in the domain by using a ring structure. The KCC is a built-in process that runs on all domain controllers. It analyzes the replication topology within a site every 15 minutes to ensure that the replication topology is working. If you add or remove a domain controller from the network or from a site, the KCC reconfigures the topology to reflect the change.

On each domain controller, the KCC creates replication routes by creating one-way, inbound connection objects that define connections from other domain controllers. For domain controllers in the same site, the KCC creates connection objects automatically without administrative intervention. If you have more than one site, you can configure site links between sites. Then, a single KCC in each site creates connections automatically between these sites as well.

Intersite Topology Generator

A fundamental concept in the generation of the topology within a site is that each server does its part to create a site-wide topology. Similarly, the generation of the topology between sites depends on each site doing its part to create a forest-wide topology between sites. As part of this effort, one domain controller per site assumes the role of the intersite topology generator. The KCC on this domain controller is responsible for creating connections between the domain controllers in its site and the domain controllers in other sites. This includes the inbound replication-connection objects for all domain controllers in the site in which the domain controller is located.

If the intersite topology generator assesses the topology and determines that its own site is the only site, it performs no further processing because no connections between sites are possible.

At 30-minute intervals, the current intersite topology generator notifies every other domain controller in the site of its existence. It does this by writing the **interSiteTopologyGenerator** attribute on the **NTDS Settings** object, in its domain controller object, which is located in the Configuration directory partition.

As the **interSiteTopologyGenerator** attribute replicates to other domain controllers, the KCC on each of these computers monitors this attribute to verify that it has been written. If 60 minutes elapse without a modification, a new intersite topology generator takes over.

It is possible for an administrator to modify a connection object on one domain controller and for the KCC to modify it on another domain controller before the initial change replicates. Overwriting such a change can occur within the local site or when a connection object changes in a remote site. By default, the KCC runs every 15 minutes. If the connection-object change does not replicate to another domain controller before the KCC on that domain controller runs, the KCC on that domain controller might modify the same connection object. In this case, ownership of the connection object belongs to the KCC because the write that is applied is the latest write to the connection object.

If you want to ensure that the KCC does not overwrite your modification to an intersite connection object, make the modification on the computer that has the intersite topology generator role in the site of the modified connection object.

Options for Designing Replication Topologies

In Windows Server 2012, you can use the following topologies for Active Directory intersite replication:

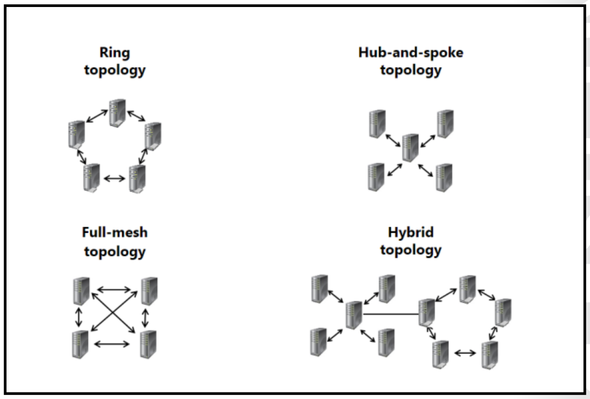

- Ring. In a ring topology, the files replicate from one server to another in a circular configuration, and each server connects to those on either side of it. Choose a ring topology if it resembles the layout of your physical network. For example, if each server is in a different site and has existing connectivity with neighboring servers, you can use the ring topology so that each server connects only to neighboring servers.
- Hub and spoke. In a hub-and-spoke topology, you designate one site as the hub. Other sites, called spokes, connect to the hub. You typically use this topology for WANs that consist of faster network connections between major computing hubs and slower links connecting branch offices. In this topology, AD DS replicates from the hub servers to the spoke servers and vice versa, but files do not replicate directly between two spoke servers. If you choose this topology, you must choose which site you want to be the hub. If you want to set up multiple hubs, use a custom topology. As a best practice, use multiple hubs to avoid a single point of failure.
- Full mesh. In this topology, every site connects to every other site. You typically use this topology if high-speed links are available between all sites. A change in AD DS on a domain controller in one site replicates directly to all bridgehead servers in all other sites, and from the bridgehead servers to all other domain controllers in each site. Because every site connects to every other site, the propagation

of change orders for replicating AD DS can impose a heavy burden on the network. You can reduce unnecessary traffic by using a different topology or by deleting the connections that you do not need.

- Hybrid. A hybrid topology is a combination of a hub and spoke and a full mesh topologies. One example of a hybrid topology is a redundant hub-and-spoke topology. In this configuration, a hub site might contain two domain controllers that connect by a high-speed link. Each of these two hub servers might connect with four branch domain controllers in a hub-and-spoke arrangement.

  In the case of the intrasite replication topology, replication by default works in a ring topology. Although you design an intersite replication topology based on your network infrastructure and routing capabilities, we do not recommend that you change the default behavior of the intrasite replication topology.

Considerations for Choosing a Replication Protocol

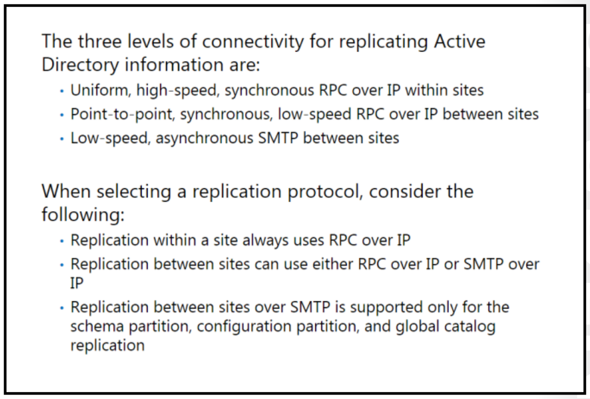

The three levels of connectivity for replicating Active Directory information are:

- Uniform, high-speed, synchronous RPC over IP within sites
- Point-to-point, synchronous, low-speed RPC over IP between sites
- Low-speed, asynchronous SMTP between sites

When selecting a replication protocol, consider the following:

- Replication within a site always uses RPC over IP
- Replication between sites can use either RPC over IP or SMTP over IP
- Replication between sites over SMTP is supported only for the schema partition, configuration partition, and global catalog replication

There are three levels of connectivity for replicating Active Directory information:

- Uniform, high-speed, synchronous remote procedure call (RPC) over IP within sites
- Point-to-point, synchronous, low-speed RPC over IP between sites
- Low-speed, asynchronous Simple Mail Transfer Protocol (SMTP) between sites

In Active Directory Sites and Services, site links are located in a container named IP, which itself is in the Inter-Site Transports container. Changes are replicated between domain controllers by using one of two protocols:

- Directory Service Remote Procedure Call (DS-RPC). DS-RPC displays in the Active Directory Sites and Services snap-in as IP, and you use it for all intrasite replication. It is the default and preferred protocol for intersite replication.
- Intersite messaging service (ISM) SMTP. Use ISM-SMTP, (also known simply as *SMTP*), only if network connections between sites are unreliable or are not always available.

Typically, you can assume you will use IP for all intersite replication. Very few organizations use SMTP for replication. This is because of the administrative overhead required to configure and manage a certification authority (CA), and because SMTP replication is not supported for the domain naming context. That is, if a site uses SMTP to replicate to the rest of the enterprise, that site must be its own domain.

The RPC intersite and intrasite transport (RPC over IP within sites and between sites) and the SMTP intersite transport (SMTP over IP between sites only) correspond to synchronous and asynchronous communication methods, respectively. Synchronous communication favors fast, available connections, while asynchronous communication is better for slow or intermittent connections.

IP Transport

The IP transport (RPC over IP) provides synchronous inbound replication. In the context of Active Directory replication, synchronous communication implies that after the destination domain controller sends the request for data, it waits for the source domain controller to receive the request, and then it constructs and sends the reply before it requests changes from any other domain controllers. In this way, inbound replication is sequential. Therefore, in synchronous transmission, the reply is received within a short time. Because of this, the IP transport is appropriate for linking sites in fully routed networks.

SMTP Transport

The SMTP transport (SMTP over IP) provides asynchronous replication, in which the destination domain controller does not wait for the reply and as a result may have multiple asynchronous requests outstanding at any particular time. Thus, in asynchronous transmission, the reply is not necessarily received within a short time. Because of this, asynchronous transport is appropriate for linking sites in networks that are not fully routed and that have particularly slow WAN links.

Considerations for Selecting a Replication Protocol

When choosing a replication protocol, consider the following:

- Replication within a site always uses RPC over IP protocol.
- Replication between sites can use either RPC over IP, or SMTP-over-IP protocol.
- Replication between sites over SMTP is supported for only the schema partition, Configuration partition, and global catalog replication. This means that domains can span sites only when point-to-point, synchronous RPC is available between sites. Replication between sites over SMTP also requires an enterprise CA, if you use it over site links. Domain controllers in the same domain must replicate by using the RPC over IP transport.

You can establish replication over the SMTP protocol only under specific conditions. If you have a scenario in which SMTP is the only option for replication between sites, you must fulfill several requirements. The KCC does not create connections that use SMTP unless you do the following:

- Install Internet Information Services (IIS) on both bridgehead servers.
- Install and configure an enterprise CA on your network. The CA provides the certificates and keys required to sign and encrypt SMTP messages that domain controllers exchange, which helps ensure the authenticity of directory updates. Specifically, a domain-controller certificate must be present on the replicating domain controllers. The replication request message is not encrypted, because that message contains no directory data. However, the replication reply message, which contains directory data, is encrypted by using a key length of 128 bits.
- Connect the sites by using SMTP site links. (Upcoming topics provide more detail about site links.)
- Ensure that the site link path between the sites has a lower cost than any IP/RPC site link that can reach the SMTP site.
- Do not attempt to replicate writable replicas of the same domain (although replication of the global catalog partial replicas is supported).
- Configure each domain controller to receive email.

You also must determine whether email routing is necessary. If the two replicating domain controllers have direct IP connectivity and can send email to each other, you do not have to perform any further configuration. However, if the two domain controllers must deliver email to each other through email gateways, then you must configure the domain controller to use the email gateway.

Planning Global Catalog and RODC Replication

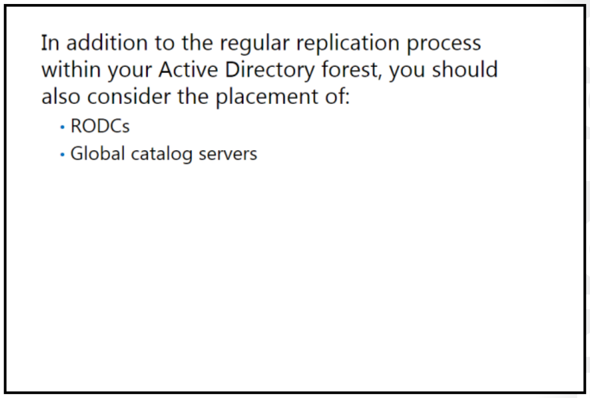

In addition to the regular replication process within your Active Directory forest, you should also consider the placement of:

- RODCs
- Global catalog servers

In addition to the regular replication process between domain controllers in the domains and the forest, you also should understand some specific types of replication when you create a replication design. The replication types include:

- RODC replication. An RODC must replicate domain data from a domain controller that is running Windows Server 2008 or newer. Therefore, in forests that contain Windows Server 2003 domain controllers, replication is among the most important considerations for determining where to place RODCs. The site topology and network constraints might affect the placement of an RODC and appropriate writable domain controllers. Typically, this scenario requires that you place a writable domain controller that is running at least Windows Server 2008 in the nearest site in the topology. The nearest site, in this case, is the site that has the lowest-cost site link for the site that contains the RODC.
- Global catalog replication. Global catalog servers are required for users to log on to domains. During the logon process, Windows contacts a global catalog server to obtain a list of universal groups to which the user belongs. If you place a domain controller in a site, but do not place a global catalog server in the same site, then the user authentication has to go outside of their local site for part of the logon process. However, if you place a global catalog server in each Active Directory site, user authentication can stay within the site for logon.

  You also can use global catalog servers for directory searches in AD DS. The global catalog servers contain a read-only copy of all objects in the forest. Additionally, applications such as Exchange Server create significantly more load on the global catalog servers. When you plan for global catalog server placement, ensure that you understand how to use the application, and that you are familiar with its Active Directory requirements.

  Global catalog replication helps ensure that users throughout the forest have fast access to information about every forest object. The default attributes that make up the global catalog provide a baseline of the most commonly searched attributes. These attributes replicate to the global catalog as part of normal Active Directory replication.

  The KCC generates the replication topology for the global catalog automatically. However, the global catalog replicates only to other domain controllers that you designate as global catalog servers.

Note: The attributes that you mark for inclusion in the global catalog will affect the replication of global catalog data.

You should place at least one global catalog server in each Active Directory site to prevent authentication traffic from traversing your WAN. For redundancy, consider placing two global catalog servers in each site.

Planning for SYSVOL Replication

Domain controllers use a shared folder named SYSVOL to store logon scripts and Group Policy Objects (GPOs). This folder then replicates to other domain controllers. Windows Server versions prior to Windows Server 2008 use the File Replication Service (FRS) to replicate SYSVOL. Windows Server 2012 uses DFS Replication service. DFS Replication can take place in a mixed domain with domain controllers built on Windows Server 2003 R2 or newer. However, a mixed domain could limit the domain mode that you can use. For example if you have a Windows 2003 R2 domain controller, you cannot raise the domain to a 2008 functional mode.

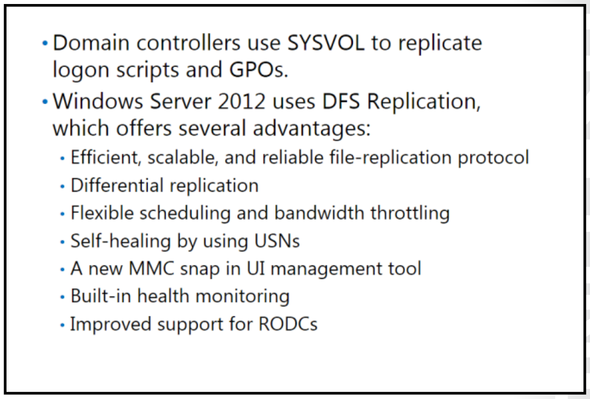

- Domain controllers use SYSVOL to replicate logon scripts and GPOs.
- Windows Server 2012 uses DFS Replication, which offers several advantages:
  - Efficient, scalable, and reliable file-replication protocol
  - Differential replication
  - Flexible scheduling and bandwidth throttling
  - Self-healing by using USNs
  - A new MMC snap in UI management tool
  - Built-in health monitoring
  - Improved support for RODCs

DFS Replication offers several advantages over FRS, including:

- An efficient, scalable, and reliable file-replication protocol that helps to ensure data consistency in multimaster replication scenarios.
- Differential replication of changes to files by using the remote differential compression (RDC) algorithm, which increases efficiency in branch-office scenarios.
- Flexible scheduling and bandwidth-throttling mechanisms.
- Self-healing from update sequence number (USN) journal wraps, and from database corruptions, resulting in minimal end-user intervention and monitoring requirements.
- The Microsoft Management Console (MMC) snap-in, which provides ease of administration.
- Built-in health monitoring tools for ease of monitoring deployments.
- Improved support for RODCs.

Considerations for Designing Site Links and Bridgehead Servers

The KCC assumes that within a site all domain controllers can communicate with each other. Between sites, however, you can represent the network paths over which replication should occur by creating site link objects. A site link object can contain two or more sites. The intersite topology generator builds connection objects between servers in each of the sites to enable intersite replication.

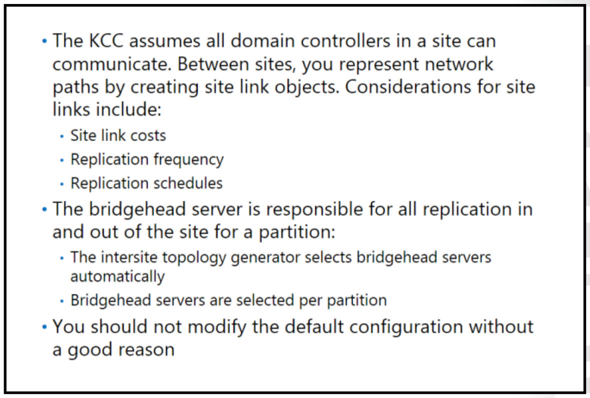

- The KCC assumes all domain controllers in a site can communicate. Between sites, you represent network paths by creating site link objects. Considerations for site links include:
  - Site link costs
  - Replication frequency
  - Replication schedules
- The bridgehead server is responsible for all replication in and out of the site for a partition:
  - The intersite topology generator selects bridgehead servers automatically
  - Bridgehead servers are selected per partition
- You should not modify the default configuration without a good reason

Understanding Site Links

A site link represents an available path for replication. A single site link does not control the network routes that replication uses, nor is it aware of routes at the network level. If you create a site link and add sites to it, you are in a way telling AD DS that it can replicate between any of the sites associated with the site link. If you create a site link and add two or more sites to it, you are presuming that there is a network route between these sites.

The intersite topology generator creates connection objects, and those objects determine the actual replication path. Although the replication topology that the intersite topology generator builds, effectively replicates AD DS it might not be efficient given your network topology. When you are designing sites, you should be fully aware of your routing topology and configure the site links accordingly.

When you create a forest, AD DS creates the default site link object DEFAULTIPSITELINK. By default, each new site that you add is associated with the DEFAULTIPSITELINK object. However, you should change this default behavior and create site links that reflect your physical network topology.

After you create site links, the intersite topology generator uses the topology to build an intersite replication topology that connects each site, and builds connection objects to configure the intersite replication paths. The intersite topology generator creates these connection objects automatically, and although you can create connection objects manually, there are few scenarios where you need (or would want) to do this.

When designing the site links, you also should plan the configuration of the site link attributes. These attributes can greatly influence replication over site links. For each site link, you can configure the site link cost, the replication frequency, and the replication schedule.

Site Link Costs

If the site link traffic has more than one route that it can take, you can use site link costs to manage the flow of replication traffic. You can configure the site link cost to indicate that a link is faster, more reliable, or preferred. Use higher costs for slow links and lower costs for fast links. AD DS replicates by using the connection with the lowest cost.

By default, AD DS configures all site links with a cost of 100. Site link costs have no specified unit; instead, you must establish metrics and define units for the value of the site link cost. For example, the metric for a site link cost usually is the actual link bandwidth. However, the site link cost also can be the real cost that you are paying for that physical link.

Replication Frequency

Intersite replication is based only on polling; there is no notification. Every three hours, by default, a replication partner polls its upstream replication partners to determine whether changes are available. If you want changes to the directory to replicate more quickly, you can change the polling interval for each site link.

Replication Schedules

By default, replication occurs 24 hours a day. However, you can restrict intersite replication to a specific time of day by changing the schedule attributes of a site link. Be careful when scheduling site link availability. If you schedule windows of availability that will not overlap, replication will not happen.

If you do not require link scheduling, select the **Ignore Schedules** option in the properties of the IP transport protocol. This option causes AD DS to ignore any schedules for site link availability, which helps ensure replication 24 hours a day over all site links.

When you design the site links, schedule the replication based on available bandwidth. For multiple site links, replication must be available during at least one common time period.

The intersite topology generator creates a replication topology between sites on a site link. However, having each domain controller in one site replicate to each domain controller in another site would be inefficient and would consume significant bandwidth. To make replication more efficient, the intersite topology generator selects one domain controller to be the main server that connects the other servers, so information passes from one server to another. This server is called the *bridgehead server*.

Understanding Bridgehead Servers

The bridgehead server is responsible for all replication into and out of the site for a partition. For example, if a data center site contains five domain controllers, one of the domain controllers is designated as the bridgehead server for domain-naming context. All changes to the domain partition within the data center replicate to all domain controllers in the site. When the changes reach the bridgehead server, those changes replicate to bridgehead servers in branch offices, which in turn replicate the changes to domain controllers in their sites. Similarly, any changes to the domain-naming context in branch offices replicate from the branches' bridgehead servers to the bridgehead server in the data center, which in turn replicates the changes to other domain controllers in the data center.

Selecting the Bridgehead Servers

The intersite topology generator selects bridgehead servers automatically, and then creates the intersite replication topology to ensure that changes replicate effectively between bridgehead servers that share a site link. Bridgehead servers are selected per partition, so it is possible that one domain controller in a site might be the bridgehead server for the schema, and another might be the bridgehead server for configuration. Typically, 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. When this is the case, bridgehead servers are chosen for those partitions as well.

You also can designate one or more preferred bridgehead servers by using the Active Directory Sites and Services snap-in. You can configure more than one preferred bridgehead server for a site, but the intersite topology generator selects and uses only one as the bridgehead server. If that bridgehead server fails, the intersite topology generator uses one of the other preferred bridgehead servers.

If you specify one or more bridgehead servers but none of them is available, the intersite topology generator does not select another server automatically, and replication does not occur for the site even if other servers could act as bridgehead servers. In an ideal situation, you would not configure preferred bridgehead servers. However, performance considerations might suggest that you assign the bridgehead server role to domain controllers that have more system resources. Firewall considerations might also require that you assign a single server to act as a bridgehead server, instead of allowing AD DS to select and possibly reassign bridgehead servers.

Windows Server 2008 R2 introduced load balancing, which distributes the workload evenly among bridgehead servers. In environments that are running earlier versions of Windows Server, inbound connections from sites might overload one domain controller in the hub site with requests. This could happen previously even when the connections to the hub site were load–balanced. The new bridgehead server selection technology improves the process of selecting a bridgehead server, and it avoids overloading a single server.

The intersite topology generator for the specified transport selects the preferred bridgehead servers automatically. In general, you should not modify the default configuration without having a good reason. However, if necessary you can disable the bridgehead server that the intersite topology generator selects, and then select your preferred bridgehead server manually. If you opt to select a bridgehead server manually:

- Ensure that you assign at least two or more bridgehead servers for any domain directory partition that meets the following conditions:
    - For any domain directory partition that has a replica in any other site
    - For any application directory partition that has a replica in another site
    - For the schema and configuration directory partitions, if no domains in the site have replicas in other sites
- Select the global catalog server as one of the preferred bridgehead servers, if the site has a global catalog server.

Considerations for Designing Site Link Bridging

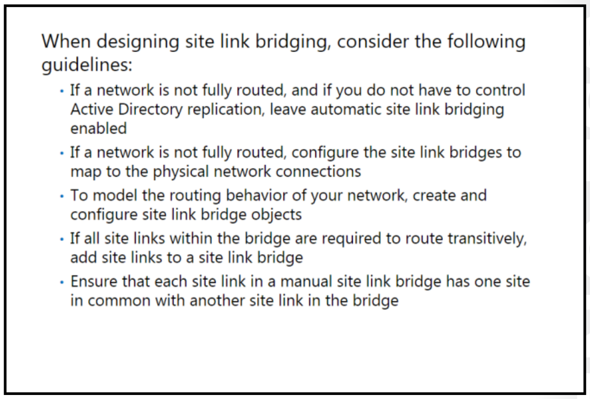

When designing site link bridging, consider the following guidelines:

- If a network is not fully routed, and if you do not have to control Active Directory replication, leave automatic site link bridging enabled
- If a network is not fully routed, configure the site link bridges to map to the physical network connections
- To model the routing behavior of your network, create and configure site link bridge objects
- If all site links within the bridge are required to route transitively, add site links to a site link bridge
- Ensure that each site link in a manual site link bridge has one site in common with another site link in the bridge

In many environments, particularly those with simpler network topologies, site links are sufficient to manage intersite replication. For more complex networks, however, you can configure additional elements and replication properties, such as automatic site link bridging and site link bridges.

Automatic Site Link Bridging

By default, site links are bridged. For example, if Site01 and Site02 are linked, and Site02 and Site03 are linked, then Site01 and Site03 are linked transitively. This means theoretically that the intersite topology generator can create a connection object directly between a domain controller in Site01 and a domain controller in Site03, and bypass the hub-and-spoke network topology.

You can disable automatic site link bridging by opening the properties of the IP transport in the Inter-Site Transports container, and clearing the Bridge All Site Links check box. You should do this only if your network is not fully routable, and transitivity does not apply at the network level.

Site Link Bridges

When a site link bridge connects two or more sites, it creates a transitive link.

You should create site link bridges only if you want to specify manually which sites are linked transitively at the network level, and which are not.

When designing site link bridging, consider the following guidelines:

- If the network is fully routed and if you do not need to control the flow of Active Directory replication, leave automatic site link bridging enabled for all site links by selecting the Bridge All Site Links option.
- If a network is not fully routed, clear the Bridge All Site Links option for the IP transport, and then configure the site link bridges to map to the physical network connections.
- If you need to model the routing behavior of your network, you can create and configure site link bridge objects.
- If all site links within the bridge are required to route transitively, you should add site links to a site link bridge.
- You must ensure that each site link in a manual site link bridge has one site in common with another site link in the bridge, so that the bridge can compute the cost from sites in one link to the sites in other links of the bridge.

Note: Site link bridges are necessary only if you clear the Bridge All Site Links option for the transport protocol. Remember that AD DS enables site link transitivity by default, in which case, site link bridges have no effect.

Demonstration: Configuring Active Directory Replication

In this demonstration you will see how to:

- Configure site links.
- Move a domain controller to a new site.

Demonstration Steps

Configure site links

1. On LON-DC1, in Server Manager, open Active Directory Sites and Services.
2. Create a new site link named **LONDON-PARIS**.
3. Add **AdatumHQ** and **PARIS** as **Sites in this site link**
4. Open the **LONDON-PARIS properties** dialog box. Configure the appropriate cost (**80**) and time interval (**60**) for the site link.
5. Configure the appropriate schedule for the site link.

Move a domain controller to a new site

1. In the navigation pane, expand **AdatumHQ** (or another site with a domain controller to move), and then expand the **Servers** folder.
2. Right-click **PAR-DC1**, and then click **Move**. Designate the appropriate site.
3. After moving the **PAR-DC1** domain controller, configure the preferred bridgehead server settings.

Lesson 3
Designing the Placement of Domain Controllers

When designing domain controllers, you must factor in where you place domain controllers, RODCs, global catalog servers, and operations master role holders within Active Directory sites. The placement design helps you determine the number of domain controllers and their hardware requirements for each domain in each site. You also must consider the impact on this design of Active Directory disaster recovery. In this lesson, you will learn about designing the placement of domain controllers.

Lesson Objectives

After completing this lesson, you will be able to:

- Plan hardware requirements for domain controllers.
- Describe considerations for deploying domain controllers on the Server Core installation.
- Describe considerations for planning domain controller locations.
- Describe considerations for planning global catalog server locations.
- Describe considerations for planning operations master server locations.
- Describe guidelines for monitoring Active Directory domain controllers.
- Describe the considerations for supporting branch offices.
- Explain the considerations for deploying domain controllers on Windows Azure™ virtual machines.

Planning Hardware Requirements for Domain Controllers

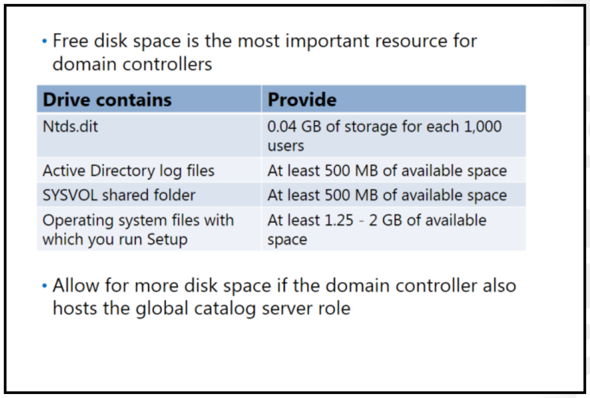

- Free disk space is the most important resource for domain controllers

Drive contains	Provide
Ntds.dit	0.04 GB of storage for each 1,000 users
Active Directory log files	At least 500 MB of available space
SYSVOL shared folder	At least 500 MB of available space
Operating system files with which you run Setup	At least 1.25 - 2 GB of available space

- Allow for more disk space if the domain controller also hosts the global catalog server role

Although AD DS is not a resource-demanding role, it is important that you plan hardware resources appropriately for servers that will be domain controllers.

Note: When considering virtualization of domain controllers, consider that most of the hardware requirements are the same for deploying domain controllers on to physical machines or in virtual machines.

Free disk space is the most important resource for domain controllers. However, you also should plan for the central processing unit (CPU), random access memory (RAM), and network adapter resources. In most cases, these requirements are the same as for installing Windows Server 2012.

At a minimum, a domain controller requires available free disk space for the Active Directory database, Active Directory log files, the SYSVOL shared folder, and the operating system.

Use the following guidelines to determine how much disk space to allocate for your AD DS installation:

- On the drive containing the Active Directory database (or the NT Directory Services directory information tree (ntds.dit) file), provide 0.04 gigabytes (GB) of storage for each 1,000 users. For example, if you have a forest with two domains (domain A and domain B), which have 10,000 and 5,000 users respectively, you should provide a minimum of 0.4 GB of disk space for each domain controller that hosts domain A. You also should provide a minimum of 0.2 GB of disk space for each domain controller that hosts domain B. Available space must equal at least 10 percent of your existing database size, or at least 250 megabytes (MB), whichever is greater.

Note: Additional space requirements also can depend on the size and number of the objects that are being recycled. For example, in a production Windows Server 2012 domain that is using the default **deletedObjectLifetime** and **recycledObjectLifetime** values of 180 days, the Active Directory Recycle Bin feature increases the Active Directory database size by 20 percent.

- On the drive containing the Active Directory log files, provide at least 500 MB of available space.
- On the drive containing the SYSVOL shared folder, provide at least 500 MB of available space.
- On the drive containing the operating system files with which installed Windows Server 2012, provide at least 1.25 to 2 GB of available space.
- Allow for more disk space if the domain controller also hosts the global catalog server role. This is especially important in multiple domain environments.

If you are upgrading existing domain controllers to Windows Server 2012, review and document the existing hardware configuration for each domain controller that you plan to upgrade. Use this information to identify the domain controllers in your environment that you can upgrade, and the domain controllers that do not meet the hardware requirements necessary to run Windows Server 2012.

Note: Remember that you can install Windows Server 2012 only on 64-bit hardware.

Note: When planning for domain controller deployment, avoid deploying software other than Active Directory–related software on domain controllers. This is for security and performance reasons. If you have a scenario that requires you to deploy additional software to the domain controller, your hardware requirements might change.

Considerations for Deploying Domain Controllers on Server Core

As a best practice, you should implement the maximum security available for servers that are domain controllers. This is because domain controllers typically store sensitive information, such as user passwords. Although the role-based configuration of Windows Server 2012 reduces a server's security surface by installing only the components and services that are required, you can further reduce its security surface by choosing a Server Core installation of Windows Server 2012.

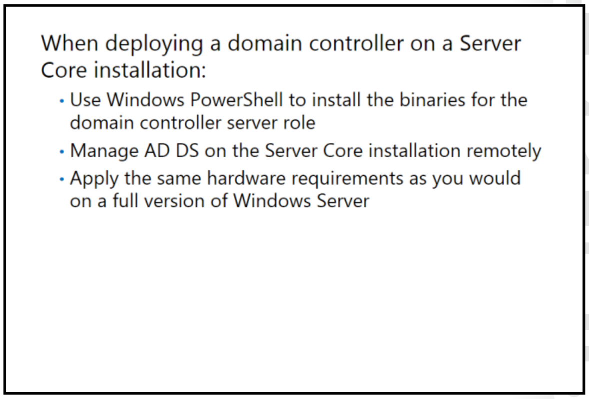

When deploying a domain controller on a Server Core installation:

- Use Windows PowerShell to install the binaries for the domain controller server role
- Manage AD DS on the Server Core installation remotely
- Apply the same hardware requirements as you would on a full version of Windows Server

Note: A Server Core installation provides a minimal installation of the Windows Server 2012 operating system. A Server Core installation does not install the Windows Explorer interface, which means that you administer a Server Core installation remotely by using GUI tools. To configure and manage a server locally, you must use command-line tools. Windows Server 2012 provides the Server Configuration tool (sconfig.cmd) that you can use to administer a Server Core installation.

Server Core installations of Windows Server 2012 consume only about 3 GB of disk space and approximately 256 MB of memory. A Server Core installation limits the server roles and features that you can add, but it does improve the server's security and manageability by reducing its attack surface. The number of services and components running at any one time is limited, so there are fewer opportunities for a hacker to compromise the server. Additionally, update management is much simpler on a Server Core installation, because there are a lower number of components to update.

The AD DS role can run on a Server Core installation. However, you cannot install or configure this role in a same way as you would on a full version of Windows Server. If you decide to implement a domain controller on a Server Core installation, you should be aware of following considerations:

- You must use Windows PowerShell® to install the binaries for the domain controller server role.
- You must install Active Directory tools on another server or workstation to manage AD DS remotely. Before you can manage the server core remotely, you must enable the remote management feature on the computer with the Server Core installation.
- The hardware requirements for the domain controller remain the same as when you are deploying it on Server Core installation.

Considerations for Planning Domain Controller Locations

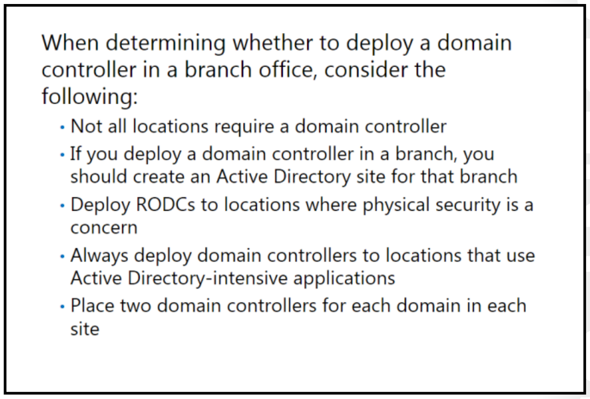

When determining whether to deploy a domain controller in a branch office, consider the following:

- Not all locations require a domain controller
- If you deploy a domain controller in a branch, you should create an Active Directory site for that branch
- Deploy RODCs to locations where physical security is a concern
- Always deploy domain controllers to locations that use Active Directory-intensive applications
- Place two domain controllers for each domain in each site

In organizations with just one location, placement of domain controllers is not as complicated of a design issue. You only need ensure that all clients and other servers have high-speed connections to one or more domain controllers. However, if your organization has more than one physical location, and if the locations are connected with a WAN link that is slower than 100 Mbps Ethernet, where you place your domain controllers is particularly important. This is because replication between domain controllers causes network traffic that requires bandwidth, which is typically limited between locations. Additionally, authentication traffic between clients and domain controllers also requires bandwidth.

When determining whether to deploy a domain controller in a branch office, consider the following:

- Not every physical location requires an onsite domain controller. Locations with very few users, or locations that connect with a high-speed WAN link to a location with a domain controller may still work well even without domain controllers on site. Therefore, place the domain controllers in organizational locations with significant numbers of users to minimize authentication traffic over the WAN link.
- If you deploy a domain controller in a branch location, you should create an Active Directory site for that location, associate it with an IP subnet, and then move the domain controller to that site. This is required if you want to direct branch-office clients and other Active Directory–aware services to the locally installed domain controller.
- Do not deploy an Active Directory domain controller in locations where you cannot guarantee physical security. Consider deploying an RODC instead.

- Always deploy domain controllers in locations where you have services or applications that intensively use AD DS, such as Exchange Server and Microsoft Lync® Server.
- Place two domain controllers for each domain in a site, and if users from multiple domains exist within that site, assign the global catalog server role to at least one of them.

Considerations for Planning Global Catalog Server Locations

The *global catalog* is a partition that stores commonly used information about every object in an Active Directory forest. In multiple domain environments, when a user in Domain B looks for an object in Domain A, the global catalog provides the query's results. To optimize efficiency, the global catalog does not contain every attribute of every forest object. Instead, it contains a subset of attributes that are useful for searching across domains. This is why the global catalog is also called the *partial attribute set*. In terms of its role in supporting searches, the global catalog is an index for the Active Directory database.

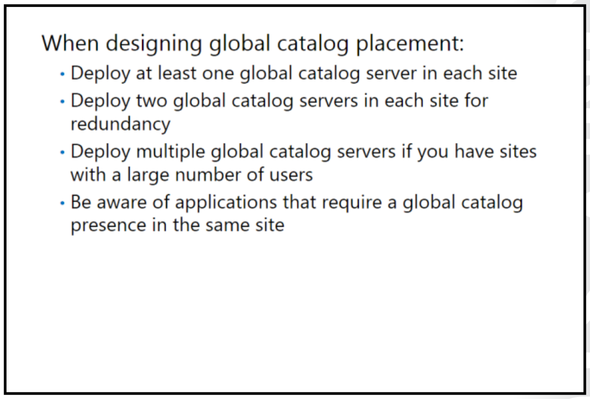

When designing global catalog placement:

- Deploy at least one global catalog server in each site
- Deploy two global catalog servers in each site for redundancy
- Deploy multiple global catalog servers if you have sites with a large number of users
- Be aware of applications that require a global catalog presence in the same site

In addition to the global catalog improving the efficiency of AD DS, it is required for applications such as Exchange Server and Microsoft Office Outlook®. Therefore, the global catalog must be made available for these and other applications. Only a domain controller can host the global catalog, but ideally, every domain controller should be a global catalog server.

If you choose to configure all domain controllers as global catalog servers, you no longer have to consider placement of the infrastructure operations master. This is because its role is no longer necessary in a domain where all domain controllers are global catalog servers.

The potential downside to this type of configuration is Active Directory replication, because the global catalog is a partition that you must replicate. In a single-domain forest, configuring all domain controllers as global catalog servers adds minimal overhead because the domain controllers already maintain a full set of attributes for all domain and forest objects. However, in a large, multiple domain Active Directory forest, overhead increases because of change replication to the partial attribute set of other domains' objects. Therefore, many organizations decide that Active Directory replication is efficient enough to replicate the global catalog without significant impact to their networks, and that the benefits of having domain controllers at remote locations far outweigh any impact to the network.

In general, when designing global catalog placement, you should follow these guidelines:

- Deploy at least one global catalog server in each Active Directory site. This is especially important if there is a slow or unreliable connection to the global catalog in another site.
- Deploy two global catalog servers in each Active Directory site for redundancy. If you have more than one domain controller per site, we recommend that all of them run the global catalog server role.
- If you have sites with a large number of users, deploy multiple global catalog servers to improve performance.
- Be aware of applications that require a global catalog presence in the same Active Directory site. Some examples of these applications are Exchange Server and DFS.

Considerations for Planning Operations Master Server Locations

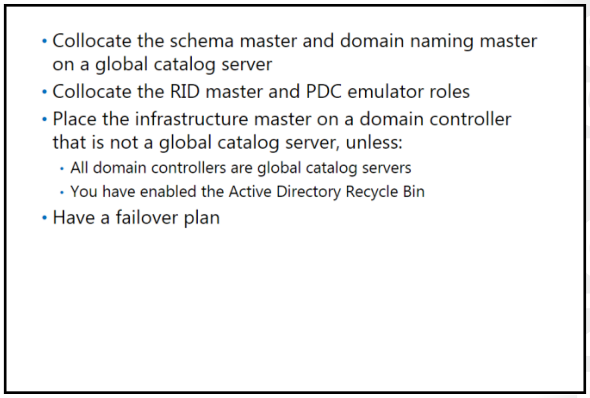

- Collocate the schema master and domain naming master on a global catalog server
- Collocate the RID master and PDC emulator roles
- Place the infrastructure master on a domain controller that is not a global catalog server, unless:
  - All domain controllers are global catalog servers
  - You have enabled the Active Directory Recycle Bin
- Have a failover plan

When you create the forest root domain with its first domain controller, the domain controller performs all five operations master roles:

- Schema operations master (schema master)
- Domain naming operations master (domain naming master)
- Infrastructure operations master (infrastructure master)
- RID operations master (RID master)
- PDC emulator operations master (pdc emulator master)

Together, these five roles are also known as operations master roles (or *flexible single master operations* or *FSMO)*.

As you add domain controllers to the domain, you can transfer the operations master role assignments to other domain controllers. This balances the load among domain controllers and optimizes placement of a single master operation. Additionally, spreading the operations master roles among multiple domain controllers reduces the possibility of losing all roles due to an outage.

The best practices for placing the operations master roles are as follows:

- Collocate the schema master and domain-naming master. Place the schema master and domain naming master roles on a single domain controller that is a global catalog server. Although these roles are used rarely, you should secure the domain controller that is hosting them. You must host the domain-naming master on a global catalog server, because when you add a new domain, the master must ensure that there is no object of any type with the same name as the new domain. The global catalog's partial replica contains the name of every object in the forest. The impact of these operations master roles on the network is light unless you are making schema modifications.
- Collocate the RID master role and PDC emulator role. Place the RID master and the PDC emulator roles on a single domain controller. If the load mandates that the roles are on two separate domain controllers, you should connect those two systems physically, and you should create explicit connection objects in AD DS so that they are direct replication partners. They should also be direct replication partners with domain controllers that you have selected as standby operations master servers.
- Place the infrastructure master on a domain controller that is not a global catalog server. You should place the infrastructure master on a domain controller that is not a global catalog server, but which you connect physically to a global catalog server. Ensure that the infrastructure master has explicit connection objects in AD DS to that global catalog server, so that they are direct replication partners. You can place the infrastructure master on the same domain controller that acts as the RID master and PDC emulator.

  If you use best practice methods and set all domain controllers in a domain as global catalog servers, then you do not have to figure out which domain controller is the infrastructure master. When all domain controllers are global catalog servers, all domain controllers have up-to-date information about every object in the forest, which eliminates the need for the infrastructure master role. Additionally, if you enable the Active Directory Recycle Bin, you do not need to worry about the placement of the infrastructure master.

- Have a failover plan. Prepare a plan for transferring operations master roles to other domain controllers in the event that one operations master server is offline. This includes a regular role transfer plan and a plan for seizing roles.

Guidelines for Monitoring Active Directory Domain Controllers

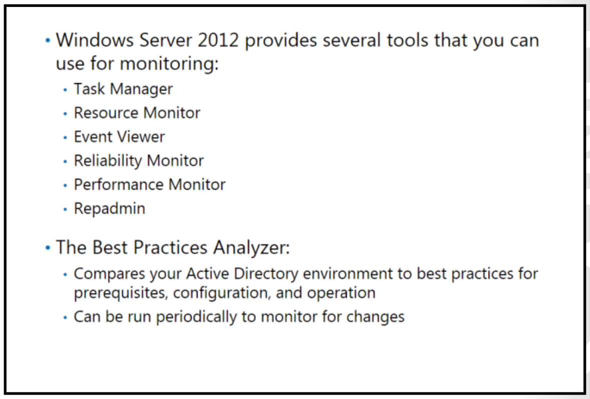

- Windows Server 2012 provides several tools that you can use for monitoring:
  - Task Manager
  - Resource Monitor
  - Event Viewer
  - Reliability Monitor
  - Performance Monitor
  - Repadmin
- The Best Practices Analyzer:
  - Compares your Active Directory environment to best practices for prerequisites, configuration, and operation
  - Can be run periodically to monitor for changes

Windows Server 2012 provides several tools that you can use for various types of monitoring. Most of these tools are available by default as Windows Server tools, and you can use them for real-time and historical monitoring of AD DS and other services. The Windows Server 2012 tools that administrators use most commonly are:

- Task Manager
- Resource Monitor
- Event Viewer
- Reliability Monitor
- Performance Monitor
- Repadmin

When monitoring AD DS, the most useful tools are Event Viewer and Performance Monitor. Event Viewer returns detailed information in various Active Directory–related logs, and Performance Monitor provides information about Active Directory performance.

Performance Monitor uses performance counters for many granular components, roles, services, and features. Windows Server components can register performance counters with Performance Monitor during installation. For example, when you add the Active Directory role to a server, the server registers the NT Directory Service performance object, which itself exposes dozens of counters that pertain to Active Directory performance.

With Performance Monitor, you can create a collection of counters interactively that monitor in real time. Alternately, you can save counters as part of a data collector set that you can reuse at any time to view real-time performance, or that you can launch in the future to log performance.

On a domain controller, you should monitor at least the following performance counters:

- NTDS\ DRA Inbound Bytes Total/sec
- NTDS\ DRA Inbound Object
- NTDS\ DRA Outbound Bytes Total/sec
- NTDS\ DRA Pending Replication Synchronizations
- NTDS \ Kerberos Authentications/sec
- NTDS\ NTLM Authentications

To establish an effective monitoring framework, follow these guidelines:

- Monitor early to establish baselines. Monitoring performance while a system is performing normally is critical. This enables you to establish a baseline so that you can document the ranges of the performance counters that you expect. To create a baseline, monitor performance counters during both busy and idle times.
- Monitor often to identify potential problems. Establish a regular monitoring routine, perhaps by using data collector sets that you schedule, and then compare the logs and reports with your baselines. Look for unusual deviations from values that you expect, to identify potential problems before those problems manifest in AD DS.
- Know how to monitor and interpret performance before a problem arises. When a problem arises, you must capture and interpret performance counters and events. Therefore, learn how to use the tools to determine which counters provide the most meaningful insight into the organization's performance. Establish data collector sets to help you capture performance during a performance-related problem, and then export those data collector sets to back them up in case the system that is down is the system that you normally use for monitoring.
- Capture appropriately. Do not perform excessive monitoring. If you attempt to monitor every counter, event trace, and registry entry, you will create such a huge amount of performance information that it will be difficult to process and thus identify the problem or problems. Additionally, monitoring performance degrades system performance. Align your monitoring activities with your business requirements for monitoring.

Best Practices Analyzer

Following best practices is critical, but sometimes organizations still do not implement them. Windows Server 2012 includes the Best Practices Analyzer (BPA), which is a server management tool that helps administrators implement best practices when configuring Active Directory services. BPA is available for the following server roles:

- AD DS
- Active Directory Certificate Services (AD CS)
- DNS Server
- Remote Desktop Services

BPA can help you implement best practices when you are configuring your Active Directory environment. Because AD DS is installed on your Windows Server 2012 server, BPA scans the Active Directory server role domain controllers, and it reports best practice violations.

You also can modify BPA reports by filtering or excluding results from BPA reports that you do not want to see. You can perform BPA tasks by using either the Server Manager GUI or cmdlets in the Windows PowerShell command-line interface.

Deploying Read-Only Domain Controllers in Branch Offices

In branch-office scenarios, the hub site is the central location for many IT services. In larger organizations, the hub site may include a robust data center. Branch offices, however, are often smaller sites in which no data center exists. Additionally, there may be no physically secure facility in which to house branch office servers, and there may be few (if any) local IT staff to support the servers.

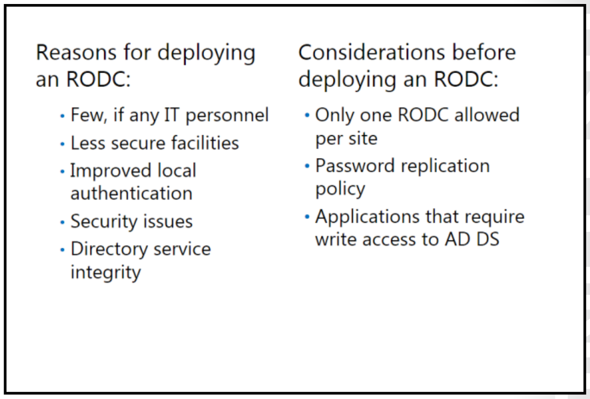

Reasons for deploying an RODC:

- Few, if any IT personnel
- Less secure facilities
- Improved local authentication
- Security issues
- Directory service integrity

Considerations before deploying an RODC:

- Only one RODC allowed per site
- Password replication policy
- Applications that require write access to AD DS

Reasons for Deploying RODCs

If you do not deploy a domain controller in a branch office, then you must direct authentication and service-ticket activities to the hub site over the WAN link. When a user first tries to access a specific service, the user's client requests a service ticket from the domain controller. Because users typically connect to multiple services during a workday, service-ticket activity is a regular occurrence. As a result, authentication and service-ticket activity over the WAN link between a branch office and a hub site can result in slow or unreliable performance.

If you place a domain controller in the branch office, authentication occurs more efficiently, but there are several potentially significant concerns:

- A domain controller maintains a copy of all attributes of all objects in its domain, including secure information such as user passwords. If a hacker accesses a domain controller, they could identify valid user names and passwords, at which point your entire domain is compromised. You then would have to reset the passwords for every user account in the domain. Because server security at branch offices is often less than ideal, a branch office domain controller poses a considerable security risk.
- Changes to the Active Directory database on a branch office domain controller replicate to the hub site and to the organization's other domain controllers. Therefore, corruption to the branch-office domain controller poses a risk to the integrity of the organization's AD DS. For example, a branch office administrator who performs a restore of the domain controller from an outdated backup could potentially cause significant issues for the entire domain.
- A branch-office domain controller may require maintenance, such as a new device driver. To perform maintenance on a standard domain controller, you must log on as a member of the Administrators group on the domain controller, which means effectively you are an administrator of the domain. It may not be appropriate to grant that level of capability to a branch-office support team.

Windows Server 2012 includes the RODC, which addresses potential branch-office issues. An RODC is a domain controller that you typically place in a branch office, and that maintains a copy of all objects and attributes in the domain except for secure information such as password-related properties. If you do not configure caching, the RODC receives logon requests from branch office users and then forwards them to a domain controller in the hub site for authentication.

You can configure a password replication policy for the RODC, which specifies the user accounts that the RODC caches. If the user who logs on is included in the password replication policy, then the RODC caches that user's credentials. Then the next time the user requests authentication, the RODC can perform the task locally. As users who are included in the password replication policy log-on, the RODC builds its cache of credentials so that it can perform authentication locally for those users.

Typically, you add users to the password replication policy who are located in the same physical site as the RODC. Because RODCs maintain only a subset of user credentials, security exposure is limited if a RODC is compromised. Only the user accounts that the RODC caches must have their passwords reset.

The RODC replicates changes to AD DS from the hub site's domain controllers. Replication is one way, and the RODC does not replicate changes to any other domain controller. This eliminates the exposure of the Active Directory services to corruption due to changes made to a compromised branch-office domain controller. Finally, RODCs have the equivalent of a local Administrators group. You can give one or more local support personnel the ability to maintain an RODC without granting them the equivalent Domain Admins rights.

RODC Limitations and Considerations

To reduce security risks and administrative costs, some domain controller options that are available for writable domain controllers are not available on an RODC. Before deciding to deploy a RODC, you must be aware of these limitations:

- RODCs cannot be operations master role holders. Operations master role holders must be able to write some information to the Active Directory database. Because of the read-only nature of the RODC's Active Directory database, it cannot act as an operations master role holder.
- RODCs cannot be bridgehead servers. Bridgehead servers specifically replicate changes from other sites. Because RODCs perform only inbound replication, they cannot act as a bridgehead server for a site.
- You can have only one RODC per site, per domain.
- Because Active Directory changes cannot be written directly to the RODC, no replication changes originate at the RODC. This means that any changes or corruption that a hacker might make at branch locations cannot replicate from the RODC to the forest. This also reduces the workload of the hub's bridgehead servers, and the effort required to monitor replication. RODC unidirectional replication applies to both AD DS and DFS replication.
- RODCs cannot properly support any application that needs to update Active Directory interactively, such as Exchange Server. If you are planning to deploy Exchange Server or similar applications on a site with RODC, you also should deploy a writable domain controller.
- You can install the DNS server service on RODCs. RODCs can replicate all application directory partitions that DNS uses, including ForestDNSZones and DomainDNSZones. If you install a DNS server on an RODC, clients can query it for name resolution just as they would query any other DNS server. Similar to the Active Directory database, the DNS server on an RODC is read-only, and therefore, it does not support client updates directly, though it can host other DNS zones that are writable.

Considerations for Deploying Domain Controllers on Windows Azure Virtual Machines

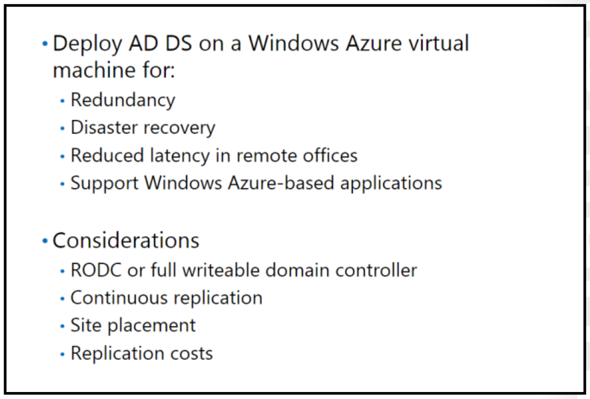

- Deploy AD DS on a Windows Azure virtual machine for:
  - Redundancy
  - Disaster recovery
  - Reduced latency in remote offices
  - Support Windows Azure-based applications
- Considerations
  - RODC or full writeable domain controller
  - Continuous replication
  - Site placement
  - Replication costs

Expanding AD DS onto Windows Azure virtual machines provides benefits that are not available when you use on-premises installations. By placing a domain controller on Windows Azure virtual machines, you can use two main benefits of cloud-based systems: high availability, and accessibility from anywhere.

Reasons for Placing a Domain Controller in Windows Azure

There are many reasons for placing a domain controller on Windows Azure virtual machines, including:

- Redundancy. You can put a domain controller in Windows Azure to authenticate users and act as a backup in case a single on-premises domain controller fails.
- Disaster recovery. If something should happen to the on-premises domain controllers, you can use a Windows Azure–based domain controller as a starting point for a disaster recovery. For example, if a large-scale disaster occurs, and most or your entire on-premises network is unavailable, your cloud-based infrastructure (including AD DS) is available as a starting point for your disaster recovery efforts. Disaster recovery does not include recovering inadvertent changes, such as removing or changing of attributes, or deleting an individual user.
- Decrease sign-in latency for remote offices. You may have remote offices without on-premises domain controllers for various reasons. A Windows Azure–based domain controller may be closer to the remote location and therefore help to reduce sign-in latency.
- Support for cloud-based applications that require AD DS. If you have deployed cloud-based applications that require authentication, you may want to deploy a domain controller in Windows Azure to provide the authentication.

Considerations for Placing a Domain Controller in Windows Azure

Installing a domain controller in Windows Azure is similar to installing a domain controller on an on-premises virtual machine. The primary difference is that you do not have control over the hardware maintenance. When deploying a domain controller in a virtualized environment, consider the following:

- An Active Directory domain controller replicates continuously. Ensure that all domain controllers are performing inbound replications. If inbound replication fails, the following event will be logged in the directory service log:
  - Event ID: 2042
  - Source: NTDS Replication
  - Type: Error
  - Description: It has been too long since this machine last replicated with the named source machine. The time between replications with this source has exceeded the tombstone lifetime. Replication has been stopped with this source.
- Do not pause a virtual domain controller for any long periods, or you can introduce errors due to replication being stopped for a long time.

When placing a domain controller in Windows Azure, you need to plan carefully the subnet and site relationships for the Windows Azure-based domain controllers. The decisions you make directly affect the overall cost of maintaining your Windows Azure environment. When planning your replication policies consider the following:

- Do you plan to place a writable domain controller or an RODC in Windows Azure? Replication to the Windows Azure domain controller does not incur additional charges, which means the RODC has a lower cost. However, an RODC might not be desirable, depending on the reasons for placing a domain controller in Windows Azure—for example, if you intend to use the Windows Azure–based domain controller for disaster recovery.
- How frequently do you schedule replication? Replication from the Windows Azure domain controller incurs additional costs. The more frequent the replication, the higher the cost. However, if you make the frequency too long, the domain controller may never synchronize fully with the on-premises domain controllers and data could be lost in a disaster recovery scenario. Additionally, if you allow the Windows Azure–based domain controller to exceed the tombstone lifetime, you may have to rebuild the Windows Azure-based domain controller completely.

- What replication policies do you use? When having up-to-date objects is not a necessity, you should create time-based replication policies, not event-driven replication policies. If you use event-driven replication policies, you may find that the costs for replication traffic vary greatly from month to month.

When determining the site configuration to use with a Windows Azure–based domain controller, you have two options: either add the Windows Azure subnet to an existing site, or create an additional site for the Windows Azure subnet. By creating a Windows Azure subnet, your Windows Azure systems use the Windows Azure–based domain controller, and other sites use the Windows Azure–based domain controller based on the subnet costs that you have defined.

Defining your link costs is important so that you use the Windows Azure–based domain controller under planned circumstances. For example, if you plan for remote sites to use your WAN connections for authentication and to use Windows Azure–based Active Directory only when the WAN links are down, the links to the Windows Azure–based site need to have a higher cost than the links for WAN-connected sites. If the link for the Windows Azure–based site has a low cost, you can use the Windows Azure–based site for authentication instead of the WAN–based site.

Lesson 4
Virtualization Considerations for Domain Controllers

Running a domain controller as a virtual machine is not the same as running other roles in the virtual environment. When designing domain controllers in virtual environments, you must consider several things. In this lesson, you will learn about designing and securing domain controllers that are running in virtual machines.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe considerations for virtualizing domain controllers.
- Explain how to secure virtualized domain controllers.
- Describe considerations for deploying domain controllers as virtual machines.
- Explain how to clone domain controllers.

Considerations for Virtualizing Domain Controllers

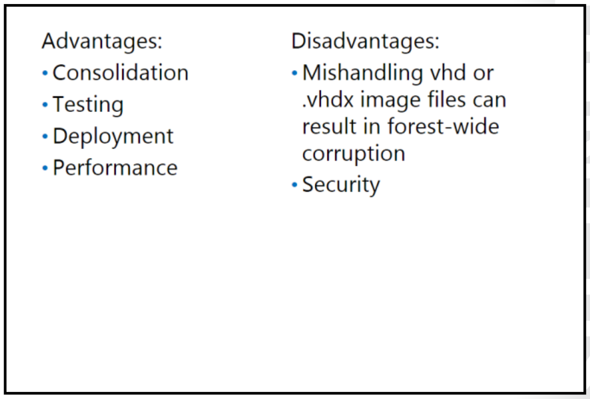

Advantages:
- Consolidation
- Testing
- Deployment
- Performance

Disadvantages:
- Mishandling vhd or .vhdx image files can result in forest-wide corruption
- Security

Before you virtualize a domain controller, you should carefully consider both the advantages and disadvantages of using virtual machines. When you use server virtualization, the differences between physical machines and virtual machines decrease. However, there are still important factors that you should consider, which can help you determine whether you should virtualize a domain controller. In addition to Windows Server 2012 Hyper-V®, you also can use other, non-Microsoft supported virtualization platforms for domain controller virtualization.

Advantages of Domain Controller Virtualization

Virtualized domain controllers offer the following advantages:

- Consolidation. With Hyper-V, you can consolidate multiple domain controllers with other application servers onto fewer servers. If you have multiple domain controllers in a hub site that serves a small domain with a presence in one or more other sites, you do not have to deploy multiple physical domain controllers to provide for bridgehead servers and redundancy support.
- Testing. Because testing environments rarely require high-performance domain controllers, you can set up a test configuration that represents your production environment using much less hardware than was necessary prior to virtualization. You can represent several domain controllers on a few physical computers with the same Active Directory configuration (including domains and sites) as your production domain controllers. You can test complex changes—such as schema changes or large-scale Active Directory site or replication changes—before you make these changes in your production environment.
- Deployment. The hardware environment that Hyper-V emulates remains unchanged throughout a virtual machine's use, regardless of the underlying native hardware to which you might transfer the

virtual machine. Therefore, you can vary server hardware in a preproduction environment to build and test the images even though the intended production server may be significantly different.

- Performance. If you configure your virtual environment correctly, you can expect a decrease in performance of only 2 – 12 percent on a single domain controller that runs in a virtual machine under heavy load. Compare this to one domain controller that runs on native hardware with the same specifications.

Disadvantages of Domain Controller Virtualization

There are risks associated with running domain controller virtual machines in production environments, which is why you should manage them carefully when you deploy them. When you consider whether to run domain controllers in virtual machines, remember that mishandling virtual hard disk files (.vhd or .vhdx) can result in forest-wide corruption. Mishandling a .vhd or .vhdx file can include starting an older copy of a .vhd or .vhdx file, or copying a .vhd or .vhdx file to a different network location, and then running multiple copies simultaneously.

Virtualizing domain controllers can result in the following issues:

- Potential data corruption or an improper restore process. In a production environment, deliberate or inadvertent misuse of .vhd or .vhdx files can result in data corruption that affects the entire forest. In most cases, the operating system detects improper restore conditions and stops replication. However, this does not guarantee data protection in all circumstances.
- Security and management of image files. Anyone who manages .vhd or .vhdx files or has any access to the files must be highly trusted in your organization, and must be a member of trusted security groups in the forest. Any user who can copy a .vhd or .vhdx file effectively owns the forest and its data. A hacker or unauthorized administrator could use a copied virtual machine file to compromise passwords or corrupt the forest. Additionally, unlike the theft of a physical computer, you are less likely to detect the copied theft of a .vhd or .vhdx file. Therefore, you should secure the .vhd or .vhdx files for the host operating system and the guest operating system with the same physical restrictions and software restrictions that you use to secure physical domain controllers.

Securing Virtualized Domain Controllers

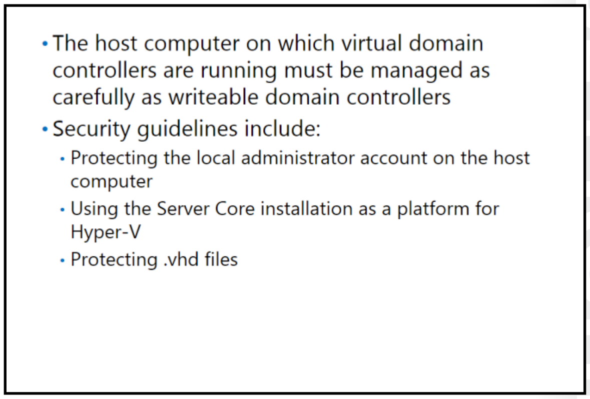

- The host computer on which virtual domain controllers are running must be managed as carefully as writeable domain controllers
- Security guidelines include:
  - Protecting the local administrator account on the host computer
  - Using the Server Core installation as a platform for Hyper-V
  - Protecting .vhd files

You must manage the host computer on which virtual domain controllers run as carefully as you manage a writeable domain controller, even if that computer is only a domain-joined or workgroup computer. This is an important security consideration because a mismanaged host is vulnerable to an elevation-of-privilege attack, which occurs when a malicious user gains access and system privileges that you did not authorize or assign legitimately. A malicious user can use this type of attack to compromise all the virtual machines, domains, and forests that this computer hosts.

When you are planning to virtualize domain controllers, keep the following security considerations in mind:

- The local administrator of a computer that hosts virtual, writeable domain controllers should have the equivalent credentials to the default domain administrator of all the domains and forests to which those domain controllers belong.

- To avoid security and performance issues, we recommend using a host that is running a Server Core installation of Windows Server 2012, with no applications installed other than Hyper-V. This configuration limits the number of applications and services that you can install on the server, which should result in increased performance, and fewer applications and services that a malicious user could use to attack the computer or network. If a separate management network exists, we also recommend that you connect the host only to the management network.

Security of .vhd or .vhdx Files

A .vhd or .vhdx file of a virtual domain controller is equivalent to the physical hard drive of a physical domain controller. As such, you should protect the .vhd file the same way that you would secure the hard drive of a physical domain controller. You should ensure that only reliable and trusted administrators can access the domain controller's .vhd or .vhdx files.

RODCs

Because RODCs are often used in locations where physical security cannot be guaranteed, you must take additional measures to mitigate the effects of an attack on an RODC. We recommend using BitLocker® Drive Encryption on the virtual host to protect .vhd or .vhdx files on the host in case of the physical disk's theft.

Considerations for Deploying Virtualized Domain Controllers

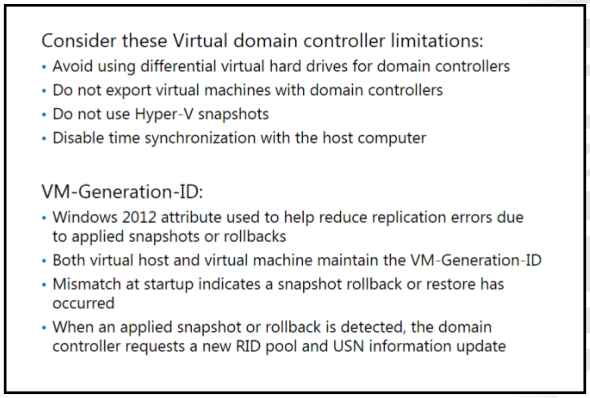

Consider these Virtual domain controller limitations:

- Avoid using differential virtual hard drives for domain controllers
- Do not export virtual machines with domain controllers
- Do not use Hyper-V snapshots
- Disable time synchronization with the host computer

VM-Generation-ID:

- Windows 2012 attribute used to help reduce replication errors due to applied snapshots or rollbacks
- Both virtual host and virtual machine maintain the VM-Generation-ID
- Mismatch at startup indicates a snapshot rollback or restore has occurred
- When an applied snapshot or rollback is detected, the domain controller requests a new RID pool and USN information update

Virtualization platforms such as Hyper-V offer a number of features that facilitate managing, maintaining, backing up, and migrating computers. However, there are some common deployment practices and features that you should not use for virtual domain controllers. The following list describes the practices to avoid when deploying domain controllers:

- Do not implement differencing disk virtual hard drives on a virtual machine that you are configuring as a domain controller. This makes it too easy to revert to a previous version and it decreases performance.
- Do not clone an operating system installation without using the Sysprep (sysprep.exe) tool, because the security identifier (SID) of the computer will not update.
- To help prevent a potential USN rollback, do not deploy additional domain controllers using copies of a .vhd or .vhdx file that represents an already deployed domain controller. Additionally, to help avoid potential USN rollbacks:
  - Do not use the Hyper-V Export feature to export a virtual machine that is running a domain controller.
  - Do not use Hyper-V snapshots on virtual domain controllers, except in lab or testing environments.
- For virtual machines that you configure as domain controllers, disable time synchronization with the host, and accept the default Windows Time service (W32time) domain-hierarchy time synchronization.

  Host-time synchronization makes it possible for guest operating systems to synchronize their system clocks with the system clock of the host operating system. Because domain controllers have their own time-synchronization mechanism, you must disable host-time synchronization on virtual machines that you configure as domain controllers. If domain controllers synchronize time both from their own source and from the host, the domain controller clock can change frequently. Because many domain controller tasks are related to the system time, a change in the system time could cause lingering objects to be left in the directory and replication may stop.

To disable host-time synchronization, in Hyper-V Manager, access the virtual machine settings in the Integration Services section, and then clear the Time Synchronization check box.

VM-Generation-ID

Windows Server 2012 introduces a new attribute that both Hyper-V and AD DS use, called **VM-Generation-ID**. This attribute helps safeguard virtual domain controllers from replication issues introduced via snapshots or cloning. Whenever a Windows Server 2012 domain controller starts the virtual machines, it compares **VM-Generation-ID** to the virtual host's **VM-Generation-ID**. If there is a mismatch, it is because either a snapshot has been applied, or cloning has taken place. When a mismatch is detected, the domain controller requests a fresh RID pool and USN information to protect AD DS.

Cloning Domain Controllers

Windows Server 2012 introduces virtualized domain controller cloning. Cloning a virtualized domain controller can present numerous challenges. For example, two domain controllers with the same name, invocation ID, and SID cannot coexist in the same forest. For Windows Server versions prior to Windows Server 2012, you could create virtualized domain controllers by deploying a base-server image on which you ran Sysprep, and then manually promoting it to become a domain controller. Windows Server 2012 provides specific virtualization capabilities to Active Directory virtualized domain controllers to resolve those issues.

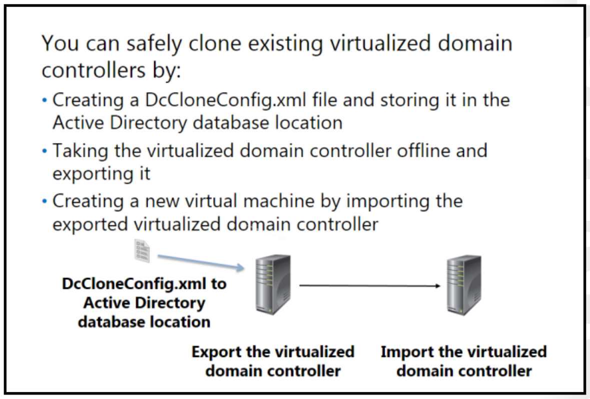

You can safely clone existing virtualized domain controllers by:

- Creating a DcCloneConfig.xml file and storing it in the Active Directory database location
- Taking the virtualized domain controller offline and exporting it
- Creating a new virtual machine by importing the exported virtualized domain controller

DcCloneConfig.xml to Active Directory database location
Export the virtualized domain controller
Import the virtualized domain controller

Virtualized domain controllers in Windows Server 2012 have two new capabilities:

- You can clone domain controllers safely to deploy additional capacity and save configuration time.
- Accidental restoration of domain controller snapshots does not disrupt the Active Directory environment.

Planning for Domain Controller Cloning

If you plan to clone virtual domain controllers, your environment must meet the following requirements.

- Use a virtualization product that supports the **Generation ID** property. This property is currently supported on Hyper-V in Windows 2012, and on VMware vSphere 5.0 patch 4 and newer vSphere versions.
- The domain controller must be running Windows Server 2012.
- The PDC emulator must be available on a Windows Server 2012 domain controller before beginning the domain controller cloning process.
- The forest functional level must be at least Windows Server 2003.
- The schema version must be Windows Server 2012 (version 56).
- The domain controller to be cloned must be in the Cloneable Domain Controllers group.

Safe Cloning

A cloned domain controller uses the Sysprep tool automatically (based on settings in DefaultDCCloneAllowList.xml). It also promotes itself to being a domain controller with the existing local Active Directory data as installation media.

Safe Backup and Restore

Rolling back to a previous snapshot of a virtualized domain controller is problematic because AD DS uses multi-master replication that relies on transactions being assigned USNs. The virtualized domain controller attempts to assign USNs to prior transactions that have already been assigned to valid transactions. This causes inconsistencies in the Active Directory database.

Windows Server 2012 implements a process called USN rollback protection. With USN rollback protection, the virtualized domain controller replicates and must be forcibly demoted or manually restored non-authoritatively. Windows Server 2012 now detects rollbacks, and non-authoritatively synchronizes the delta of changes between a domain controller and its partners for both AD DS and SYSVOL. You can now use snapshots without risk of permanently disabling domain controllers and requiring manually forced demotion, metadata cleanup, and re-promotion.

Creating a Virtualized Domain Controller Clone

To create a virtualized domain controller clone in Windows Server 2012, perform the following high-level steps:

1. Create a DcCloneConfig.xml file that contains the unique server configuration.
2. Copy this file into the Active Directory database location. By default, this location is C:\Windows\NTDS.
3. Take the virtualized domain controller offline and then either export it or copy it.
4. Create a new virtual machine by importing the exported virtual machine. This virtual machine is promoted automatically as a unique domain controller.

Lesson 5
Designing Highly Available Domain Controllers

As a key service within Windows-based network environments, AD DS must be made available as much as possible. To provide the maximum level of availability, you must carefully design the availability of all components that affect AD DS. By doing this correctly, you can make AD DS highly available.

In this lesson, you will learn about design concepts for making AD DS highly available.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe how to plan for high availability.
- Describe components of an AD DS high availability design.
- Describe considerations for designing highly available domain controllers.
- Describe considerations for designing highly available global catalog servers.
- Describe considerations for designing a highly available DNS infrastructure.
- Describe considerations for designing a highly available network infrastructure.
- Describe considerations for designing backup and recovery in AD DS.

Planning for High Availability

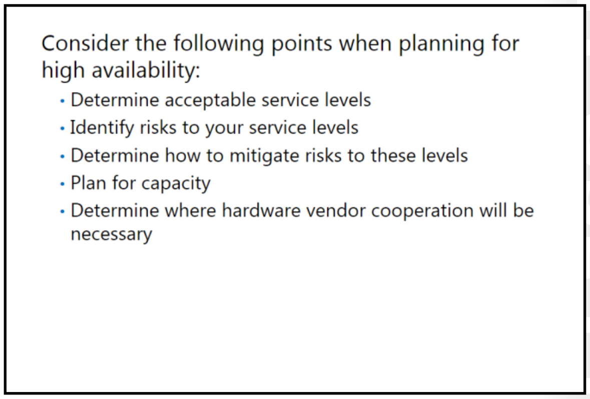

Consider the following points when planning for high availability:

- Determine acceptable service levels
- Identify risks to your service levels
- Determine how to mitigate risks to these levels
- Plan for capacity
- Determine where hardware vendor cooperation will be necessary

Availability refers to a level of service that applications, services, and systems provide. Availability is measured as a percentage of time that a service or system is available. Highly available systems have minimal downtime, whether planned or unplanned, and are available more than 99 percent of the time. For example, a system that is unavailable for 8.75 hours per year would have a 99.9 percent availability rating. Lower acceptable downtimes have greater costs. Therefore, the level of high availability should depend on an organization's requirements and budget.

To improve availability, you must implement fault-tolerant mechanisms that mask or minimize how failures of the service's elements and dependencies affect the system. You can achieve fault tolerance by implementing redundancy to single points of failure.

Because many services depend on AD DS, high availability is crucial for AD DS. For example, when AD DS is offline, users cannot log on to the network. In addition, access to resources is impaired or impossible, services such as Exchange Server do not work, and authentication does not function.

When determining business requirements regarding high availability for AD DS, you should consider priorities and budget. You can effectively support business requirements by negotiating each service's priority. This identifies where to target resources. Some measures that you can use to help you define the priority of services are as follows:

- The number of users that are impacted, and the severity of impact
- The number of external customers that are impacted, and the severity of that impact
- The number of other services that are impacted, and the severity of that impact

Consider the following points when planning a high availability strategy:

- Determine acceptable service levels. What is an acceptable service level for AD DS?
- Identify risks to your service levels. What are the most probable risks to your system?
- Determine how to mitigate risks to these levels. How can you minimize the probability of risk?
- Plan for capacity. You must not have single point of failure. For example, if one domain controller fails, you must have built-in capacities to ensure that another domain controller is available.
- Determine when hardware vendor cooperation will be necessary. What is the procedure for hardware parts replacement? Is there a service level agreement (SLA) with the hardware vendor? What is the maximum time for hardware replacement?

Components of an Active Directory High Availability Design

To make AD DS highly available, you must have several components present and available in the network infrastructure. These components include:

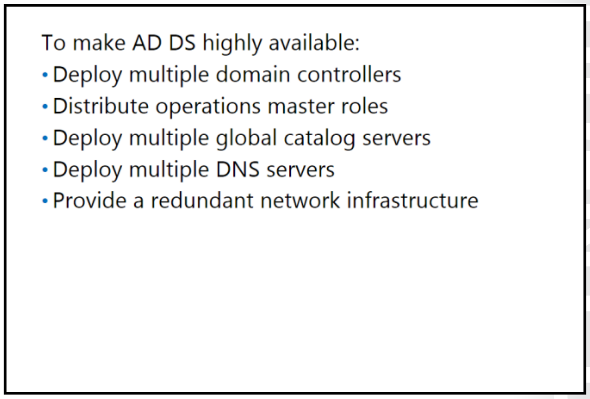

To make AD DS highly available:
- Deploy multiple domain controllers
- Distribute operations master roles
- Deploy multiple global catalog servers
- Deploy multiple DNS servers
- Provide a redundant network infrastructure

- Multiple domain controllers. You should have multiple domain controllers and global catalog servers. Client computer–initiated DNS lookups locate domain controllers and global catalog servers. If DNS lists multiple domain controllers and global catalog servers, then the client computer selects a domain controller or global catalog server in the local site first. If no domain controller or global catalog server is available in the local site, then the client computer uses a domain controller or global catalog server in a remote site. This process is automatic, and it does not require configuration.

  However, there are some exceptions to this process. Automatic failover to alternate domain controllers and global catalog servers does not work when Exchange Server runs on a domain controller or global catalog server. If Exchange Server runs on a domain controller, it only performs Active Directory queries locally. This is one reason we do not recommend installing Exchange Server on domain controllers.

- Distributed operations master roles. As described previously, you should try to distribute operations master roles to avoid losing all operations master roles if the hosting domain controller fails. For most operations master roles, it is not critical if the role holder is offline for a short time. However, you should have a domain controller available in the same Active Directory site as the operations master role holder, to ensure that you can seize the role on the second server if the current role holder fails.
- Multiple global catalog servers. Multiple global catalog servers are necessary for many features in AD DS, because they contain a subset of all information from AD DS in the Active Directory forest. The global catalog is used intensively when users log on to the network locally and externally, such as to access Microsoft Outlook Web App, and when users perform searches on AD DS or Exchange Server.

- Multiple DNS servers. To make internal DNS highly available, you must use multiple DNS servers with DNS information synchronized between them. By default, the DNS zones for AD DS are Active Directory–integrated, and are replicated between all DNS servers in the Active Directory forest.

  You also must configure member servers and clients to use multiple DNS servers. Dynamic Host Configuration Protocol (DHCP) configures most clients with IP address information automatically. You can add multiple DNS servers to the DHCP server configuration, and DHCP configures those DNS servers automatically on all DHCP clients. When you use multiple DNS servers, the DNS clients access the primary DNS server until it is unavailable, at which point they fail over to the next listed server.

- Redundant network infrastructure. This includes backup links between sites, redundant switches that connect domain controllers and clients, and a redundant cooling system.

You can make your Active Directory environment highly available by configuring these components properly.

Considerations for Designing Highly Available Domain Controllers

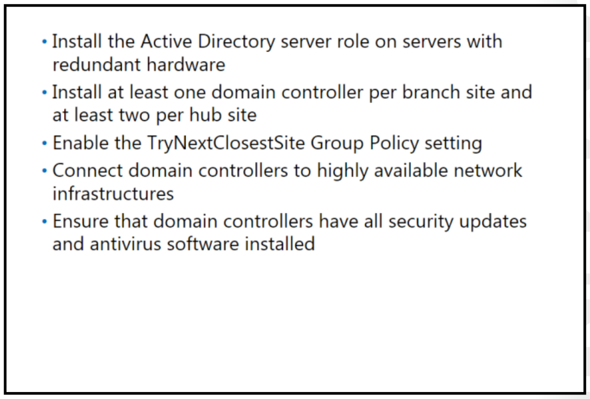

- Install the Active Directory server role on servers with redundant hardware
- Install at least one domain controller per branch site and at least two per hub site
- Enable the TryNextClosestSite Group Policy setting
- Connect domain controllers to highly available network infrastructures
- Ensure that domain controllers have all security updates and antivirus software installed

Designs for highly available domain controllers differ depending on the network infrastructure. Consider the following points when you are planning high availability for domain controllers:

- Install the Active Directory server role on servers with redundant hardware. By making hardware redundant, you increase the availability of the domain controller itself. If you cannot use redundant hardware in the machine on which the domain controller runs, ensure that you have spare parts available so that you can replace them quickly.

- Install at least one domain controller per branch site, and at least two per hub site. The most efficient way to make AD DS highly available is to deploy more than one domain controller. We recommend having at least two domain controllers in each hub site. Additionally, if you configure branch sites as Active Directory sites, you should deploy at least one domain controller per branch site.

- Enable the TryNextClosestSite Group Policy setting. If you have a domain controller that runs Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012, you can enable client computers that run the Windows Vista®, Windows 7, Windows 8, Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012 operating systems to locate domain controllers more efficiently by enabling the TryNextClosestSite Group Policy setting. This setting improves the Domain Controller Locator by helping streamline network traffic, particularly in larger enterprises that have many branch offices and sites.

  Additionally, this setting can affect how you configure site link costs because it affects the order in which you locate domain controllers. For enterprises that have many hub sites and branch offices, you can reduce Active Directory network traffic significantly by ensuring that clients fail over to the next closest hub site when they cannot find a domain controller in the closest hub site. If you do not configure the TryNextClosestSite setting, clients will randomly choose a site and domain controller for authentication if a domain controller is not available in their site.

Note: As a best practice, if you enable the TryNextClosestSite setting, you should simplify your site topology and site link costs as much as possible. In organizations with many hub sites, this can simplify the plans that you make for managing situations in which clients in one site need to fail over to a domain controller in another site.

- Connect domain controllers to highly available network infrastructures. If possible, try to have highly available network infrastructure components so that you avoid a single point of failure in that segment.
- Ensure that you install all security updates and antivirus software on all domain controllers. When you are working with domain controllers, it is critical that you keep up to date with security updates and antivirus protection. You should ensure that you update domain controllers regularly.

Note: High availability techniques in single-site environments might differ from techniques that you implement in multisite Active Directory environments.

Considerations for Designing Highly Available Global Catalog Servers

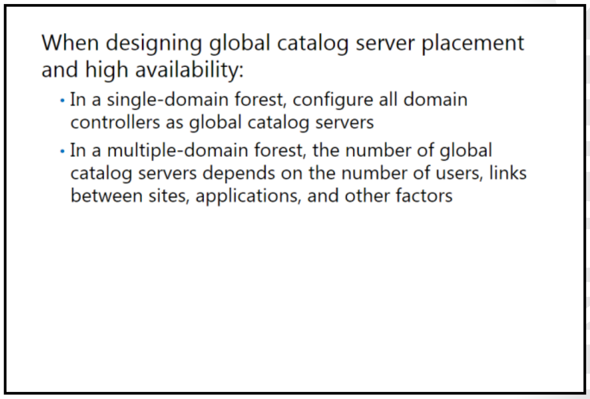

When designing global catalog server placement and high availability:

- In a single-domain forest, configure all domain controllers as global catalog servers
- In a multiple-domain forest, the number of global catalog servers depends on the number of users, links between sites, applications, and other factors

Global catalog servers facilitate the lookup of information to all domains in a forest, specifically to domains outside the current domain. The global catalog is a subset of information from each domain that replicates to every global catalog server in the forest. Applications such as Exchange Server rely heavily on the global catalog for relevant information. The logon process also uses the global catalog to enumerate universal group memberships.

Because any domain controller also can be a global catalog server, high availability can be relatively easy to achieve. Keep in mind, however, that deploying multiple global catalog servers increases network traffic between domain controllers, which can be a problem in scenarios where multiple sites exist and links are slow.

Certain applications, such as Exchange Server, Microsoft Message Queuing (also known as *MSMQ*), and applications using the DCOM do not deliver adequate response over latent WAN links. Therefore, you require a highly available global catalog infrastructure to provide low-query latency. Determine whether any applications that perform poorly over a slow WAN link are running in locations or whether the locations require Exchange Server.

If your locations include applications that do not deliver adequate response over a WAN link, you reduce query latency by placing a global catalog server at the location.

All global catalog services physically reside on one or more domain controllers. There is no way to separate global catalog functionality from a domain controller.

When designing global catalog placement and high availability, follow these guidelines:

- Single forest, single domain environment. In a single-domain forest, configure all domain controllers as global catalog servers. Because every domain controller stores the forest's only domain directory partition, configuring each domain controller as a global catalog server does not require any additional disk space usage, CPU usage, or replication traffic. In a single-domain forest, all domain controllers act as virtual global catalog servers in that they can all respond to any authentication or service request.

- Single forest, multiple domains. The number of global catalogs depends on the number of users, links between sites, applications, and other factors. However, you must have at least one global catalog server per domain. If you use an application that uses AD DS intensively, consider deploying more than one global catalog server per domain.

As a best practice, always examine global catalog availability and domain controller availability. If possible, make all domain controllers global catalog servers, and you will achieve high availability for this service.

Considerations for Designing a Highly Available DNS Infrastructure

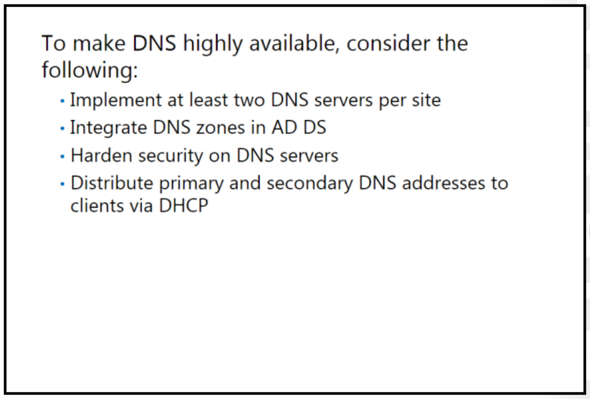

To make DNS highly available, consider the following:

- Implement at least two DNS servers per site
- Integrate DNS zones in AD DS
- Harden security on DNS servers
- Distribute primary and secondary DNS addresses to clients via DHCP

In an AD DS environment, DNS is a critical service that enables client computers and servers to locate domain controllers and various network services, and to perform name resolution for both internal and external names.

You can build a highly available DNS system in much the same way that you would build a highly available file server, except that the quantity of data typically is much smaller.

Your highly available DNS solution must ensure that the client requesting name resolution can always find a DNS server that contains your zone data. Therefore, your primary concern is the availability of the DNS server. The most complex highly available DNS solution should have two or three servers with complete copies of all the host records that you want to publish.

To create redundant DNS servers, you install the DNS service on two or more servers. You can specify that the DNS data reside in AD DS, in which case Active Directory replication performs the DNS transfers, and you do not have to specify primary and secondary DNS servers.

Additionally, you can create intermediary DNS servers that simply cache responses from your DNS servers without holding a copy of the DNS database. These caching servers reduce the load on your primary and secondary DNS servers by reducing the number of queries that ultimately reach those servers. Other DNS servers on the Internet may cache your DNS records, which increases the resolving capacity of your system at no cost.

Consider the following points when you are designing a highly available DNS solution:

- Implement at least two DNS servers per site.
- Integrate DNS zones in AD DS.
- Harden security on DNS servers by implementing additional security features and technologies, such as DNS Security Extensions (DNSSEC) and cache pollution protection.
- Distribute primary and secondary DNS addresses to clients via DHCP.

Considerations for Designing a Highly Available Network Infrastructure

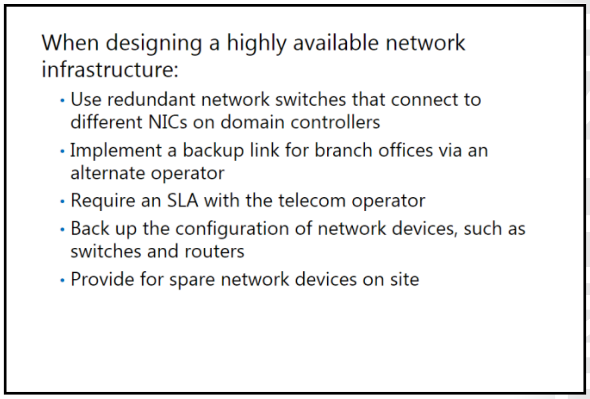

When designing a highly available network infrastructure:

- Use redundant network switches that connect to different NICs on domain controllers
- Implement a backup link for branch offices via an alternate operator
- Require an SLA with the telecom operator
- Back up the configuration of network devices, such as switches and routers
- Provide for spare network devices on site

When planning a highly available network, you must consider both the local area network (LAN) and WAN. Each of these networks has specific high availability requirements that you must analyze separately.

Highly Available LANs

You must introduce redundant components to make a LAN highly available, which typically requires that you use redundant switches. This ensures that the failure of a single switch does not affect overall network traffic. You must configure redundant network interface cards (NICs) on the computer to make the network connectivity for any individual computer fault tolerant.

Highly Available WANs

In many cases, the WAN service provider is responsible for level of WAN connectivity availability. Your organization typically agrees to the availability level as part of the SLA with the service provider. You also can create a WAN with private links between your locations.

WAN links that experience frequent outages can cause significant productivity loss to users if the location does not include a domain controller that can authenticate the users. If your WAN link availability is not operating at 100 percent, and if your remote sites cannot tolerate a service outage, place a regional domain controller in locations where the users log on or access Exchange Server when the WAN link is down.

If your WAN link availability is highly reliable, placing a domain controller at the location depends on the logon performance requirements over the WAN link. Factors that influence logon performance over the WAN include link speed and available bandwidth, the number of users and usage profiles, and the amount of logon network traffic versus replication traffic.

Additionally, consider implementing high availability in your network infrastructure design plan. Because domain controllers, servers, and clients connect via network infrastructure, you must have redundant network equipment.

When designing a highly available network infrastructure:

- Use redundant network switches that connect to different NICs on domain controllers.
- Implement a backup link for branch offices via an alternative telecommunications provider.
- Require an SLA with the telecommunications provider.
- Back up the configuration of network devices, such as switches and routers.
- Provide for spare network devices on site.

Considerations for Backup and Recovery in AD DS

Options for AD DS Backup

Windows Server Backup enables you to back up and restore a server, its roles, and its data. Windows Server Backup is installed as a feature in Server Manager.

Note: Windows Server Backup MMC is accessible via the Tools list in Server Manager. However, you first must manually add the feature.

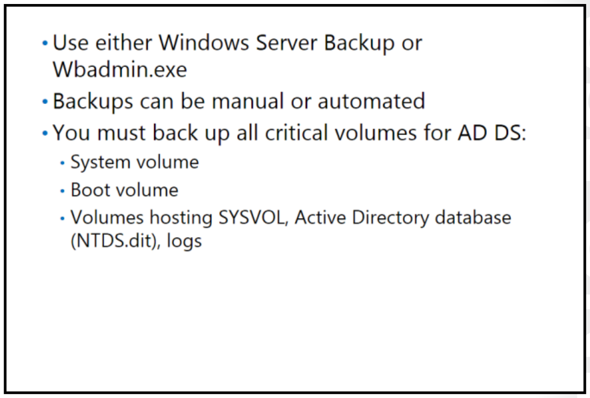

- Use either Windows Server Backup or Wbadmin.exe
- Backups can be manual or automated
- You must back up all critical volumes for AD DS:
  - System volume
  - Boot volume
  - Volumes hosting SYSVOL, Active Directory database (NTDS.dit), logs

Windows Server Backup provides a snap-in administrative tool and the Wbadmin command-line tool (Wbadmin.exe). Both the snap-in and Wbadmin enable you to perform manual or automatic backups to an internal or external disk volume, a remote share, or optical media.

Note: Windows Server Backup no longer supports backing up to tape.

In earlier versions of Windows Server, backing up Active Directory involved creating a backup of the system state. In Windows Server 2012, the system state still exists, but it is physically larger. Because of interdependencies between server roles, physical configuration, and AD DS, the system state is now a subset of a Full Server backup, and in some configurations, it might be just as large as a full server back up. To back up a domain controller, you must back up all critical volumes fully.

Windows Server Backup enables you to perform the following types of backups:

- Full server
- Selected volumes
- System state
- Individual files or folders

When you use Windows Server Backup to back up the critical volumes on a domain controller, the backup includes all data that resides on the volumes that host the following:

- Boot files, which consist of the Bootmgr file and the Boot Configuration Data (BCD) store
- Windows operating system and registry
- SYSVOL tree
- Active Directory database (Ntds.dit)
- Active Directory database log files

To perform a backup, you must first install the Windows Server Backup feature. You can then use the Windows Server Backup console to create backup jobs. You use the Actions pane in the Windows Server Backup console to start a wizard to perform a scheduled backup, or to perform a one-time backup manually. The wizard prompts for a backup type, backup selection, backup destination, and schedule (if performing a scheduled job).

Active Directory Recycle Bin

The Active Directory Recycle Bin was introduced in Windows 2008 R2. Previously, you could only access this feature by using Windows PowerShell cmdlets and the LDAP tool. In Windows Server 2012, you can now access the Active Directory Recycle Bin from the Active Directory Administrative Center. This simplifies recovery for Active Directory objects that were deleted and now need to be retrieved. The Active Directory Administrative Center lets administrators enable the Active Directory Recycle Bin, and locate or restore deleted objects in the domain.

The Active Directory Recycle Bin has the following characteristics:

- You must enable the Active Directory Recycle Bin manually.
- Once you enable the Active Directory Recycle Bin, you cannot disable it.
- The Active Directory Recycle Bin cannot restore sub-trees of objects in a single action. For example, if you delete an organizational unit (OU) with nested OUs, users, groups, and computers, restoring the base OU does not restore the child objects. You must do this in subsequent operations.
- The Active Directory Recycle Bin requires at least Windows Server 2008 R2 forest functional level.
- You must be a member of the Enterprise Admin group to use the Active Directory Recycle Bin.
- The Active Directory Recycle Bin increases the size of the Active Directory database (ntds.dit) on every domain controller in the forest. The Active Directory Recycle Bin consumes more disk space over time because it preserves objects and all attribute data.
- The Active Directory Recycle Bin preserves objects for an amount of time that matches the tombstone lifetime of the forest. This is 180 days by default.
- After you enable the Active Directory Recycle Bin, you can view deleted restorable objects in the Deleted Objects folder.

Options for Restoring AD DS

When a domain controller or its directory is corrupted, damaged, or fails, you can use one of the following restore options: typical, authoritative, full server, and alternate location.

Typical Restore

The first domain controller or directory restore option is called typical restore, or *nonauthoritative* restore. In a normal restore operation, you restore a backup of AD DS as of a known good date. Effectively, you move the domain controller back in time. When AD DS restarts on the domain controller, the domain controller contacts its replication partners and requests all subsequent updates. The domain controller synchronizes with the rest of the domain by using standard replication mechanisms.

A typical restore is useful when the directory on a domain controller was damaged or corrupted, but the problem has not yet spread to other domain controllers. This method does not work if you are trying to restore a deleted object and the deletion has replicated to the other domain controllers.

Authoritative Restore

If the typical restore does not work, you can perform an authoritative restore. In an authoritative restore, you restore the last known good version of AD DS, just as you do in a typical restore. However, before restarting the domain controller, you mark the objects that you want to recover as authoritative so that they replicate from the restored domain controller to its replication partners. These objects may include deleted objects, objects that were modified inadvertently, or an object that you want to return to a previous state for any reason. When you mark objects as authoritative, Windows increments the version number of all object attributes to be so high that the version is guaranteed to be higher than the version number of the deleted object on all other domain controllers. When you restart the restored domain controller, it replicates from its replication partners all changes to the directory. It also notifies its partners that it has changes, and the version numbers of the changes ensure that partners replicate the changes throughout the directory service.

Full Server Restore

The third option for restoring the directory service is to restore the entire domain controller. You do this by booting to the Windows Recovery Environment (Windows RE) and restoring a full server backup of the domain controller. By default, this is a typical restore. If you also must mark objects as authoritative, you must restart the server in the Directory Services Restore Mode, and then set those objects as authoritative before starting the domain controller into normal operation.

Alternate Location Restore

Finally, you can restore a system state backup to an alternative location. This enables you to examine files, and potentially mount the NTDS.dit file as described in the previous lesson. You should not copy the files from an alternative restore location over the production versions of those files. Also, do not perform a selective restore of AD DS.

Lab: Designing and Implementing an Active Directory Domain Services Physical Topology

Scenario

A. Datum Corporation has been experiencing some issues with its current Active Directory deployment. Because of these issues, the Active Directory design team must review the current Active Directory site and replication topology, and then provide recommendations for making changes. You are tasked with investigating the potential cause of the issues, and then designing a solution. You have gathered documentation about the A. Datum network, and are working towards a potential network design.

Supporting Documentation

Email from Bill Malone, Senior Manager at A. Datum:

Brad Sutton

From:	Bill Malone [Bill@adatum.com]
Sent:	06 Sep 16:22
To:	Brad@Adatum.com
Subject:	A. Datum Physical Topology
Attachments:	Adatum.vsd; Location_details.docx

Brad,

I managed to track down the reports that your predecessor produced. Rather than bury you in paperwork, I've produced a summary of the problems that we've been experiencing.

- Occasionally, it takes a very long time for users to log on, especially for users in London and Paris.
- Messages sent through Exchange Server can take several minutes to arrive.
- The Exchange Server administration team has expressed concerns with the current Active Directory deployment because of the sporadic performance issues with email flow. The team has provided performance counters that clearly show that most of the Exchange Server performance issues relate directly to slow responses from domain controllers.
- IT management has requested that the design ensure improved performance of Exchange Server. Currently, A. Datum deploys Exchange servers in Paris and London.
- The Active Directory design team also must provide a site design for various new branch offices. These offices connect to the Paris office with high-speed links. Because these offices are relatively small (but growing), they do not have dedicated IT staff. The IT staff from Paris might not deploy new domain controllers in these locations, but they are considering alternatives in case the link between Paris and elsewhere go down unexpectedly.
- A production facility will be located in Rome. Users in Rome will use the link to Paris heavily during work hours for database access and for file sharing. Currently, an Active Directory site exists for the London location. At each location, except for the new locations, at least one domain controller exists.

I have attached a network diagram that your predecessor produced, along with details about the distribution of staff across the locations.

Regards,

Bill

Adatum.vsd

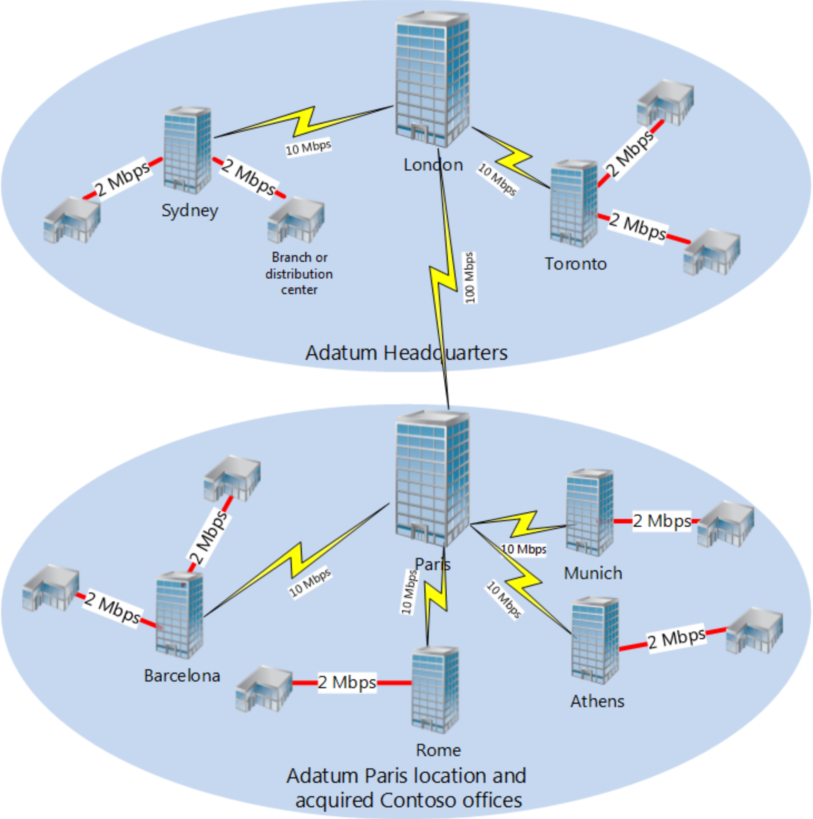

Location_details.docx

The following table lists the current locations for London.

Location	Function	Characteristics
London, England	Main headquarters: Production, distribution, and back office	Employees: 15,500
Toronto, Canada	Main regional hub office: North American production and distribution facility	Employees: 7,800
Sydney, Australia	Main regional hub office: Pacific production and distribution facility	Employees: 3,500

The following table lists the current locations for Paris.

Location	Function	Characteristics
Paris, France	Head office: Planned role as sales, marketing, and distribution center for Europe	Current employees: 1,800 (Planned employees: 4,700)

Location	Function	Characteristics
Rome, Italy	Regional hub office	Employees: 250
Barcelona, Spain	Regional hub office	Employees: 200
Munich, Germany	Regional hub office	Employees: 200
Athens, Greece	Regional hub office	Employees: 200

The following table lists regional branches and distribution centers.

Location	Number of servers	Number of users (total across all branches)	Branches
Munich	One at each branch	100	3
Barcelona	One at each branch	250	5
Rome	One at each branch	250	5
Athens	One at each branch	75	2
Sydney	One at each branch	100	2
Toronto	One at each branch	200	3

Objectives

At the end of this lab, you will be able to:

- Design an Active Directory sites strategy.
- Plan the placement of domain controllers within those sites.
- Implement the Active Directory sites design.

Lab Setup

Estimated Time: 75 minutes

Virtual machines	20413C-LON-DC1 20413C-LON-RTR 20413C-LON-SVR4
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 **20413C-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.

4. Sign in using the following credentials:

- User name: **Administrator**
- Password: **Pa$$w0rd**
- Domain: **Adatum**

5. Repeat steps 2 through 4 for 20413C-LON-RTR and 20413C-LON-SVR4.

Exercise 1: Designing Active Directory Sites and Replication

Scenario

Currently, the combined A. Datum Corporation forest is configured as part of a single site. You must produce a new site design that addresses the concerns of management, resolves issues that have been identified, and provides for future expansion of the messaging system across Europe.

Initial A. Datum Site Design	
Document Reference Number: BS0907/1	
Document Author Date	Brad Sutton 7th September
Requirements Overview Design an Active Directory site strategy to support the following objectives: • Improve logon times at all locations. • Improve mail flow throughout the organization.	
Additional Information • Exchange Server currently is deployed only in A. Datum at London and Paris. • Currently the only domain controllers are in the London main headquarters location. • London and Paris are connected by a 100 Mbps link. • Both Sydney and Toronto are connected to London by 10-Mbps links. • All regional hubs in Europe are connected to Paris by 10-Mbps links. • All branch offices and distribution centers in A. Datum are connected to the nearest regional hub office by a 2-Mbps link. • All branches have local file servers. • Only the default site object exists in AD DS at present.	
Proposals 1. How can you address problems that users currently have with logon times? 2. How can you address problems that users currently have with Exchange Server?	

Initial A. Datum Site Design
3. Do you need to create new Active Directory sites? If yes, for which offices?
4. Do you need to restructure any existing sites?
5. What else do you need to plan for the new sites?
6. Are there any alternative solutions?

The main tasks for this exercise are as follows:

1. Read the supporting documentation

2. Update the proposal document with your planned course of action

3. Examine the suggested proposals in the Lab Answer Key

4. Discuss your proposed solution with the class, as guided by your instructor

▶ Task 1: Read the supporting documentation

- Read the supporting documentation provided.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the Initial A. Datum Site Design document.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones in the Lab Answer Key.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you should have created a suitable Active Directory site design for A. Datum Corporation.

Exercise 2: Planning the Placement of Domain Controllers and Active Directory Replication

Scenario

Your Active Directory site design has been accepted, and now you need to create a design for deploying Active Directory domain controllers in the environment.

Note: You may use the information supplied in the preceding exercise to help to address the questions in this exercise.

Domain Controller Placement Strategy

Document Reference Number: BS0907/2

Document Author Date	Brad Sutton 7th September

Requirements Overview

Plan a domain controller placement strategy to support the following objectives:

- One of the critical business requirements for A. Datum is the availability of authentication services. In all of the company's locations, users must be able to log on to AD DS if a failure occurs for a single server or a network link. Currently, it can sometimes take a long time for users to log on to the network.
- As part of the expansion into Europe, A. Datum has acquired Contoso, which has a number of regional hubs across Europe. Contoso also plans to open small branch offices. These offices will employ between 20 and 50 employees, and will provide basic file and print services, and access to several critical business applications. A specific OU in the Adatum.com domain represents each new office. The OU for each office should contain all the user and computer objects for the corresponding sales office. These offices do not have dedicated server rooms and do not have onsite network administrators. Employees at these locations who have higher levels of IT knowledge are responsible for maintaining local IT resources and assisting with deploying IT services.
- A permanent 10-Mbps link connects the new offices to the main office. Contoso would like to rapidly deploy the IT services that those offices require, and at a minimal cost. The Active Directory design team must provide recommendations and a design for deploying the required IT services to these offices.

Identify important requirements that influence the placement design for the AD DS domain controller. Use the scenario information to determine the following:

- Number of users at each new location
- Available bandwidth between the head office and the new locations
- Authentication and Active Directory search requirements
- Security considerations for new office locations
- Availability of IT staff
- Avoidance of a single point of failure in AD DS

Additional Information

- A. Datum's current Active Directory environment consists of a single Active Directory forest that only has domain controllers running Windows Server 2012. The forest root domain is deployed in London on four domain controllers, and all five operations master roles are located on one server. This placement means that when you design your deployment for Active Directory domain controllers, you must avoid a single point of failure in AD DS as much as possible.
- A permanent 2-Mbps link connects the new offices to the appropriate regional hub office.
- All branches have local file servers.

Domain Controller Placement Strategy

Proposals

1. How will you address problems that users currently have with logon times?

2. How will you avoid having a single point of failure in AD DS?

3. How will you provide permanent authentication and Active Directory search services that do not depend on links between locations?

4. How will you provide management and maintenance of IT services in new sales locations?

5. How will you address security considerations in the new sales offices?

6. How will you achieve rapid deployment of IT services in new sales locations?

7. Are there any alternative solutions?

Active Directory Replication Considerations

Since you are adding additional sites, you also should decide if you need to modify the replication configuration. Use the scenario information to determine the following:

- Current physical network topology
- Available bandwidth between locations
- Link usage between London and Paris
- Current problems that users have

As part of your proposal, answer the following questions.

1. Which sites should you link with Active Directory site links?

Domain Controller Placement Strategy
2. Do you need to configure any additional site link attributes? If yes, which attributes should you configure, and how?
3. Do you need to configure bridgehead servers?
4. Do you need to configure site link bridging?

The main tasks for this exercise are as follows:

1. Read the supporting documentation

2. Update the proposal document with your planned course of action

3. Examine the suggested proposals in the Lab Answer Key

4. Discuss your proposed solution with the class, as guided by your instructor

▶ Task 1: Read the supporting documentation

- Read the supporting documentation provided.

▶ Task 2: Update the proposal document with your planned course of action

1. Answer the questions in the proposals section of the Domain Controller Placement Strategy document.

2. Answer the questions in the Active Directory Replication Considerations section of the Domain Controller Placement Strategy document.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones in the Lab Answer Key.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you should have created a suitable domain controller placement strategy.

Exercise 3: Implementing Active Directory Sites and Domain Controllers

Scenario

You have received approval from A. Datum management to go ahead with implementing the Active Directory sites and domain controllers. As part of this assignment, you need to install an additional domain controller. In this exercise, you will implement a part of your Active Directory site and replication design. This involves creating a new domain controller.

The main tasks for this exercise are as follows:

1. Install the Paris domain controller

2. Create the Active Directory sites

3. Configure site links

4. Move the new domain controller to the appropriate site

5. To prepare for the next module

▶ Task 1: Install the Paris domain controller

1. On LON-SVR4, Rename the computer **PAR-DC1**, and then restart it.

Note: The lab steps will refer to 20413C-LON-SVR4 as PAR-DC1.

2. Switch to PAR-DC1 (20413C-LON-SVR4), and then sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
3. On PAR-DC1, use Server Manager to install AD DS.
4. When the Active Directory binaries have installed, use the Active Directory Domain Services Configuration Wizard to install and configure **PAR-DC1** as an additional domain controller for Adatum.com.
5. After the server restarts, sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.

▶ Task 2: Create the Active Directory sites

1. On PAR-DC1, open the Server Manager console and then open Active Directory Sites and Services.
2. Create a new site with the following configuration:

- Name: **Paris**
- Site link object: **DEFAULTIPSITELINK**

3. If necessary, on LON-DC1 open Active Directory Sites and Services.
4. Create a new subnet with the following configuration:

- Prefix: **10.10.0.0/16**
- Site object: **Paris**

5. In the navigation pane, click the **Subnets** folder.
6. In the details pane, verify that the two subnets display, and are associated with their appropriate site.

▶ Task 3: Configure site links

1. On PAR-DC1, if necessary, open Active Directory Sites and Services.
2. Create a new site link named **LONDON-PARIS**, and then configure it with the following settings:

- Sites: **AdatumHQ, Paris**
- Cost: **80**
- Replication: Every **60** minutes

▶ Task 4: Move the new domain controller to the appropriate site

1. On PAR-DC1, if necessary, open Active Directory Sites and Services.
2. Move **PAR-DC1** from the **AdatumHQ** site to the **Paris** site.
3. Verify that PAR-DC1 is now located under the Servers node in the Paris site.
4. Configure PAR-DC1 as the preferred bridgehead server for the **IP** transport.

▶ Task 5: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machines** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for 20413C-LON-RTR and 20413C-LON-SVR4.

Results: After completing this exercise, you should have successfully configured Active Directory sites, domain controllers, and replication.

Question: What was your approach to the Active Directory site and replication design?

Question: How did you address the Active Directory domain controller planning exercise?

Question: How does this physical Active Directory design compare with your organization’s Active Directory implementation?

Module Review and Takeaways

Review Questions

Question: In a multisite enterprise, why is it important that you identify and associate all subnets with a site?

Question: What is the purpose of a bridgehead server?

Question: Which protocol can you use as an alternative to Active Directory replication? What is the disadvantage of using it?

Module 10

Planning and Implementing Storage and File Services

Contents:

Module Overview

With the growing importance and use of networked applications, high-performance and resilient storage has become increasingly important. When planning storage for your networked applications, you must first select a suitable storage technology that includes both local and remote storage solutions. A common remote storage technology supported by Windows Server® 2012 is Internet small computer system interface (iSCSI). In this module, you will learn about storage area networks (SANs), Storage Spaces, and file services, and how to implement them.

Objectives

After completing this module, you will be able to:

- Plan and implement iSCSI SANs.
- Plan and implement Storage Spaces.
- Optimize files services for branch offices.

Lesson 1
Planning and Implementing iSCSI SANs

After you decide to implement a distributed storage architecture, you must then choose an appropriate storage technology. Windows Server 2012 provides the components necessary for you to implement an iSCSI storage infrastructure. iSCSI storage is an inexpensive and simple way to configure connections to remote disks.

Many application requirements dictate that remote storage connections must be redundant in nature for fault tolerance or high availability. You can implement a fault-tolerant storage solution with iSCSI by using the underlying network fault tolerance, which is often less expensive than maintaining fault-tolerant storage technologies such as SANs.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the components that make up a SAN.
- Describe iSCSI.
- Compare and contrast the benefits and limitations of iSCSI SANs with Fibre Channel SANs.
- Describe an iSCSI target server and a Microsoft iSCSI initiator.
- Describe the options for implementing high availability for iSCSI.
- Describe the options for providing iSCSI security.
- Implement iSCSI.
- Describe the considerations for implementing iSCSI storage.

SAN Components

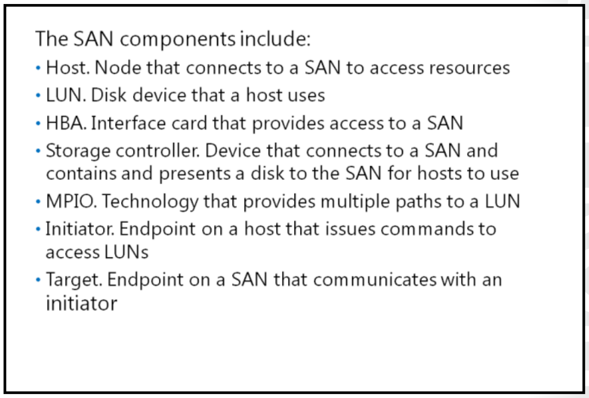

The SAN components include:

- Host. Node that connects to a SAN to access resources
- LUN. Disk device that a host uses
- HBA. Interface card that provides access to a SAN
- Storage controller. Device that connects to a SAN and contains and presents a disk to the SAN for hosts to use
- MPIO. Technology that provides multiple paths to a LUN
- Initiator. Endpoint on a host that issues commands to access LUNs
- Target. Endpoint on a SAN that communicates with an initiator

The following components are associated with SANs:

- *Host*. A host node connects to a SAN to access resources.
- *Logical unit number* (LUN). A LUN is a disk device that a host uses. A LUN is the identifier of a SCSI logical unit, and by extension, of a Fibre Channel or iSCSI logical unit.
- *Host bus adapter* (HBA). HBA refers to a Fibre Channel interface card that is placed in a server to provide access to a SAN (Fibre Channel, iSCSI, or serial-attached SCSI (SAS)). It is the functional equivalent of the network adapter in a traditional Ethernet network.
- *Storage controller*. A storage controller is a device that connects to the SAN, and that contains and presents disks to the SAN for hosts to use.

- *Mulitpath I/O* (MPIO). MPIO refers to the technology that provides multiple paths from a host to a storage LUN. These paths enable path redundancy, and aggregated bandwidth between the CPU where the application is executing, and the target where the data is stored.
- *Initiator*. The initiator refers to the software or hardware on the host computer that issues commands to access a LUN.
- *Target*. This refers to the storage controller that connects to an initiator.

What Is iSCSI?

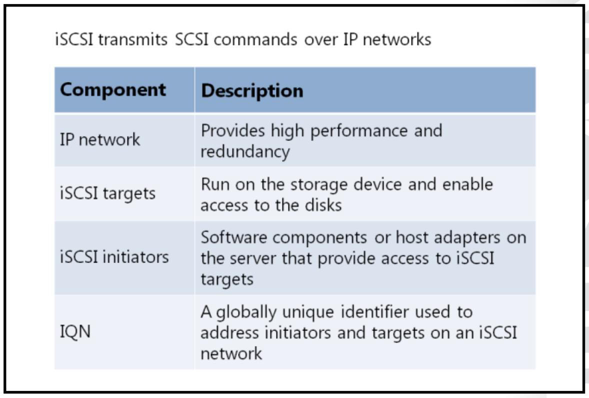

iSCSI transmits SCSI commands over IP networks

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

iSCSI is a protocol that supports access to remote, SCSI-based storage devices over a TCP/IP network. iSCSI carries standard SCSI commands over IP networks to facilitate data transfers over the intranet, 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 relies on standard Ethernet networking architecture. Specialized hardware (such as HBAs or network switches) is optional. iSCSI uses TCP/IP (typically, Transmission Control Protocol (TCP) port 3260). This means that iSCSI 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 high performance, local storage bus subsystem architecture, and emulates it over LANs and WANs, thereby creating a SAN.

Unlike some SAN technologies, iSCSI requires no specialized cabling. You can run it over the existing switching and IP infrastructure. However, you can increase iSCSI SAN deployment performance 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 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 obtaining storage access. 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 unique identifiers that address initiators and targets on an iSCSI network. When you configure an iSCSI target, you must configure the IQN for the iSCSI initiators that connect 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, you can always identify iSCSI endpoints (both target and initiator) by their IP addresses.

Comparing iSCSI SANs with Fibre Channel SANs

Fibre Channel SANs have been used in the industry for a long time and became a standardized technology in the early 1990's. Fibre Channel SANs use dedicated HBAs, which manage storage traffic based on the Fibre Channel protocol. Because all storage traffic is handled at the HBA level, the CPU is free to provide processing cycles to applications and services that are running on the host. Fibre Channel switches can work at speeds of up to 20 Gbps. Due to the use of dedicated HBAs, switches, and Fibre Channel connections, Fibre Channel SANs typically provide a higher degree of performance when compared to other storage technologies.

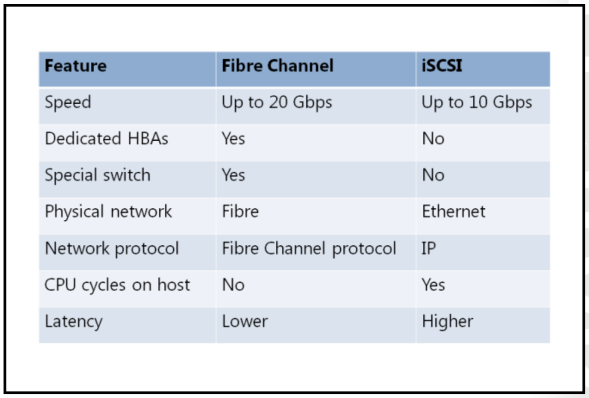

Feature	Fibre Channel	iSCSI
Speed	Up to 20 Gbps	Up to 10 Gbps
Dedicated HBAs	Yes	No
Special switch	Yes	No
Physical network	Fibre	Ethernet
Network protocol	Fibre Channel protocol	IP
CPU cycles on host	No	Yes
Latency	Lower	Higher

However, the use of dedicated components makes a Fibre Channel solution a lot more expensive than other technologies. Furthermore, investing in a Fibre Channel connection between a SAN and hosts that are physically far away from the SAN becomes even more expensive. This makes it almost impossible to maintain connectivity over to a SAN from hosts located at other sites. You can overcome some of this limitation by using Fibre Channel protocol distance gateways, but these are also expensive. Finally, Fibre Channel SANs work with World Wide Node Names and Port Names, which are similar to media access control (MAC) addresses, for LUN connectivity. Managing World Wide Node and Port names is time consuming, particularly when many changes are being made to the storage fabric.

Note: *Fabric* refers to the grouping of specific resources that are shared across a data center, such as storage and network. You will see references to storage fabric, network fabric, and compute fabric whenever you are discussing data center concepts, and more specifically private and public cloud environments.

Benefits of the iSCSI SANs

Fibre Channel had become a standard in enterprise storage before iSCSI technology was introduced. With the introduction of iSCSI in the early 2000's, the iSCSI technology soon became a viable option to small and medium enterprises that were not able to afford a Fibre Channel solution.

iSCSI does not require special hardware such as HBAs and switches because it does not use Fibre Channel protocol. Instead, it uses TCP/IP for all storage communication by encapsulating the SCSI commands in an IP frame. Because most iSCSI initiators are implemented through software, iSCSI uses more CPU cycles than any Fibre Channel implementation. Although you can overcome this limitation by using iSCSI HBAs, not every single implementation of iSCSI supports HBAs. With the advance of Ethernet, you can use 10-

Gbps networks dedicated for iSCSI traffic, and achieve performance similar to Fibre Channel at a fraction of the cost.

The following table describes the differences between Fibre Channel SANs and iSCSI SANs.

Feature	Fibre Channel SAN	iSCSI SAN
Speed	Up to 20 Gbps	Up to 10 Gbps
Dedicated HBA	Yes	In some implementations
Special switch for storage traffic	Yes	No
Physical network	Fibre	Ethernet
Network protocol	Fibre Channel protocol	IP
CPU cycles on host required	No	Yes
Latency	Lower	Higher

iSCSI Target Server and iSCSI Initiator

In order to implement iSCSI, you must understand more about the iSCSI target server and the iSCSI initiator.

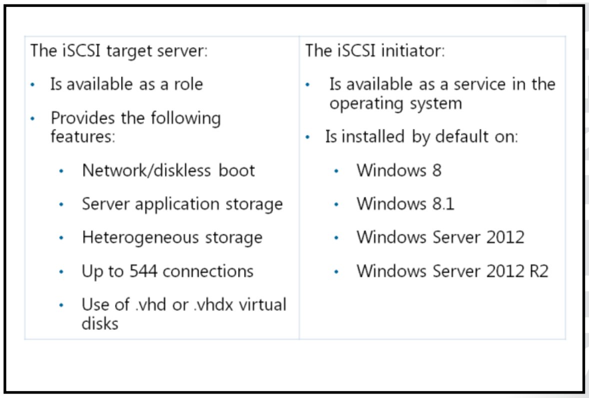

The iSCSI target server:

- Is available as a role
- Provides the following features:
  - Network/diskless boot
  - Server application storage
  - Heterogeneous storage
  - Up to 544 connections
  - Use of .vhd or .vhdx virtual disks

The iSCSI initiator:

- Is available as a service in the operating system
- Is installed by default on:
  - Windows 8
  - Windows 8.1
  - Windows Server 2012
  - Windows Server 2012 R2

iSCSI Target Server

You can manage an iSCSI SAN by using the Microsoft® iSCSI Target Server role in Windows Server 2012 R2 and Windows Server 2012. The iSCSI Target Server role 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 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 Windows Server Hyper-V® server farm, or 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 also can 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 operating system environments.

- Test lab environments. The iSCSI target server role enables 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 includes the following new and improved features in the iSCSI Target Server role:

- Virtual disks. iSCSI target servers can now use the .vhdx file format for LUNs that are being exposed to hosts.
- Integration with Microsoft System Center 2012 R2 Virtual Machine Manager (VMM). You can use VMM to manage the storage fabric by using a Storage Management Interface Specification (SMI-S) provider. Windows Server 2012 R2 includes a Storage Management Interface Specification provider for iSCSI, which you can use to better manage iSCSI target servers from VMM.
- Multiple sessions. A single iSCSI target server can manage up to 544 simultaneous connections. In Windows Server 2012, this limit was 256.

iSCSI target servers that provide block storage use 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 must have shared storage for the cluster—either hardware Fibre Channel storage, or a SAS 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 that is installed by default since Windows Server 2008 and Windows Vista®. To connect your computer to an iSCSI target, you start the iSCSI Initiator, and then configure the initiator.

The new features in Windows Server 2012 iSCSI initiator include:

- Authentication. You can enable Challenge Handshake Authentication Protocol (CHAP) to authenticate initiator connections, or you can enable reverse CHAP to enable the initiator to authenticate the iSCSI target.
- Query initiator computer for ID. This feature is only supported with Windows 8 or Windows Server 2012.

Additional Reading: For more information about the introduction of iSCSI targets in Windows Server 2012, see http://go.microsoft.com/fwlink/?linkid=279916.

Options for Implementing High Availability for iSCSI

You can integrate the iSCSI target server and iSCSI initiator settings into more advanced configurations.

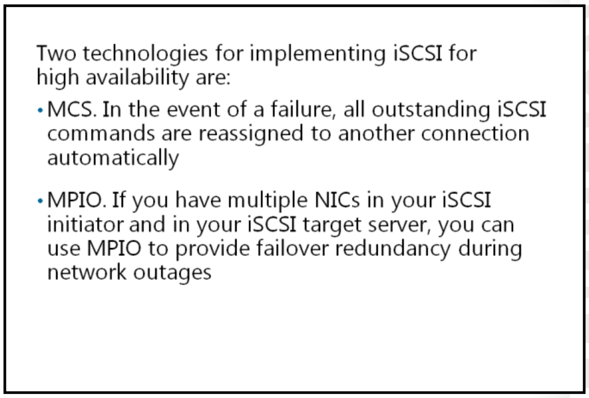

Two technologies for implementing iSCSI for high availability are:

- MCS. In the event of a failure, all outstanding iSCSI commands are reassigned to another connection automatically
- MPIO. If you have multiple NICs in your iSCSI initiator and in your iSCSI target server, you can use MPIO to provide failover redundancy during network outages

Configuring iSCSI for High Availability

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

Although MCS and MPIO achieve similar results, these two technologies use different approaches to achieve high availability for iSCSI storage connections:

- MCS is a feature of the iSCSI protocol that:
  - Enables multiple TCP/IP connections from the initiator to the target for the same iSCSI session.
  - Supports automatic failover. If a failure 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. MPIO:
  - Provides failover redundancy during network outages, if you have multiple network interface cards (NICs) in your iSCSI initiator and iSCSI target server.
  - Requires a device-specific module (DSM) if you want to connect to a non-Microsoft SAN device that is connected to the iSCSI initiator. The Windows Server operating system includes a default MPIO DSM that is installed as the MPIO feature within Server Manager.
  - Is widely supported. Many SANs can use the default DSM without any additional software, while others require a specialized DSM from the manufacturer.
  - Is more complex to configure, and is not as fully automated during failover as MCS.

iSCSI Security Options

iSCSI uses a regular Ethernet network to transport storage information. As a result, anyone with access to the network segment that iSCSI traffic uses can try to read the data being transported, or initiate a man-in-the-middle attack to access data directly from the iSCSI target server.

You can prevent these attacks by using different settings at the iSCSI target server level. You can choose to use Internet Protocol security (IPsec), create a list of access servers, or implement authentication methods in your organization.

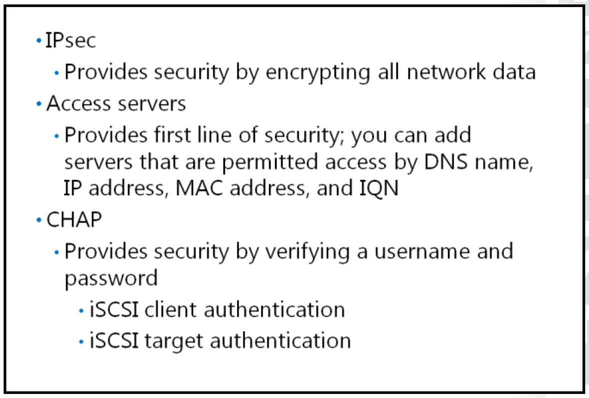

- IPsec
  - Provides security by encrypting all network data
- Access servers
  - Provides first line of security; you can add servers that are permitted access by DNS name, IP address, MAC address, and IQN
- CHAP
  - Provides security by verifying a username and password
    - iSCSI client authentication
    - iSCSI target authentication

IPsec

Because iSCSI uses an Ethernet network to transport all storage traffic, you can use IPsec on the network segment that iSCSI is using to encrypt traffic. You may consider using IPsec to prevent attacks from within the network. For example, you can place a system that is running a network protocol analyzer in the same network segment as the iSCSI traffic, and monitor all traffic that is gaining access to your data.

Access Servers

When configuring iSCSI targets, the first line of security is the list of access servers. You add to the list all servers that are permitted to access the iSCSI target. You can add servers by:

- Domain Name System (DNS) name. You can select servers by searching in Active Directory® Domain Services, or by typing their host name. This is the easiest form of access server setting to break. This is because someone can simply place a computer on the iSCSI segment that has the same name as the intended iSCSI client to receive iSCSI traffic.

- IP address. You can type the IP address of the iSCSI client that will have access to the iSCSI target. Although this is safer than specifying a DNS name, it is still easy to mimic. This is because someone can still place a system on the same network segment that the iSCSI traffic is using, and use the IP address of a known iSCSI client.
- MAC address. You can type the MAC address of the iSCSI client that will have access to the iSCSI target. This is safer than using DNS names and IP addresses; however, this method is still prone to attacks.
- IQN. You can use IQNs to specify what iSCSI clients can access an iSCSI target. IQNs are created automatically by the iSCSI target when you specify a DNS name, IP address, or MAC address for clients. They have a very distinct format; that is, all IQNs start with the letters iqn. For example, you know that iqn.2013-01.com.adatum:storage.server1.disk1 is a valid IQN because:
  - The first set of characters is always iqn.
  - The second set of characters represent the year and month that the naming authority took ownership of the domain in use.
  - The third set of characters show the domain name in reverse (com.adatum).
  - The final set of characters, after the colon are optional characters to identify the iSCSI disk.

Although using access servers provides a layer of security, this is a very thin security layer, because experienced hackers can easily find the DNS name, IP address, MAC address, and even IQN being used to start an iSCSI session with an iSCSI target. Then, they can change their system settings to use the information they discovered and start a man-in-the-middle attack.

Authentication

To prevent these forms of attacks, you can implement authentication at the iSCSI target level. There are two forms of authentication you can configure at the iSCSI target level:

- iSCSI client authentication. This forces the iSCSI client to provide authentication information to the iSCSI target server.
- iSCSI target authentication. This forces the iSCSI target to provide authentication information to the iSCSI client.

Note: Clients and targets authenticate using CHAP, which consists of verifying a username and password. As a best practice, ensure that you use strong passwords for authentication.

Demonstration: Implementing iSCSI

In this demonstration, you will:

- Add the iSCSI Target Server role.
- Create two iSCSI virtual disks and an iSCSI target.
- Connect to the iSCSI target.
- Verify the presence of the iSCSI drive.

Demonstration Steps

Add the iSCSI Target Server role

1. On LON-DC1, switch to Server Manager.

2. Use the Add Roles and Features Wizard to browse to **File And Storage Services (Installed)\File and iSCSI Services**, and install the **iSCSI Target Server**. Accept the default values.

Create two iSCSI virtual disks and an iSCSI target

1. On LON-DC1, in Server Manager, in the navigation pane, click **File and Storage Services**, and then click **iSCSI**.
2. In the iSCSI VIRTUAL DISKS pane, click **TASKS**, and then in the **TASKS** drop-down list 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-SVR1**
   - Access servers: **172.16.0.11**
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 box, click **New iSCSI Virtual Disk**.
6. Create a virtual disk that has the following settings:
   - Name: **iSCSIDisk2**
   - Disk size: **5 GB**
   - iSCSI target: **LON-SVR1**
7. On the **View results** page, wait until creation completes, and then close the **View Results** page.

Connect to the iSCSI target

1. Sign in to LON-SVR1 with the 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-DC1-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

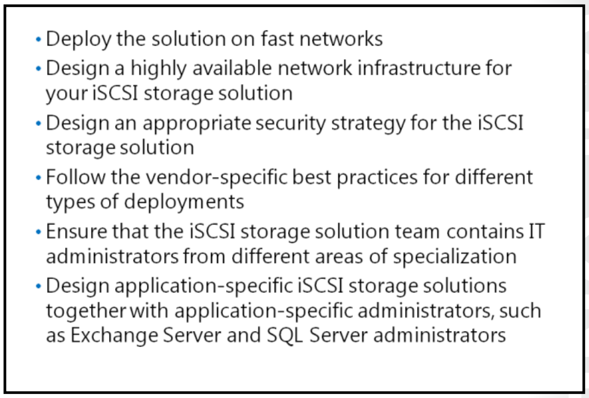

- Deploy the solution on fast networks
- Design a highly available network infrastructure for your iSCSI storage solution
- Design an appropriate security strategy for the iSCSI storage solution
- Follow the vendor-specific best practices for different types of deployments
- Ensure that the iSCSI storage solution team contains IT administrators from different areas of specialization
- Design application-specific iSCSI storage solutions together with application-specific administrators, such as Exchange Server and SQL Server administrators

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

- Deploy the iSCSI solution on at least 1-Gbps networks.
- High availability design for network infrastructure is crucial because data from servers to iSCSI storage is transferred through network devices and components. (High availability considerations were explained earlier in this module.)
- Design an appropriate security strategy for the iSCSI storage solution. (Security considerations and recommendations were explained earlier in this module.)
- Read the vendor-specific best practices for different types of deployments and applications that use iSCSI storage solution, such as Exchange Server and Microsoft SQL Server®.
- IT personnel who will be designing, configuring, and administering the iSCSI storage solution must include IT administrators from 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
Planning and Implementing Storage Spaces

Managing direct-attached storage (DAS) on a server can prove to be a tedious task for administrators. To overcome this problem, many organizations use SANs that physically group disks together. However, SANs require special configuration and sometimes special hardware, and are therefore expensive. To overcome these issues, Microsoft introduced Storage Spaces as a new feature in Windows Server 2012 and Windows Server 2012 R2. This new feature pools disks together in a similar way to SANs, and presents them to the operating system as a single disk that can span multiple physical disks in a pool. This lesson explains how to configure and implement Storage Spaces.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the use of Storage Spaces.
- Plan physical storage for Storage Spaces.
- Plan for Storage Spaces.
- Plan for storage tiers.
- Configure Storage Spaces.
- Plan high availability for Storage Spaces.
- Compare and contrast the benefits and limitations of iSCSI with Storage Spaces.
- Plan storage optimization in Windows Server 2012 R2.

Storage Spaces Overview

Storage Spaces is a storage virtualization feature that is built into Windows Server 2012, Windows 8, and newer Windows Server and Windows client operating systems. You can use Storage Spaces to add physical disks of any type and size to a storage pool, and to create highly-available virtual disks from the storage pool.

To better understand Storage Spaces, you must learn some of the terminology used when dealing with this feature:

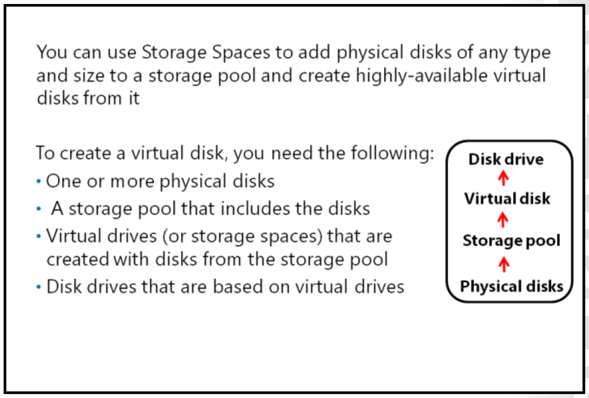

You can use Storage Spaces to add physical disks of any type and size to a storage pool and create highly-available virtual disks from it

To create a virtual disk, you need the following:

- One or more physical disks
- A storage pool that includes the disks
- Virtual drives (or storage spaces) that are created with disks from the storage pool
- Disk drives that are based on virtual drives

- Physical disk. These are physical disks, such as SAS disks, that are attached to your server. You can use any DAS or iSCSI disks in your storage pools.
- Storage pool. This is a collection of one or more physical disks that you can use to create virtual disks. You can add to a storage pool any available physical disk that is not formatted or attached to another storage pool.
- Virtual disk. This resembles a physical disk from the perspective of users and applications. However, virtual disks are more flexible because they include thin provisioning (or just-in-time (JIT) allocations) and resiliency to physical disk failures with built-in functionality such as mirroring.

- Disk drive. This is a volume that you can access from your operating system, for example by using a drive letter. After you create a virtual disk, it is presented to the operating system as a disk drive.

You can add physical disks to a server and create a storage pool, virtual disks, and disk drives. However, when you add physical disks to a pool, they must satisfy the following requirements:

- One physical drive is required to create a storage pool.
- A minimum of two physical drives are required to create a resilient mirror virtual disk.
- A minimum of three physical drives are required to create a virtual disk with resiliency through parity.
- Three-way mirroring requires at least five physical drives.
- Drives must be blank and unformatted; no volume must exist on them.
- Drives can be attached by using different bus interfaces including iSCSI, SAS, Serial ATA (SATA), small computer system interface (SCSI), and USB. You cannot use SATA, USB, or SCSI disks in a failover cluster.

Planning Physical Storage for Storage Spaces

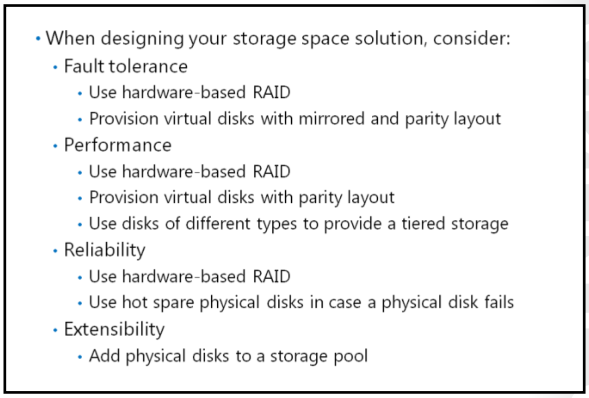

- When designing your storage space solution, consider:
  - Fault tolerance
    - Use hardware-based RAID
    - Provision virtual disks with mirrored and parity layout
  - Performance
    - Use hardware-based RAID
    - Provision virtual disks with parity layout
    - Use disks of different types to provide a tiered storage
  - Reliability
    - Use hardware-based RAID
    - Use hot spare physical disks in case a physical disk fails
  - Extensibility
    - Add physical disks to a storage pool

When creating pools, Storage Spaces can use any DAS device and iSCSI disks. You can use SATA and SAS drives (or even older integrated drive electronics (IDE) and SCSI drives) that are connected internally to the computer, or that are connected through USB.

When designing your storage space solution, you must consider the following factors:

- Fault tolerance. Do you want data to be available in case a physical disk fails? If so, you must use multiple physical disks and provision virtual disks by using mirroring or parity. You can also use hardware-based redundant array of independent disks (RAID) systems to provide the necessary availability and reliability, and then provision virtual disks by using the simple layout.
- Performance. You can improve performance for read and write actions by using hardware-based RAID systems, by using a parity layout for virtual disks, or by using a combination of both. You also need to consider the speed of each individual physical disk when determining performance.

Alternatively, you can use disks of different types to provide a tiered system for storage. For example, you can use solid-state drives (SSDs) for data that you require faster access to, and SATA drives for data that is not frequently accessed.

- Reliability. Hardware-based RAID systems and virtual disks in parity layout provide some reliability. You can improve that degree of reliability by using hot spare physical disks in case a physical disk fails.
- Extensibility. One of the main advantages of using Storage Spaces is the ability to expand storage in the future by adding physical disks. You can add physical disks any time after you create a storage pool to expand its storage capacity or to provide fault tolerance.

Planning Storage Spaces

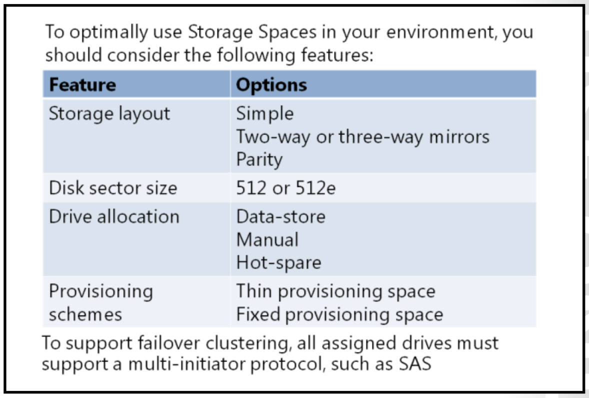

To optimally use Storage Spaces in your environment, you should consider the following features:

Feature	Options
Storage layout	Simple Two-way or three-way mirrors Parity
Disk sector size	512 or 512e
Drive allocation	Data-store Manual Hot-spare
Provisioning schemes	Thin provisioning space Fixed provisioning space

To support failover clustering, all assigned drives must support a multi-initiator protocol, such as SAS

Before you implement virtual disks to configure Storage Spaces, you must consider the features and the options they provide as described in the following table.

Feature	Description
Storage layout	This feature defines the number of disks from the storage pool that are allocated. Valid options are: • Simple. A simple space has data striping but no redundancy. In data striping, logically sequential data is segmented across all disks in a way that enables different physical storage drives to access these sequential segments. Striping makes it possible to access multiple segments of data simultaneously. Do not host important data on a simple volume, because it provides no failover capabilities if the disk in which the data is stored fails. • Two-way and three-way mirrors. Mirror spaces maintain two or three copies of the data that they host (two data copies for two-way mirrors and three data copies for three-way mirrors). Duplication occurs with every write, to ensure that all data copies are always current. Mirror spaces also stripe the data across multiple physical drives. Mirror spaces are preferable because of their greater data throughput and lower access latency. They also do not introduce a risk of corrupting at-rest data, and they do not require the additional journaling stage when writing data. • Parity. A parity space resembles a simple space. Data, along with parity information, is striped across multiple physical drives. Parity enables Storage Spaces to continue to service read and write requests even when a drive has failed. Parity is always rotated across available disks to enable I/O optimization. A storage space requires a minimum of three physical drives for parity spaces. Parity spaces have increased resiliency through journaling.
Disk sector size	A storage pool's sector size is set the moment it is created. If the list of drives being used contains only 512 and 512e drives, the pool defaults to 512e. However, if the list contains at least one 4-kilobyte (KB) drive, the pool sector size is defaulted to 4 KB. Optionally, an administrator can explicitly define the sector size that all contained spaces in the pool will

Feature	Description
	inherit. After an administrator defines this, Windows operating systems will only allow you to add drives that have a compliant sector size, that is: 512 or 512e for a 512e storage pool, and 512, 512e, or 4 KB for a 4-KB pool.
Cluster disk requirement	Failover clustering prevents interruption to workloads or data if there is a computer failure. For a pool to support failover clustering, all assigned drives must support a multi-initiator protocol, such as SAS.
Drive allocation	Drive allocation defines how the drive is allocated to the pool. Options are: • Datastore. This is the default allocation when any drive is added to a pool. Storage Spaces can select available capacity on datastore drives automatically for both storage space creation and JIT allocation. • Manual. Administrators can choose to specify manual as the usage type for drives added to a pool. A manual drive is not used automatically as part of a storage space unless it is specifically selected during creation of that storage space. Administrators can use this usage type to specify particular types of drives for use only by certain storage spaces. • Hot spare. Reserve drives that are not used in the creation of a storage space are sometimes referred to as a *hot spare*. If a failure occurs on a drive that is hosting storage space columns, a reserve drive is called on to replace the failed drive.
Provisioning schemes	You can provision a virtual disk by using one of two methods: • Thin provisioning space. Thin provisioning is a mechanism that enables storage to be easily allocated on a just-enough and JIT basis. Storage capacity in the pool is organized into provisioning slabs that are not allocated until the point in time when datasets grow to actually require the storage. Instead of the traditional fixed storage allocation method in which large pools of storage capacity are allocated but may remain unused, thin provisioning optimizes use of available storage. Organizations also may be able to reduce operating costs (such as electricity and floor space) that are associated with keeping unused drives in use. • Fixed provisioning space. In Storage Spaces, fixed provisioned spaces also use the flexible provisioning slabs. The difference is that the storage capacity is allocated at the time that the space is created.

Planning Storage Tiers

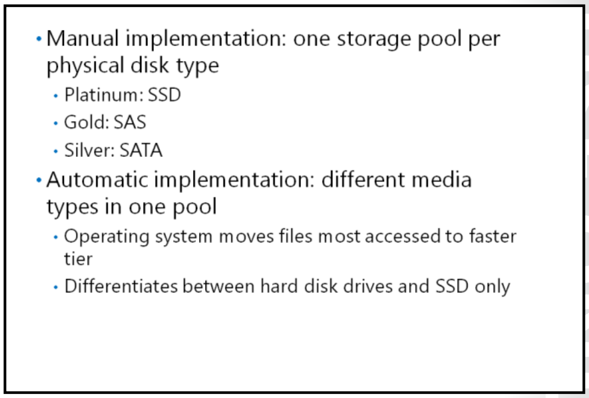

- Manual implementation: one storage pool per physical disk type
  - Platinum: SSD
  - Gold: SAS
  - Silver: SATA
- Automatic implementation: different media types in one pool
  - Operating system moves files most accessed to faster tier
  - Differentiates between hard disk drives and SSD only

When implementing Storage Spaces, you can use disks of different types, architecture, and speed. However, when combining physical disks that have different speed performance into a single pool, you may not achieve the best throughput for the virtual disks provisioned in your pool. Therefore, you should try to create different pools for each type of physical disk available. By doing so, you can create a tiered storage approach that benefits from the different disk types available, and then use tiers according to application needs.

There are two ways of implementing storage tiers in Storage Spaces: manually, or automatically.

Manual Implementation

You can achieve manual implementation by creating different pools. For example, consider a scenario where you have 5 SSDs, 10 SAS drives at 7,200 revolutions per minute (RPM), and 5 SATA drives at 5,400 RPM. You could create three pools: one with all SSD disks, one with the 7,200 RPM SAS disks, and one with the 5,400 RPM SATA disks. Your SSD disks would be part of a higher storage tier, with the SAS disks in the middle tier, and the SATA disks in the lower tier. You can use higher tier disks for storage-intense applications, similar to data files in a database server. You can use middle tier disks for applications that require better performance, but are not as storage intensive, such as log files for a database server, or file sharing solutions. Finally, you can use disks in the lower tier for other applications that do not require too much performance, such as data archival.

Note: Some organizations label their storage tiers, with names such as Platinum, Gold, and Silver.

Automatic Implementation

Beginning with Windows Server 2012 R2, you can create storage tiers within the same storage pool. You can do this if you use different types of physical disks (both hard disk drives and SSD). When the operating system recognizes that different media types are being used within the same pool, you will be able to create tiered virtual disks. You do this in the New Virtual Disk Wizard by selecting the check box, Create storage tiers on this virtual disk. When you enable storage tiers, the operating system moves files that are more frequently accessed to faster media.

Note: For storage tiers to work, you must ensure that the operating system recognizes the SSD disks as SSD, because sometimes the operating system displays the media type for SSD drives as unspecified. When this occurs, you must manually change the media type by using Windows PowerShell®.

Demonstration: Configuring Storage Spaces

In this demonstration, you will:

- Create a storage pool.
- Create a mirrored disk.

Demonstration Steps

Create a Storage Pool

1. On LON-SVR1, switch to Server Manager.
2. In the navigation pane, click **File and Storage Services**, and then in the Servers pane, click **Storage Pools**.
3. Create a storage pool with the following settings:
    - Name: **StoragePool1**
    - Select both physical disks for this storage pool
4. On the **View results** page, wait until the creation completes, and then click **Close**.

Create a Mirrored Disk

1. On LON-SVR1, in Server Manager, in the VIRTUAL DISKS pane, create a virtual disk with the following settings:
    - Storage pool: **StoragePool1**
    - Name: **Mirrored vDisk**
    - Storage Layout: **Mirror**
    - Provisioning type: **Thin**
    - Virtual disk size: **8**
2. On the **View results** page, wait until the creation completes, ensure that **Create a volume when this wizard closes** is selected, and then click **Close**.
3. In the New Volume Wizard, create a volume with the following settings:
    - Virtual disk: **Mirrored vDisk**
    - Drive letter: **F**
    - File system: **ReFS**
    - Volume label: **Mirrored Volume**
4. On the **Completion** page, wait until the creation completes, and then click **Close**.

Planning High Availability for Storage Spaces

Using hardware-based RAID arrays and mirrored or parity virtual disks in Storage Spaces provides a certain degree of fault tolerance and high availability to storage resources. However, because all physical disks connect to a single system, that system itself becomes a single point of failure. If the system to which the physical disks are connected fails, access to the storage resources ceases to exist.

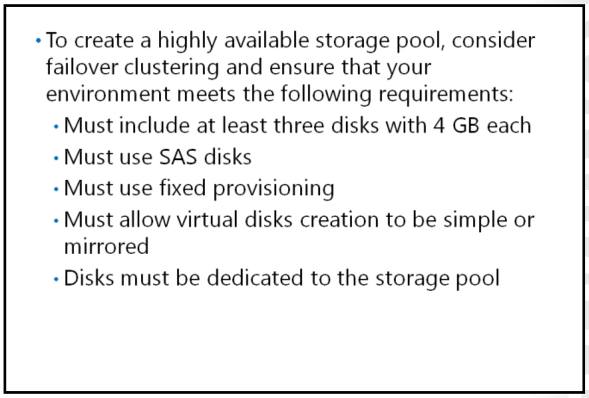

- To create a highly available storage pool, consider failover clustering and ensure that your environment meets the following requirements:
    - Must include at least three disks with 4 GB each
    - Must use SAS disks
    - Must use fixed provisioning
    - Must allow virtual disks creation to be simple or mirrored
    - Disks must be dedicated to the storage pool

In some scenarios, a single point of failure is acceptable. However, in critical systems, a single point of failure cannot be tolerated and they require a higher availability of system resources. To ensure your storage solution is highly available, you can create a storage space on a failover cluster in Windows Server 2012 and Windows Server 2012 R2.

To create a highly available storage pool by using failover clustering, your environment must meet the following requirements:

- At least three physical disks, with a minimum of 4 gigabytes (GB) of capacity on each physical disk.
- All physical disks must be SAS disks, and must be accessible to all nodes in the failover cluster.
- All physical disks must pass the Windows failover validation tests.
- All virtual disks must use fixed provisioning.
- Only simple and mirrored virtual disks can be created in the cluster.
- All physical disks must be dedicated to the storage pool.

If all the above conditions are met, you can create your clustered storage pool directly from the Failover Cluster Manager snap-in for Windows Server 2012.

Additional Reading: For more information on how to configure a clustered storage space in Windows Server 2012, visit http://go.microsoft.com/fwlink/?LinkID=391905.

Comparing iSCSI with Storage Spaces

Although some of the terminology and even feature sets are similar in both iSCSI and Storage Spaces, you can make good use of these technologies when you implement them together. To better understand how to use iSCSI and Storage Spaces together, you need to understand their differences. The following table compares the various aspects of iSCSI and Storage Spaces.

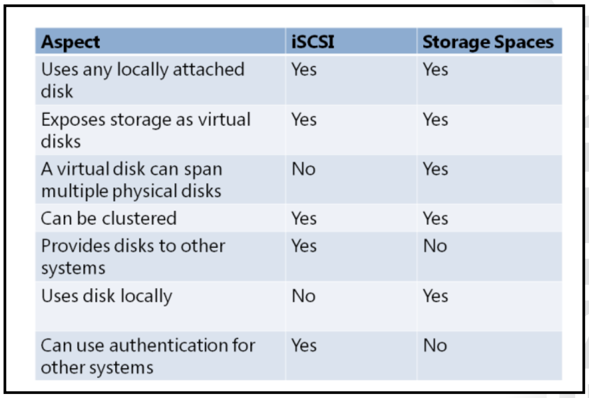

Aspect	iSCSI	Storage Spaces
Uses any locally attached disk	Yes	Yes
Exposes storage as virtual disks	Yes	Yes
A virtual disk can span multiple physical disks	No	Yes
Can be clustered	Yes	Yes
Provides disks to other systems	Yes	No
Uses disk locally	No	Yes
Can use authentication for other systems	Yes	No

Aspect	iSCSI	Storage Spaces
Uses any locally attached disk	Yes	Yes
Exposes storage as virtual disks	Yes	Yes
A virtual disk can span multiple physical disks	No	Yes
Can be clustered	Yes	Yes
Can provide disks to other systems on the network	Yes	No
Uses disks locally	No	Yes
Can use authentication to control access to disks	Yes	No

For example, consider that when you create a storage space it is accessible only to the system in which the storage space resides. You can create a shared folder and use the system as a file server. Alternatively, you can use iSCSI and expose a virtual disk to a server that then uses that disk to expose file shares. In either solution, you will not be able to fully utilize the combined feature set of Storage Spaces and iSCSI.

However, when you combine iSCSI and Storage Spaces, you can build a fully functional iSCSI SAN at a fraction of the cost of traditional SAN solutions. By using SAS disks that are shared across two or more Windows Server 2012 R2 nodes, you can create a highly available storage pool. Then, on the same cluster, you expose that pool by using iSCSI so that your iSCSI virtual disks will be on top of a storage space virtual disk that is fault tolerant and optimized for performance. By doing this, you achieve all the security features that only iSCSI provides: the ability to make virtual disks available to multiple servers by using iSCSI, and the fault tolerance and flexibility of Storage Spaces.

Planning Storage Optimization in Windows Server 2012 R2

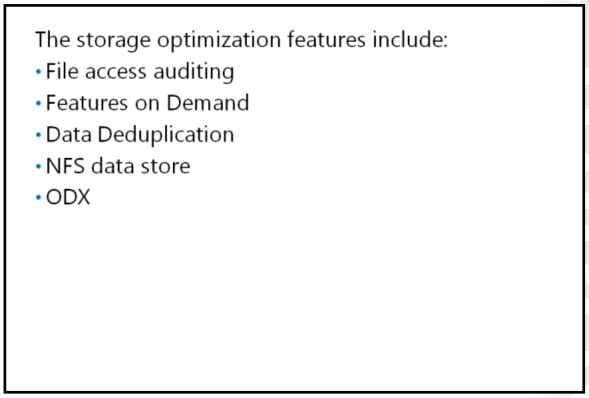

The storage optimization features include:

- File access auditing
- Features on Demand
- Data Deduplication
- NFS data store
- ODX

Windows Server 2012 includes options for storage optimization. These options provide you with an efficient way to deploy, administer, and secure your storage solutions. The storage optimization features include:

- File access auditing. File access auditing in Windows Server 2012 creates an audit event whenever users access files. As compared to previous Windows Server versions, this audit event data contains additional information about the attributes of the file that was accessed.
- Features on Demand. You can use the Features on Demand feature to save disk space. You can use this feature to remove role and feature files from the operating system. If you need to install these roles and features on the server, the installation files will be retrieved from remote locations, installation media, or the Windows Update website. You can remove feature files from both physical and virtual computers, Windows image (.wim) files, and offline virtual hard disks.
- Data Deduplication. Data Deduplication identifies and removes duplications within data, without compromising the integrity of the data. Data Deduplication is highly scalable, resource efficient, and nonintrusive. It can run concurrently on large volumes of primary data without affecting other workloads on the server. Low impact on server workloads is maintained by throttling the CPU and memory resources that are consumed. By using Data Deduplication jobs, you can schedule when Data Deduplication should run, specify the resources to deduplicate, and fine-tune file selection. When combined with Windows BranchCache®, you apply the same optimization techniques to data that is transferred over the WAN to a branch office. This results in faster file download times, and reduced bandwidth consumption.
- Network file system (NFS) data store. The NFS data store is the NFS server implementation in Windows Server 2012 operating systems. In Windows Server 2012, the NFS server supports high availability, which means that you can deploy the server in a failover clustering configuration. When a client connects to a NFS server in the failover cluster, and if that server fails, the NFS server will fail over to another node in the cluster. This enables the client to still connect to the NFS server.

- Offloaded Data Transfer (ODX). Windows ODX (also known as *copy offload*) allows for copying data on a storage array without moving the data from the storage array to the server, and then from the server back to the storage. This saves network bandwidth and ensures that data is copied much faster within the same storage array. Instead of moving files to the server, ODX moves a pointer to the files, and the server sends that back to the destination. The storage array itself executes the copy.

Additional Reading: For more information on Windows ODX, visit http://go.microsoft.com/fwlink/?LinkID=391906.

Lesson 3
Optimizing File Services for Branch Offices

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, branch offices must back up data that they maintain in their remote branch offices, which is why organizations prefer to centralize data where possible. Therefore, your challenge as an administrator is providing efficient access to network resources for users in branch offices.

Distributed File System (DFS) and BranchCache are features on Windows Server that you can use to optimize file services for smaller branch offices. Whether you are a Windows Server administrator with plenty of experience with DFS and BranchCache, or are just being introduced to the technology, you will learn to plan for the usage of these technologies in this lesson.

Lesson Objectives

After completing this lesson, you will be able to:

- Plan for DFS namespaces.
- Plan for DFS replication.
- Describe DFS storage scenarios.
- Describe DFS enhancements in Windows Server 2012 R2 and Windows Server 2012.
- Optimize file access by using BranchCache.
- Select the appropriate BranchCache mode for a solution.
- Configure BranchCache.
- Plan for implementing BranchCache.
- Compare and contrast the benefits and limitations of DFS with BranchCache.

Considerations for Planning a DFS Namespace

You can create DFS namespaces either as a domain-based or as a stand-alone namespace. The main difference between a domain-based namespace and a stand-alone namespace is in the format of the Universal Naming Convention (UNC) path that is used to access the share that the DFS namespace provides.

Domain-based namespaces use a UNC format that contains the name of the domain, for example, \\*adatum.com**files*. In this instance, the share name is *files*, and the domain name is *adatum.com*. One of the advantages of a domain-based namespace is that you are not linked directly to a particular server. This makes it easier to implement high availability for the share, even if a server that provides access to the DFS namespace is unavailable.

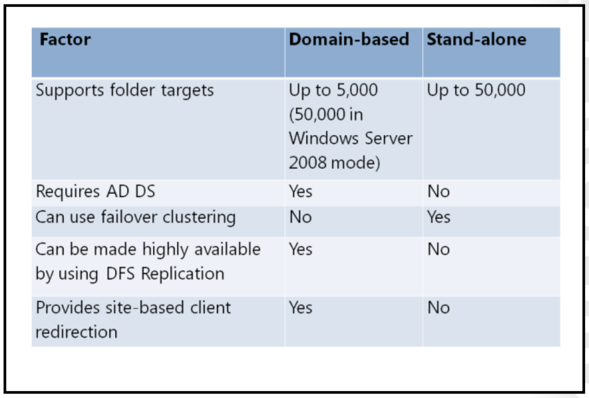

Factor	Domain-based	Stand-alone
Supports folder targets	Up to 5,000 (50,000 in Windows Server 2008 mode)	Up to 50,000
Requires AD DS	Yes	No
Can use failover clustering	No	Yes
Can be made highly available by using DFS Replication	Yes	No
Provides site-based client redirection	Yes	No

Stand-alone namespaces use a UNC format that includes the server name, for example, \\*server1**files*. In this instance, the share name is *files*, and it resides on a server named *server1*. Because the UNC uses the server name, if the server is unavailable there is no way of accessing the share by using this UNC, even if the contents of the share are replicated elsewhere.

When deciding whether to use a domain-based namespace or a stand-alone namespace, consider the factors listed in the following table.

Factor	Domain-based namespace	Stand-alone namespace
Supports folder targets	Up to 5,000 (50,000 in Windows Server 2008 mode)	Up to 50,000
Requires AD DS	Yes	No
Can use failover clustering	No	Yes
Can be made highly available by using DFS Replication	Yes	No
Provides site-based client redirection	Yes	No

Considerations for Planning DFS Replication

DFS Replication provides a way to keep folders synchronized between servers across well-connected and limited bandwidth connections. Keep in mind the following key points regarding DFS Replication:

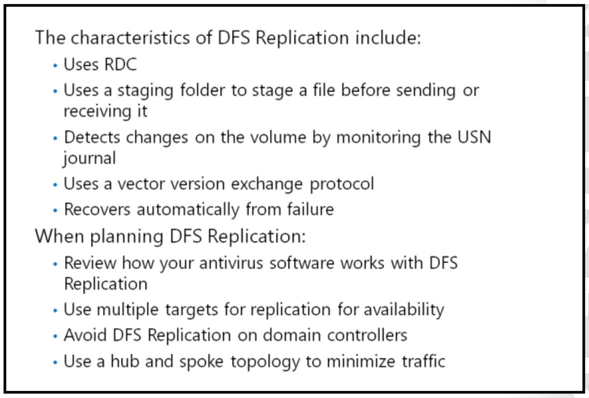

The characteristics of DFS Replication include:

- Uses RDC
- Uses a staging folder to stage a file before sending or receiving it
- Detects changes on the volume by monitoring the USN journal
- Uses a vector version exchange protocol
- Recovers automatically from failure

When planning DFS Replication:

- Review how your antivirus software works with DFS Replication
- Use multiple targets for replication for availability
- Avoid DFS Replication on domain controllers
- Use a hub and spoke topology to minimize traffic

- DFS Replication uses remote differential compression (RDC). RDC is a client-server protocol that can update files efficiently over a limited bandwidth network. RDC detects data insertions, removals, and rearrangements in files, thereby enabling DFS Replication to replicate only the changed file blocks when files are updated. RDC is only used for files that are 64 kilobytes (KB) or larger, by default.

  DFS Replication also supports cross-file RDC, which allows DFS Replication to use RDC, even when a file with the same name does not exist at the client. Cross-file RDC can determine files that are similar to the file that needs to be replicated. It also uses blocks of similar files that are identical to the replicating file to minimize the amount of data that needs to be replicated.

- DFS Replication uses a hidden staging folder to stage a file before sending or receiving it. Staging folders hold modified or new files while they wait to be replicated. When a request comes from another member of the DFS topology, the sending computer stages the file. DFS Replication reads the file from the replicated folder and creates a compressed version of the file in the staging folder. Then DFS Replication sends the compressed file to other computers in the DFS topology. If RDC is used, only a portion of the file is replicated. The receiving computer downloads the data and recreates the file in its staging folder. After replication completes, DFS decompresses the file and moves it into the replicated folder. Every replicated folder has a staging folder located under the local path of the replicated folder in the DfsrPrivate\Staging folder.

- DFS Replication detects changes on the volume by monitoring the file system update sequence number (USN) journal and replicates changes only after the file closes.
- DFS Replication uses a version vector exchange protocol to determine which files need to be synchronized. The protocol sends less than 1 KB per file across the network to synchronize the metadata associated with changed files on the sending and receiving members.
- DFS Replication uses a conflict resolution heuristic of "last writer wins" for files that are in conflict, and "earliest creator wins" for name conflicts. A file that is updated at multiple servers simultaneously is *in conflict*. Files and folders that lose the conflict resolution are moved to a folder known as the Conflict and Deleted folder. You also can configure the service to move deleted files to the Conflict and Deleted folder for retrieval, in case the file or folder should be deleted. Each replicated folder has its own hidden Conflict and Deleted folder, which is located under the local path of the replicated folder in the DfsrPrivate\ConflictandDeleted folder.
- DFS Replication can recover automatically from USN journal wraps, USN journal loss, or DFS Replication database loss.
- DFS Replication uses a Windows Management Instrumentation (WMI) provider that provides interfaces to obtain configuration and monitoring information from the DFS Replication service.

Consider the following key points regarding DFS Replication:

- DFS Replication may cause antivirus software to raise alerts based on replication behavior. Review the documentation for your antivirus software to learn how to configure it to work with DFS Replication.
- Use multiple targets for replication to provide availability of data. If a target is not working properly, clients can be redirected to another target server.
- Avoid using DFS Replication on domain controllers. Domain controllers already use DFS Replication for the SYSVOL folder. By isolating the use of DFS, troubleshooting becomes easier.
- Use a hub and spoke topology as much as possible to minimize replication traffic. When users change data in a spoke site, the data does not replicate to the hub. This allows you to manage replication better, with a single path to data replication.

DFS Data Storage Scenarios

Several key scenarios can benefit from DFS namespace and DFS Replication. Some of these scenarios include:

- Sharing files across branch offices
- Data collection
- Data distribution

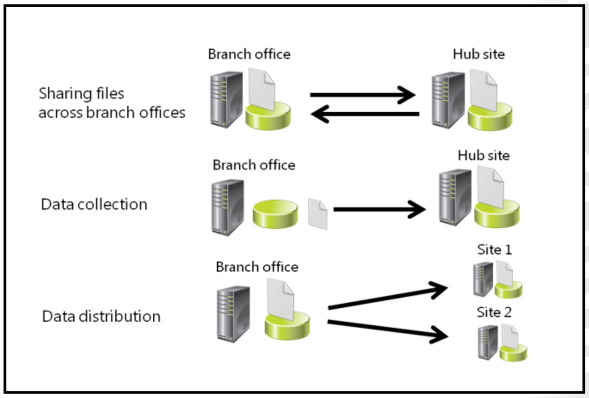

Sharing Files Across Branch Offices

Large organizations that have many branch offices often have to share files or collaborate between these locations. DFS Replication can help replicate files between branch offices or from a branch office to a hub site. Having files in multiple branch offices also benefits users who travel from one branch office to another. Using DFS Replication, the changes that users make to their files in one branch office are replicated back to their branch office.

We recommend this scenario only if users can tolerate some file inconsistencies as changes are replicated throughout the branch servers. This is because DFS Replication only replicates a file after it is closed. Therefore, DFS Replication is not a best practice for replicating database files or any files that are kept open for long periods of time.

In addition, DFS Replication is not suitable for files that multiple users may modify in different locations between replication cycles. This is because it uses a 'last write wins' scenario to determine which file version is authoritative. A user may update a file only to have their modification overwritten by another user's changes when replication occurs.

Data Collection

DFS technologies can collect files from a branch office and replicate them to a hub site. Using DFS Replication, you can replicate critical data to a hub site, and then back up the hub site using standard backup procedures. This increases the branch office data recoverability if a server fails, because files will be available in two separate locations and will be backed up. By using this method, companies can reduce branch office costs by eliminating backup hardware and onsite IT personnel expertise. Replicated data can also make branch office file shares fault tolerant. If the branch office server fails, clients in the branch office can access the replicated data at the hub site.

Data Distribution

You can use DFS namespaces and DFS Replication to publish and replicate documents, software, and other line-of-business (LOB) data throughout your organization. DFS namespaces and folder targets can increase data availability and distribute client load across various file servers.

Windows Server 2012 R2 and Windows Server 2012 Enhancements to DFS

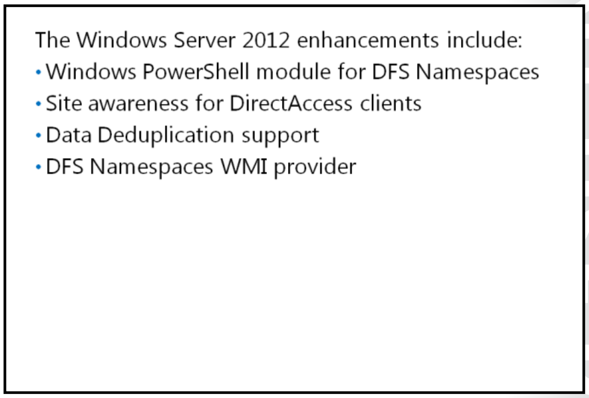

The Windows Server 2012 enhancements include:

- Windows PowerShell module for DFS Namespaces
- Site awareness for DirectAccess clients
- Data Deduplication support
- DFS Namespaces WMI provider

DFS namespaces and DFS Replication in Windows Server 2012 include the following new features:

- Windows PowerShell module for DFS Namespaces. Windows Server 2012 includes a module that you can use for managing most of the tasks related to DFS namespaces by using Windows PowerShell.
- Site awareness for DirectAccess clients. This feature allows remote computers to access the closest DFS namespace server when connected by using the Windows 7 or Windows 8 DirectAccess feature.
- Support for Data Deduplication. DFS Replication supports volumes that use Data Deduplication.
- Windows Management Infrastructure provider. You can use this feature to manage DFS namespaces by using WMI-capable tools.

Additional Reading: For more information on the enhancements to DFS Replication in Windows Server 2012, visit http://go.microsoft.com/fwlink/?LinkID=391907.

In addition to the new features in Windows Server 2012, Windows Server 2012 R2 includes the following enhancements to DFS Replication:

- Windows PowerShell module for DFS Replication. Windows Server 2012 R2 includes a module that you can use for managing most of the tasks related to DFS Replication by using Windows PowerShell.
- Windows Management Infrastructure provider. You can use this feature to manage DFS Replication by using WMI-capable tools.
- Database cloning for initial sync. This feature allows DFS to bypass initial replication when you create new replicated folders, replace servers, or recover from a disaster.
- Database corruption recovery. This feature allows DFS to rebuild corrupt databases without data loss that would be caused by an unauthoritative initial sync.
- Cross-file RDC disable. You can use this feature to disable cross-file RDC between servers.
- File staging tuning. You can this feature to configure variable file staging sizes on individual servers.
- Preserved file restoration. This feature allows DFS to recover automatically after a loss of power or DFS Replication service stoppage.
- Membership disabling improvements. This feature stops DFS Replication private folder cleanup when a server's membership is disabled in a replication folder.

Additional Reading: For more information on the enhancements to DFS Replication in Windows Server 2012 R2, visit http://go.microsoft.com/fwlink/?LinkID=391908.

Optimizing File Access by Using BranchCache

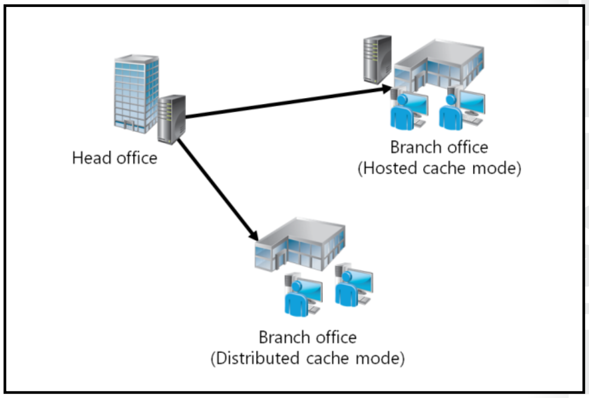

The BranchCache feature that was introduced with Windows Server 2008 R2 and Windows 7 reduces the network use on WAN connections between branch offices and headquarters by caching frequently used files locally on computers in the branch office. BranchCache improves the performance of applications that 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). This is 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 the 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 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. BranchCache functionality in Windows Server 2012 R2 includes the following improvements:

- Scalability. To allow for scalability, BranchCache allows for more than one hosted cache server per location.
- Ability to store more data. 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).
- Simpler deployment. This 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.

How Client Computers Retrieve Data by Using BranchCache

When BranchCache is enabled on both the client computer and the 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 connects to a content server in the head office that is running Windows Server 2012, and requests the content in the same way that it would when not 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 normally been sent over.
4. By 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.

Selecting a BranchCache Mode

You can configure BranchCache to use either hosted cache mode or distributed cache mode:

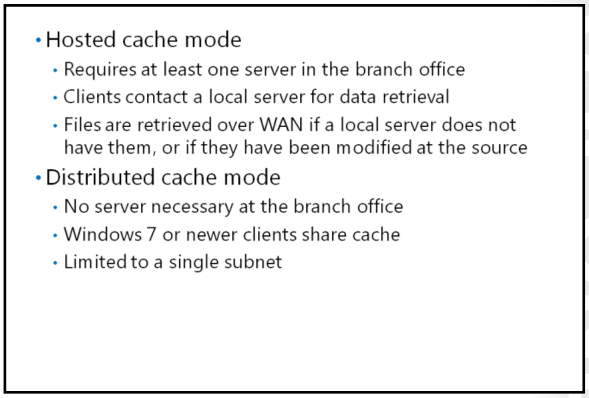

- Hosted cache mode
  - Requires at least one server in the branch office
  - Clients contact a local server for data retrieval
  - Files are retrieved over WAN if a local server does not have them, or if they have been modified at the source
- Distributed cache mode
  - No server necessary at the branch office
  - Windows 7 or newer clients share cache
  - Limited to a single subnet

- Hosted cache mode. This mode operates by deploying a server that is running Windows Server 2008 R2 or newer 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 the content from the hosted cache.
- 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 that are running Windows 8.1 or Windows 7 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.

Demonstration: Configuring BranchCache

In this demonstration, you will see how to:

- Add BranchCache for the Network Files role service.
- Configure BranchCache in the 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 Server Manager.
2. In the Add Roles and Features Wizard, install the following roles and features to the local server:
   - On the **Select server roles** page, navigate to **File and Storage Services (2 of 12 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. In the Local Group Policy Editor, expand **Computer Configuration\Administrative Templates\Network\Lanman Server**, and perform the following steps:
   - 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 **File Explorer**, and on drive C, create a folder named **Share**.
2. Configure the **Share** folder properties as follows:
   - Enable **Share this folder**.
   - Select **Enable BranchCache** in **Offline Settings**.

Considerations for Implementing BranchCache

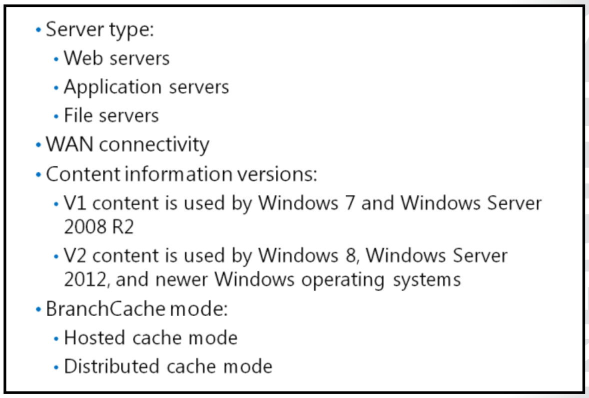

- Server type:
  - Web servers
  - Application servers
  - File servers
- WAN connectivity
- Content information versions:
  - V1 content is used by Windows 7 and Windows Server 2008 R2
  - V2 content is used by Windows 8, Windows Server 2012, and newer Windows operating systems
- BranchCache mode:
  - Hosted cache mode
  - Distributed cache mode

Consider the following factors when deploying BranchCache:

- BranchCache content servers. BranchCache was designed to enable the following content servers to share resources across WAN links:
  - Web servers. You can install the BranchCache feature on a computer running Internet Information Services (IIS) on the main office.
  - Application servers. Windows Server Update Services (WSUS) and Configuration Manager branch distribution points also can use BranchCache to share content across WAN links.
  - File servers. File servers that are running Windows Server 2008 R2 and newer are the main servers that benefit from BranchCache. They require the BranchCache for the Network Files role service.
- WAN connectivity. Remember that WAN connectivity is required for BranchCache. The hosted cache server, or client in case of distributed cache, still requires access to the server that has BranchCache enabled to verify the hash for the file being accessed. BranchCache is not a high availability technology; it is used for file distribution in high latency or expensive WAN links.
- Content information versions. BranchCache uses two types of content information:
  - V1 BranchCache content. This type of content is used by computers running Windows 7 and Windows Server 2008 R2. V1 BranchCache content uses a larger fixed size for file segments. Changes to files invalidate the segment changed, along with the end segment for the file, causing more data to be sent over the WAN.
  - V2 BranchCache content. This type of content is used by computer running Windows 8, Windows Server 2012, and newer Windows operating systems. V2 BranchCache content uses smaller, variable size segments that are more tolerant to changes in the file, causing less WAN traffic.
- BranchCache mode. You need to choose the BranchCache content information type based on the available systems on a branch network. Hosted cached is preferable, but requires a server. With Windows Server 2012 R2, you can have multiple host cache servers for availability.

Comparing DFS and BranchCache

You can use both DFS and BranchCache to make content available to branch offices while reducing the traffic over a WAN link. The main difference between the two technologies is related to how WAN is used to replicate data. In DFS Replication, all content from a folder replicates over to the branch office, whereas in BranchCache, content replicates only when a client requests it. Therefore, BranchCache uses the WAN more efficiently, and thus is recommended more for smaller branch offices that access data less frequently.

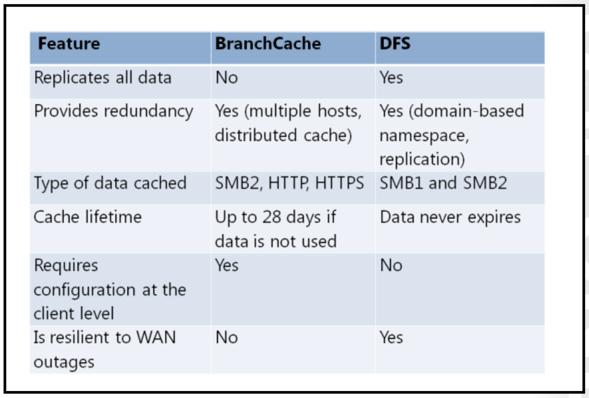

Feature	BranchCache	DFS
Replicates all data	No	Yes
Provides redundancy	Yes (multiple hosts, distributed cache)	Yes (domain-based namespace, replication)
Type of data cached	SMB2, HTTP, HTTPS	SMB1 and SMB2
Cache lifetime	Up to 28 days if data is not used	Data never expires
Requires configuration at the client level	Yes	No
Is resilient to WAN outages	No	Yes

The following table lists the main differences between BranchCache and DFS.

Feature	BranchCache	DFS
Replicates all data	No	Yes
Provides redundancy	Yes (multiple hosts, distributed cache)	Yes (domain-based namespace, replication)
Type of data cached	SMB2, HTTP, HTTPS	SMB1 and SMB2
Cache lifetime	Up to 28 days if data is not used	Data never expires
Requires configuration at the client level	Yes	No
Is resilient to WAN outages	No	Yes

Lab: Planning and Implementing Storage

Scenario

For the most part, servers in A. Datum Corporation are configured solely with DAS. As the project to migrate to Windows Server 2012 R2 gathers pace, it is important to consider alternatives to this storage technology.

You must select a storage technology that will help to improve the performance of disk-intensive applications, and provide for improved storage management.

A number of applications throughout the A. Datum organization are mission critical, and it is important that the selected storage technology provide suitable fault tolerance to enable high availability of these applications. These applications are based on SQL Server.

Objectives

After completing this lab, you will be able to:

- Plan a storage solution.
- Implement iSCSI storage.
- Configure a redundant storage space with storage pools.

Lab Setup

Estimated Time: 50 minutes

Virtual machines	20413C-LON-DC1 20413C-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, in the **Start** screen, click **Hyper-V Manager**.
2. In Hyper-V Manager, click **20413C-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in using the following credentials:
   - User name: **Administrator**
   - Password: **Pa$$w0rd**
   - Domain: **Adatum**
5. Repeat steps 2 through 4 for 20413C-LON-SVR1.

Exercise 1: Planning a Storage Solution

Scenario

SQL Server serves as the foundation for a sales and customer management application that has been migrated over from a UNIX platform within A. Datum. The entire sales force in A. Datum is now being trained on how to use this application.

The platform that supports this application has been virtualized and configured for high availability. It is important that the storage used by this application does not become a point of failure.

Supplemental documentation

Email from Dan Park:

<u>Brad Sutton</u>

From:	Dan Park [Dan@Adatum.com]
Sent:	01 Nov 13:12
To:	Brad@Adatum.com
Subject:	Sales application storage

Brad,

Our entire sales force in A. Datum is now being trained on the use of our application. Originally, this application ran on UNIX at Contoso in Paris, but it has been ported across to Windows Server 2012 R2.

I'm told that the platform that supports the sales application has been virtualized and configured for high availability. It is important that the storage used by this application does not become a point of failure. Could you look into this for us, and then report back?

Thanks,

Dan

Sales Application Storage Strategy	
Document Reference Number: BS1102/1	
Document Author Date	Brad Sutton 2nd November
Requirements Overview Plan a storage strategy to support the following objectives: • High availability. Storage must not become a point of failure. • Performance. Many users are connected to the application. • Cost. Cost of the storage solution should not be prohibitively expensive.	
Additional Information • Originally, this application ran on UNIX at Contoso in Paris, but it has been ported across to Windows Server 2012 R2.	

Sales Application Storage Strategy

Proposals

1. How will you configure storage?

2. What type of storage is indicated?

3. How will you try to ensure that the storage is made highly available?

4. How could Storage Spaces help address the requirements?

The main tasks for this exercise are as follows:

1. Read the supporting documentation.

2. Update the proposal document with your planned course of action.

3. Examine the suggested proposals in the Lab Answer Key.

4. Discuss your proposed solution with the class, as guided by your instructor.

▶ Task 1: Read the supporting documentation

- Read the documentation provided.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the Sales Application Storage Strategy document.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones in the Lab Answer Key.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you will have planned storage for the sales application.

Exercise 2: Implementing iSCSI Storage

Scenario

A. Datum has decided to implement virtual machine storage in the head office and at the two regional hubs by using an iSCSI storage deployment. You must configure the iSCSI targets and iSCSI initiators to support this deployment.

The main tasks for this exercise are as follows:

1. Install the iSCSI target.

2. Configure iSCSI targets.

3. Connect to and configure iSCSI targets.

▶ Task 1: Install the iSCSI target

1. Sign in to LON-DC1 with the user name **Adatum\Administrator** and the password **Pa$$w0rd**.
2. In Server Manager, click **Add roles and features**.
3. In the Add Roles and Features Wizard, browse to **File And Storage Services (Installed)\File and iSCSI Services**, and install the **iSCSI Target Server**. Accept the default values.
4. If prompted, restart LON-DC1, and sign back in.

▶ Task 2: Configure iSCSI targets

1. On LON-DC1, in Server Manager, in the navigation pane, click **File and Storage Services**, and then click **iSCSI**.
2. Create a virtual disk with the following settings:
   - Storage location: **C:**
   - Disk name: **iSCSIDisk1**
   - Size: **5 GB**
   - iSCSI target: **New**
   - Target name: **LON-DC1**
   - Access servers: **172.16.0.11**
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**

▶ Task 3: Connect to and configure iSCSI targets

1. Sign in to LON-SVR1 with the user name **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-DC1-target**
4. In Server Manager, on the **Tools** menu, open **Computer Management**.
5. 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.
6. Close the Computer Management console.

Results: After completing this exercise, you will have successfully implemented an iSCSI SAN.

Exercise 3: Configuring a Redundant Storage Space

Scenario

After you have configured the iSCSI components, you want to use storage pools to simplify storage configuration on the Windows Server 2012 R2 servers. To meet the requirements for high availability, you decide to evaluate redundancy features in Storage Spaces.

The main tasks for this exercise are as follows:

1. Create a storage pool by using the iSCSI disks that are attached to the server.

2. Create a mirrored disk.

3. Copy a file and verify that it displays.

4. Disconnect an iSCSI disk and verify that the file is still accessible.

▶ Task 1: Create a storage pool by using the iSCSI disks that are attached to the server

1. On LON-SVR1, switch to Server Manager.
2. In the navigation pane, click **File and Storage Services**, and then in the Servers pane, click **Storage Pools**.
3. Create a storage pool with the following settings:
   - Name: **StoragePool1**
   - Select both physical disks for this storage pool
4. On the **View results** page, wait until the creation completes, and then click **Close**.

▶ Task 2: Create a mirrored disk

1. On LON-SVR1, in Server Manager, in the VIRTUAL DISKS pane, create a virtual disk with the following settings:
   - Storage pool: **StoragePool1**
   - Name: **Mirrored vDisk**
   - Storage Layout: **Mirror**
   - Provisioning type: **Thin**
   - Virtual disk size: **8 GB**
2. On the **View results** page, wait until the creation completes, ensure that **Create a volume when this wizard closes** is selected, and then click **Close**.
3. In the New Volume Wizard, create a volume with the following:
   - Virtual disk: **Mirrored vDisk**
   - Drive letter: **F**
   - File system: **ReFS**
   - Volume label: **Mirrored Volume**
4. On the **Completion** page, wait until the creation completes, and then click **Close**.

▶ Task 3: Copy a file and verify that it displays

1. On the Start screen, type **command prompt**, and then press Enter.

2. Type the following command, and then press Enter:

```
Copy C:\windows\system32\write.exe F:\
```

3. Open File Explorer, and access **Mirrored Volume (F:)**. You should now see write.exe in the file list.

▶ Task 4: Disconnect an iSCSI disk and verify that the file is still accessible

1. Switch to LON-DC1.
2. In the iSCSI VIRTUAL DISKS pane, in the LON-DC1 list, disable the iSCSI virtual disk named **iSCSIDisk1.vhd**.
3. Switch to LON-SVR1.
4. Open File Explorer, and then open **F:\write.exe** to ensure that access to the volume is still available.
5. In Server Manager, in the STORAGE POOLS pane, on the menu bar, click the **Refresh "Storage Pools"** button. Notice the warning that displays next to Mirrored vDisk.
6. In the VIRTUAL DISK pane, right-click **Mirrored vDisk**, and then in the drop-down list box, click **Properties**.
7. In the Mirrored vDisk Properties window, in the Health pane, notice that the Health Status indicates a Warning. The Operational Status should indicate Degraded.

Results: After completing this exercise, you will have configured a redundant storage space.

▶ Task: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-SVR1**, and then click **Revert**.
3. In the **Revert Virtual Machines** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20413C-LON-DC1**.

Module Review and Takeaways

Review Question(s)

Question: Tailspin Toys needs to decide how to implement various aspects of its storage infrastructure. The company will need to store shared files in a central location, but they do not want to implement a file server at this time. What kind of storage would you recommend?

Question: Tailspin Toys plans to implement several database servers, and wants to provide disk space for the databases. The company would prefer to create a single, centrally-managed array of disks for all the databases. What kind of storage would you recommend?

Question: What are the primary benefits of a SAN storage solution over a DAS storage solution?

Module 11

Designing and Implementing Network Protection

Contents:

Module Overview

A network security design has several important aspects to it, including a consistent process for identifying threats, and monitoring and maintaining security. Security models—such as the defense-in-depth model—provide consistent frameworks for identifying network threats. If you use one of the models to identify potential security risks, you can then analyze the risks to determine how you will mitigate them.

Windows Server® 2012 and Windows Server 2012 R2 provide features and applications that you can use to help secure your network against anticipated security threats. These include Windows® Firewall, Internet Protocol security (IPsec), and Network Access Protection. You must carefully consider the security threats posed to elements of your network, and then design the implementation of these features and applications to counter these perceived threats.

Objectives

After completing this module, you will be able to:

- Describe the design process for network security.
- Explain how to identify and mitigate network security threats.
- Design and implement a Windows Firewall strategy.
- Design and implement Network Access Protection (NAP).

Lesson 1
Overview of Network Security Design

You can use threat modeling to help you identify and mitigate common network threats that may make your organization susceptible to attack. Having identified common network security threats, you should examine the relative value of your assets, and then allocate your security resources based on both the likelihood of the risk occurring, and the value of the asset. Risk analysis helps you to prioritize your efforts and spending to secure your network.

Lesson Objectives

After completing this lesson, you will be able to:

- Discuss common network vulnerabilities.
- Describe the key network security principles.
- Explain the defense-in-depth model.
- Describe how to protect against common network vulnerabilities.
- Design network perimeter security.
- Describe the process for designing network security.
- Describe the importance of network security policies and procedures.
- Implement best practices for creating a risk management plan.

What Network Threats Do Organizations Face?

You have been asked by the IT manager at Wingtip Toys to provide some advice about what network threats they might face. You have decided to write a report, listing the ten most common network threats and possible mitigations or solutions to counter these various network threats.

Question: What are the 10 most common network security threats faced by organizations?

Question: What possible mitigations or solutions exist to counter these threats?

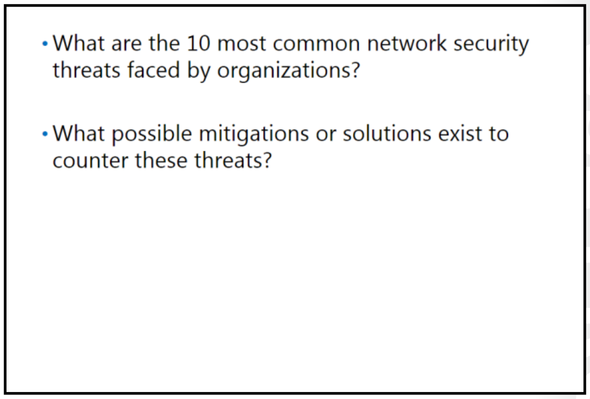

- What are the 10 most common network security threats faced by organizations?
- What possible mitigations or solutions exist to counter these threats?

Overview of Network Attacks

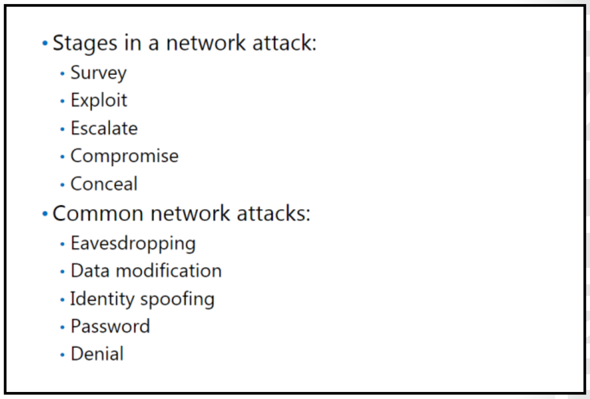

- Stages in a network attack:
  - Survey
  - Exploit
  - Escalate
  - Compromise
  - Conceal
- Common network attacks:
  - Eavesdropping
  - Data modification
  - Identity spoofing
  - Password
  - Denial

Before developing a plan to mitigate the effects of network attacks, you must understand why they occur, the stages within a network attack, and the common causes of network attacks.

Why Do Network Attacks Occur?

Some of the reasons why network attacks occur include the following:

- Profit. A common motivation for attacks is profit. Hackers may use the threat of a denial-of-service attack to extort money or implement phishing attacks to footprint your network.
- Revenge. Hackers may feel slighted by an organization and want to punish it. For example, former employees may attack their previous organizations if they consider that their employment was terminated unfairly. These hackers are particularly dangerous, because they have an in-depth knowledge of the network, and a personal motivation for attack.
- Espionage. Hackers may spy on an organization or government to obtain secrets. Such an attack is often motivated by patriotism or monetary gain.
- Publicity. Hackers may attack a network or application to seek public notoriety or to advertise their own services. Publicity-seekers often advertise their attacks.
- Personal satisfaction. Hackers may attack networks as a hobby, for the challenge, for their own amusement, or as a boost to their egos. Such hackers are dangerous because they attack networks indiscriminately.
- Political or philosophical reasons. Insiders may collect and distribute confidential data because they have political or philosophical disagreements with the organization. These types of malicious users are sometimes termed “Hacktivists.”
- Terrorism. Hackers may attack a network as part of a terrorist effort that has been sponsored by a group or state. These are the most serious types of attack because human life may be at risk.

Stages in a Network Attack

By understanding the basic steps that hackers use to target your network, you are better equipped to take defensive measures. The following are the basic steps that hackers use to target your network:

1. Survey. First, hackers usually survey the potential target to identify and assess its characteristics. These characteristics may include its supported services and protocols together with potential vulnerabilities and entry points.
2. Penetrate. After surveying a potential target, the next step is to exploit and penetrate. Hackers look for known vulnerabilities based on the list of network resources that they have gathered during survey and assessment.
3. Escalate. After compromising a network, hackers immediately attempt to escalate their privileges by accessing administrative and system accounts.
4. Compromise. After gaining access to a system, hackers take steps to make future access easier, such as planting back-door programs, using an existing account that lacks strong protection, or creating a new account.
5. Conceal. Attackers then cover their tracks by clearing logs and hiding tools.

Common Network Attacks

You should be able to identify common network attacks so that you can design security mitigations. Some common types of network attacks include the following.

- Eavesdropping. Attackers eavesdrop by using a network sniffer to capture network communication. All clear-text data is at risk. However, packet sniffing is relatively difficult on a switched network, or where physical access has been restricted.
- Data modification. Attackers perform data modification after they have captured data with a network sniffer. Most network sniffers support modifying packets and replaying them. Again, packet sniffing can be difficult to implement.
- Identity spoofing. Attackers can use identity spoofing to fool some firewalls into thinking that communication is coming from an internal source rather than an external source. They accomplish this by falsifying the source IP address.
- Password. Password-based attacks rely on users with simple passwords. After guessing the password of a user, hackers can view network resources that are accessible by that user. Using least-privilege can mitigate this type of attack because standard users should have very limited access to sensitive resources.
- Denial of service. Denial-of-service attacks prevent authorized users from accessing network services. Most denial-of-service attacks exploit software flaws. By keeping your network software up-to-date, you can help to avoid some of these types of attacks.
- Man-in-the-middle. Man-in-the-middle attacks require a computer to monitor and potentially modify network communication between two hosts.
- Compromised key. A compromised key occurs when a key that is used for encryption is known to anyone other than legitimate parties to the communication. Knowledge of the key enables unauthorized parties to view the contents of encrypted communication. This also includes unauthorized knowledge of private keys for certificates used during authentication.
- Application layer. Application layer attacks cause faults in either an operating system or an application to bypass normal access controls. A common application layer attack is a buffer overflow.
- Social networking and social engineering. The prevalence and widespread use of social networking sites, and the ongoing nuisance of unsolicited email also pose threats. Many of these types of attacks are difficult to defeat, as they rely on the inherent trust of your network users. To counter these threats, educate your users to help to ensure that they can identify potential phishing activities or other social threats.

Key Principles of Network Security

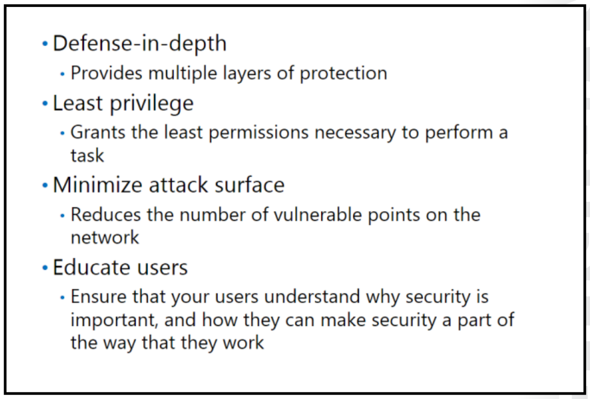

- Defense-in-depth
  - Provides multiple layers of protection
- Least privilege
  - Grants the least permissions necessary to perform a task
- Minimize attack surface
  - Reduces the number of vulnerable points on the network
- Educate users
  - Ensure that your users understand why security is important, and how they can make security a part of the way that they work

Users must have access to network resources to which they have been authorized. In addition, the network requires a secure, shared IT infrastructure. To attain these goals, apply the following four principles:

- Implement a defense-in-depth approach to security. Defense-in-depth refers to a combination of people, operations, and security technologies. Implementing a defense-in-depth approach provides multiple layers of protection to your network by defending against threats at multiple points in your network.

- Implement least privilege. Least privilege refers to granting a user, resource, or application the least amount of privilege or permissions necessary to perform their required task. Granting excessive permissions can introduce numerous vulnerabilities that hackers can potentially exploit. For example, never grant a user Full Control permissions on a folder when Modify permissions on specific files within that folder would be sufficient for the user to complete their required operational tasks.
- Minimize the attack surface. This reduces the number of possible points of entry for a hacker by removing unnecessary software, services, and devices. For example, by implementing Windows Server 2012 or Windows Server 2012 R2 role-based installation, you deploy only the server roles that your organization requires, and no more.
- Educate your users. Users must understand why certain security features are important. You may have implemented many security features to help to protect your network, but if users do not implement them (such as not posting or giving out their passwords), then your network is no longer secure.

You should apply these four principles to all aspects of your network security design.

What Is the Defense-in-Depth Model?

When you begin considering security issues, it is best to take a holistic approach. For example, when you park your car in a public place, you consider a number of factors before walking away from it: where it is parked, whether the doors are locked, and whether you have you left anything valuable lying in view. You understand the risks associated with parking in a public place, and you attempt to mitigate those risks. As with your car, you cannot properly implement security features on a computer network without first understanding the security risks posed to that network.

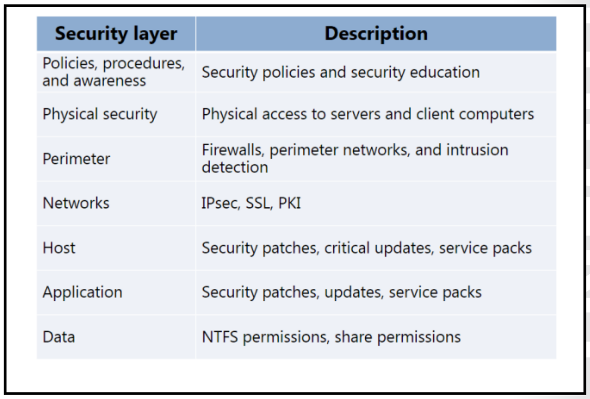

Security layer	Description
Policies, procedures, and awareness	Security policies and security education
Physical security	Physical access to servers and client computers
Perimeter	Firewalls, perimeter networks, and intrusion detection
Networks	IPsec, SSL, PKI
Host	Security patches, critical updates, service packs
Application	Security patches, updates, service packs
Data	NTFS permissions, share permissions

You can mitigate risks to your computer network by considering the risks, and then providing security at differing infrastructure layers. The term *defense-in-depth* describes the use of multiple security technologies at different points throughout your organization.

Policies, Procedures, and Awareness

You must ensure that your users adhere to organizational security policies and practice physical security measures. Enforcing a strong user password policy is not helpful if users write their passwords down and attach them to their computer screens, or share them with others. Also, make users aware of security risks posed to their computers if they visit certain types of websites. This is why organizations typically have an Acceptable Use Policy.

Physical Security

If an unauthorized person can gain physical access to your computers, then most other security measures are of little consequence. You must ensure that computers containing the most sensitive data—such as servers—are physically secure.

In addition, do not make it easy to connect to your network. If someone can plug a laptop into your network and access your intranet, you could face serious problems. Physical security includes securing the network infrastructure. The problem, of course, is that while you want to make it difficult for non-authorized people to access your computers and infrastructure, you want to make it relatively straightforward for authorized employees.

Perimeter

These days, no organization works in isolation. Organizations operate within a global community, and network resources must be available to service that global community. This might include building a website to describe your organization's services, or making internal services such as web conferencing and email accessible externally so that users can work from home or from satellite offices.

To keep your organization safe, you should create both a private network and a perimeter network with firewalls, intruder prevention, detection systems, and other components. Perimeter networks mark the boundary between public and private networks. By deploying servers in the perimeter network—such as those that publish websites or make email mailboxes available to external users—you can provide corporate services more securely across the public network.

Networks

Once you connect computers to a network, they become susceptible to any number of threats. These threats include eavesdropping, spoofing, denial of service, and replay attacks. This is especially relevant when communication takes place over public networks by users who are working from home or from remote offices. You can use a variety of technologies to help mitigate these threats.

Note: You can place highly sensitive data on servers that connect to an isolated network. However, this approach is not always appropriate for much of the data that you store on your network.

At the center of many of these risks is authentication. If two computers can identify one another, then they can communicate more securely. You can provide authentication services in a number of ways, such as digital certificates that are exchanged during initial communications. How you distribute and manage these certificates depends on your organization, but it might include implementing a public key infrastructure (PKI).

In addition to authentication, consider using encryption to ensure that data is secure while it is in transit. You can encrypt communication to the perimeter-based servers (otherwise known as *edge servers*) from the public network with tunneling technologies, and you can encrypt communication between the edge servers and the internal network with IPsec.

In addition, Secure Sockets Layer (SSL), which is widely used on the Internet, can provide for secure and authenticated communications across networks.

Note: Consider implementing network intruder detection within your network. This will help you identify inappropriate network access, potentially before data is compromised or service is disrupted.

Host

The next layer of defense is that used for the host computer. You must keep these computers secure with the latest security updates and patches. Windows Update and Windows Server Update Services (WSUS) can help to keep your Windows 8, Windows 8.1, Windows Server 2012, and Windows Server 2012 R2 computers up-to-date. In addition, consider using a host-based firewall, such as Windows Firewall, and implementing both antivirus and malware protection.

Application

Applications are only as secure as your latest security update. You should consistently use Windows Update to keep your applications up to date.

If your organization relies upon applications that are designed and written in-house, it is important to design these applications to be secure. This involves implementing secure authentication, and supporting other security features. For example, if the application requires the use of a service account, be sure to configure the account with the minimal permissions required to run the application. In addition, design the application to support encrypted network traffic.

Data

The final layer of security is data security. This might include using NTFS file system permissions and shared folder permissions to ensure that only authorized users can access files at a defined level of access. You might also be concerned with intellectual property rights and ensuring that your data is used appropriately, perhaps by using AD RMS. Finally, for data privacy, you can use both file and disk encryption technologies, such as the Encrypting File System (EFS) or BitLocker.

Mitigating Common Network Vulnerabilities

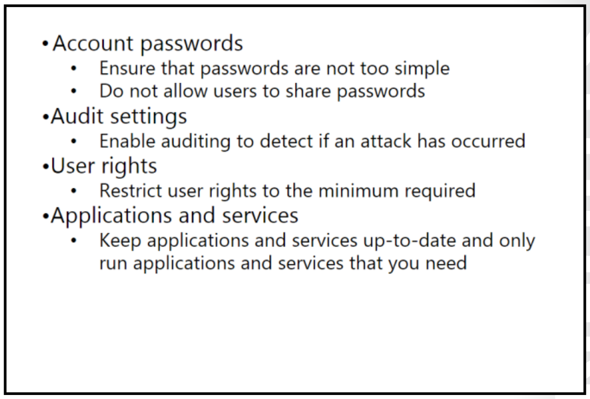

- Account passwords
  - Ensure that passwords are not too simple
  - Do not allow users to share passwords
- Audit settings
  - Enable auditing to detect if an attack has occurred
- User rights
  - Restrict user rights to the minimum required
- Applications and services
  - Keep applications and services up-to-date and only run applications and services that you need

Most successful networks attacks succeed by exploiting common and well-known vulnerabilities or weaknesses. These vulnerabilities and weaknesses can be organized into the following general categories:

- Weak passwords and authentication systems. Weak passwords and weak authentication systems enable hackers to gain access to your system by using brute-force password attacks. You can increase security by requiring complex passwords, or by introducing two-factor authentication such as smart cards. However, you must balance increased security for passwords with ease of use.
- Not tracking network access attempts. You should use audit logs to monitor which users accessed network resources, and when. If you do not have audit logs enabled or configured to collect the appropriate information, then it may be impossible to detect hackers accessing your network.
- Not implementing least privilege. The rights and permissions granted to any users should be the minimum required to perform their job. In this way, if a hacker compromises an account, the damage they can cause is minimized. This is also true for service accounts.
- Any service or application is a potential point of attack. To minimize the risk of applications and services being attacked, you should remove all unnecessary applications and services. In addition, you should apply applications and services security updates regularly.

Guidelines for Modeling and Countering Vulnerabilities

A threat model is a structured approach to predicting potential threats to information security. By discovering potential threats while performing threat modeling, you can create an accurate risk management plan. By predicting threats, you can proactively reduce your risk.

Best Practices

The following guidelines will assist you when modeling threats to your network:

- Encourage creative thinking among team members. Some suggestions, however unrealistic, may prompt others to discover other valid threats.
- Ensure that team members have all the information that they require, such as network diagrams or application source code.
- Manage discussions so that they focus on the validity of a threat to the network, and avoid disagreements about minor differences of opinion.
- When assembling your team, consider including a trusted external party who specializes in network security testing. The external party has skills that are likely not available internally, and brings a different perspective.
- Use caution when including team members who may have conflicts of interest. For example, a developer who wrote the code in the application being assessed or a manager who funded the project to create the application may overestimate the ability of the application to withstand an attack, or may be too familiar with it to be objective about its assessment.

Designing Network Perimeter Security

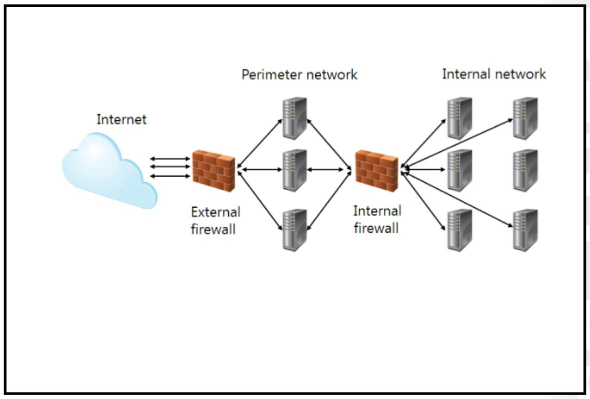

The perimeter network is the network between an external and an internal firewall. No traffic can pass directly from the Internet to the protected internal network, and no traffic can pass directly from the protected internal network to hosts on the Internet. Instead, all traffic must traverse a host on the perimeter network. Hosts on the perimeter network should not be joined to an Active Directory Domain Services (AD DS) domain. You should configure the external firewall so that traffic can only pass to hosts on the perimeter network using specific ports. For example, you should configure incoming traffic on Transmission Control Protocol (TCP) port 80 to route to a web server on the perimeter network.

Similarly, you should configure the internal firewall so that traffic traversing the internal firewall can only pass if it uses specific ports. For example, traffic on TCP port 80 from the internal network to the web server on the perimeter network should be allowed, but traffic on port 80 from the web server on the perimeter network to the internal network need not be.

The network perimeter design model has the following benefits.

- In the event that a host on the perimeter network (such as a web server) is compromised by a hacker, a firewall still blocks the hacker from accessing hosts on the internal network.
- It allows specific services to be made available to the Internet in a protected manner without exposing hosts on the internal network.
- It blocks direct communication between hosts on the Internet and hosts on the internal network, and blocks direct communication from hosts on the internal network and hosts on the Internet. This makes it very difficult for a hacker to access hosts on the internal network due to the traffic flow being restricted.

You typically deploy the following server roles on perimeter networks:

- External web server. The external web server should only contain content that the organization needs to make available to the public. Ensure that only information that should be available to the public should be published on this server and that sensitive information is not published to the server. Intranet servers should be hosted on trusted internal networks.
- Web Proxy Server. Clients on the internal network use this server to access web-related content on the Internet. It will also store cached copies of commonly accessed sites. This server may be configured to check web content for malware, and may also be used to block certain sites and content from being accessed by clients on the internal network.
- Federation server proxy. If the organization uses Active Directory Federation Services (AD FS), a federation server proxy may be deployed on the perimeter network to allow external parties to utilize AD FS.
- Exchange Edge Transport server or Simple Mail Transfer Protocol (SMTP) relay. This server routes mail traffic in to and out of the organization. This server may be configured to block unsolicited commercial email, filter malware, and filter sensitive messages being sent by people on the internal network.
- DNS forwarder. This server forwards Domain Name System (DNS) requests from DNS servers on the internal network to DNS servers on the Internet.

The Process of Designing Security

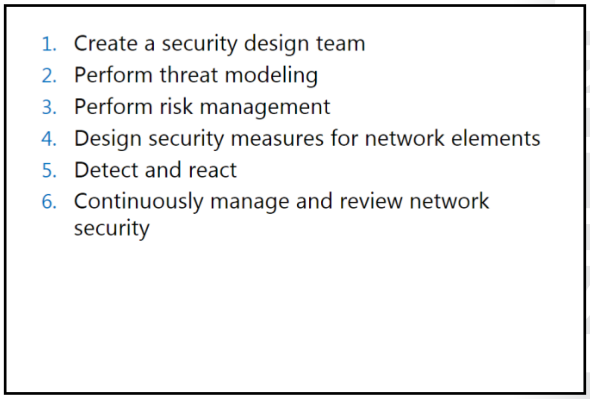

1. Create a security design team
2. Perform threat modeling
3. Perform risk management
4. Design security measures for network elements
5. Detect and react
6. Continuously manage and review network security

The stages in the network security design process are as follows:

1. Create a security design team. Solicit multiple perspectives for your security design, so that you can plan to mitigate as many vulnerabilities and threats as possible. Always keep in mind that widely consulting for information encourages acceptance of the finished plan.
2. Perform threat modeling. Threat modeling predicts threats to a given asset or resource. Knowing the threats that could potentially affect an asset helps you to design countermeasures to protect the asset.
3. Perform risk management. This analyzes the likelihood of a threat occurring, and the potential damage that a threat may cause. Risk management is a valuable tool that can help you convince management that security measures are necessary to defend a resource adequately against a threat.
4. Design security measures for your network elements. Create appropriate policies and procedures to protect your network based on the threat modeling and risk management that you performed.
5. Detect and react. Identify ways to detect intrusions and respond to security incidents in a controlled manner. Detecting an attack early is vital to limiting the damage that the attack may cause. A careful and thoughtful response to the attack can make recovery easier, and can prevent mistakes that may make the situation worse.
6. Manage and review network security on a continual basis. Create, implement, and review policies for acceptable use, network management, and secure network operation.

Network Security Policies and Procedures

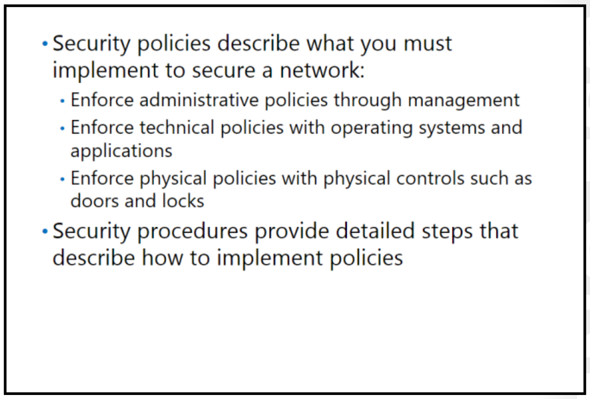

- Security policies describe what you must implement to secure a network:
  - Enforce administrative policies through management
  - Enforce technical policies with operating systems and applications
  - Enforce physical policies with physical controls such as doors and locks
- Security procedures provide detailed steps that describe how to implement policies

Security policies are individual policies and guidelines that you create to govern the secure and appropriate use of technology and processes within your organization. The following are a few types of security policies that you may already have in place:

- Administrative policies. Administrative policies are enforced by management. These policies cannot be enforced by operating systems, applications, or physical controls. An example is a nondisclosure agreement.
- Technical policies. Technical policies, such as security templates or a set of security-related configuration settings, are enforced by operating systems and applications.
- Physical policies. Physical policies, such as locks, are enforced by implementing physical controls.

Security procedures describe how to comply with security policies. Your security procedures should include the detailed steps necessary to implement your security policies. Part of applying security policies involves addressing network access and utilization policies.

Reasons for Security Policy Failure

Security policies fail for many reasons. If you anticipate these reasons, you can ensure that your policies address these concerns.

Security policies most often fail because they are:

- Not enforced. Employees tend to disregard security policies if the policies are not enforced and violators are not disciplined.
- Difficult to read. Security policies often contain legal or technical language that makes them difficult for employees to understand.
- Difficult to find. Security policies that are stored in obscure or inaccessible locations prevent many employees from following them.
- Outdated. Security policies that are not kept up-to-date quickly become obsolete when technologies and business processes change.
- Too vague. Security policies that are open to interpretation by employees often result in inconsistent deployment of security.
- Too strict. Security policies that are too strict in their enforcement or in their effect on business processes are generally not taken seriously by employees, or enforced by management.
- Not supported by management. If management does not support security policies or their implementation, employees typically do not follow them either.

Guidelines for Creating Policies and Procedures

Guidelines for creating security policies and procedures include the following:

- Ensure that your security policies serve a clear purpose and are written concisely.
- Write simple procedures and policies that demonstrate how to comply with the policies successfully.
- Obtain management support for the purpose, implementation, and enforcement of security policies.

- Distribute your security policies so that employees can refer to them easily. For example, give paper copies of policies to employees, or post the policies to convenient internal websites, and update the policies regularly.
- Before implementing security policies, ensure that they do not disrupt or hinder business processes.
- Use technology to enforce security policies. This helps prevent employees from unwittingly violating security policies. However, remember that technology is not the only method of enforcement.
- Ensure that the consequences of violating security policy are consistent with the severity of the violation and with the culture of your organization. Ensure that managers are empowered to enforce the consequences of violating security policy.

Risk Assessment and Impact

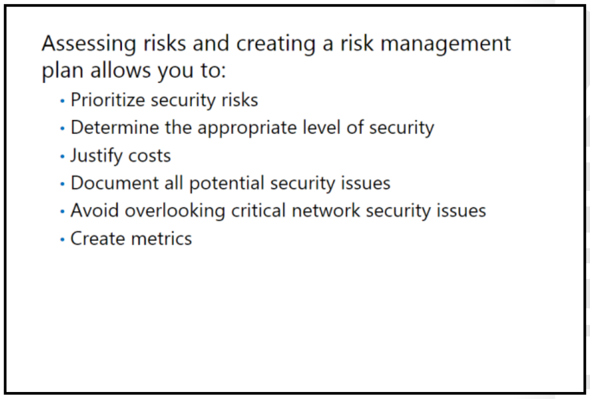

Assessing risks and creating a risk management plan allows you to:

- Prioritize security risks
- Determine the appropriate level of security
- Justify costs
- Document all potential security issues
- Avoid overlooking critical network security issues
- Create metrics

Risk assessment helps to ensure that your security plan is rational and that you apply your resources to maximize results. By assessing risks and creating a risk management plan, you can:

- Rank the security risks to your organization relative to other risks. This helps your organization determine how to allocate resources to secure the network.
- Discover the point at which incremental improvements to security become inefficient and costly.
- Use a quantitative risk analysis to justify the expense of security personnel, hardware, and software.
- Create a comprehensive list of threats and their potential impact to your network. This is necessary to allocate resources properly for network security.
- Ensure that all important threats are identified. An organization that chooses to respond to security threats randomly may overlook critical security issues on its network.
- Create metrics that help you to judge the success of your security plan. You can also use metrics to prepare compensation plans for executives and security personnel.

Note: Ensure that your risk management plan includes both proactive and reactive elements. In other words, proactively plan to mitigate security threats, and have plans ready to reactively deal with security problems should they occur.

What Is at Risk?

When assessing risk, start with a comprehensive list of assets that are at risk of attack. You can then analyze the various risks for each asset. Some categories of risk are:

- Hardware. This includes desktop and portable computers, routers and switches, storage devices, and backup media.
- Software. This includes software installation media, operating system images, custom software applications and code, and virtualized servers.
- Documentation. This includes security policies, security procedures, and network diagrams and building plans.
- Data. This includes trade secrets, employee confidential information, and customer information.

A large part of the security role is protecting public confidence and the trust of business partners. This is known as *goodwill*. While goodwill is difficult to quantify monetarily, losing goodwill can be costly for an organization.

For example, suppose an hacker defaces your organization's website and accesses confidential information. You notify customers that the hacker has stolen the private information of the website's users, including their addresses and credit card numbers. In addition to incurring direct financial losses from lost business, your organization also suffers a loss of goodwill because the company's image becomes tarnished.

Lesson 2
Designing and Implementing a Windows Firewall Strategy

When planning network security, you must consider security requirements for your perimeter network, and the threat of attacks from the Internet. You must also consider how to apply security on internal networks. A common way to protect internal networks is by using Windows Firewall.

The Windows Firewall feature affords protection by providing both inbound and outbound firewall rules, and by authenticating and optionally encrypting network traffic through connection security rules. Windows Firewall is a host-based firewall for both Windows Server and Windows client operating systems. Windows Firewall uses connection security rules to implement IPsec.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe scenarios addressed by Windows Firewall.
- Describe the benefits and recommended usage for IPsec.
- Describe how to use connection security rules to authenticate and encrypt network traffic.
- Select appropriate authentication options and methods.
- Implement best practices when designing network security rules.
- Describe how to configure connection security rules.

Scenarios Addressed by Windows Firewall

Windows Firewall is a host-based firewall for Windows Server and Windows client operating systems. Windows Firewall includes support for different network types, outbound rules, and connection security rules. As part of your security design, you must consider what rules you should implement, and how you will deploy them to hosts on your network.

Question: What scenarios can Windows Firewall can help to address?

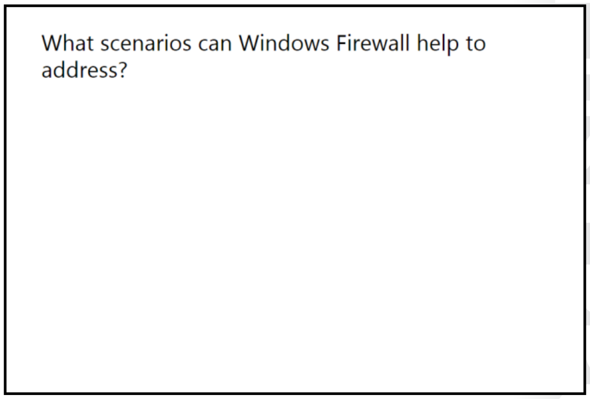

What scenarios can Windows Firewall help to address?

Configuring Inbound and Outbound Rules

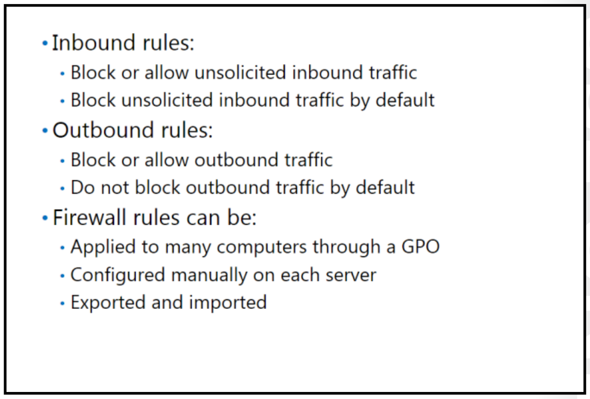

- Inbound rules:
  - Block or allow unsolicited inbound traffic
  - Block unsolicited inbound traffic by default
- Outbound rules:
  - Block or allow outbound traffic
  - Do not block outbound traffic by default
- Firewall rules can be:
  - Applied to many computers through a GPO
  - Configured manually on each server
  - Exported and imported

Rules are criteria that define what inbound or outbound traffic you will allow or block through the Windows Firewall.

Inbound Rules

Inbound rules explicitly allow or block traffic that matches the rules' criteria. For example, you can configure a rule to allow traffic that is secured by IPsec for Remote Desktop through the firewall, but block the same traffic if it is not secured by IPsec. You would need to separately configure an IPsec rule to secure the traffic.

When you first deploy Windows Server 2012, Windows Firewall blocks all unsolicited inbound traffic. Unsolicited traffic is traffic sent to the computer without there being a request sent from the computer. To allow a particular type of unsolicited inbound traffic, you must create an inbound rule that describes that traffic. For example, if you want to run a web server, you must create a rule that allows unsolicited inbound network traffic on TCP port 80. You can configure the default action that Windows Firewall with Advanced Security takes, for example, whether to allow or block connections when no inbound rule applies.

Outbound Rules

Windows Firewall allows all outbound traffic unless a rule blocks it. Outbound rules explicitly allow or deny traffic originating from a computer that matches a rule's criteria. For example, you can configure a rule to explicitly block outbound traffic from a computer by IP address through the firewall, but allow the same traffic for other computers.

Rule Types

There are four different types of inbound and outbound firewall rules:

- Program rules. These rules control connections for a program. Use this type of firewall rule to allow a connection based on the program that is trying to connect. These rules are useful when you are not sure of the port or other required settings, because you only specify the path to the program's executable (.exe) file.
- Port rules. These rules control connections for a TCP or User Datagram Protocol (UDP) port. Use this type of firewall rule to allow a connection based on the TCP or UDP port number over which the computer is trying to connect. You specify the protocol and the individual or multiple local ports to which the rule applies.
- Predefined rules. These rules control connections for a Windows-based experience. Use this type of firewall rule to allow a connection by selecting one of the programs or experiences from the list. Network-aware programs that you install typically add their own entries to this list so that you can enable and disable them as a group.
- Custom rules. Configure these rules as necessary. Use this type of firewall rule to allow a connection based on criteria that the other three types of firewall rules do not encompass.

Consider the scenario in which you want to create and manage tasks on a remote computer by using the Task Scheduler user interface. Before connecting to the remote computer, you must enable the Remote Scheduled Tasks Management firewall exception on the remote computer. You can do this by using the predefined rule type on an inbound rule. Alternatively, you might want to block all web traffic on the default TCP web server port 80. In this scenario, you create an outbound port rule that blocks the specified port.

Managing Inbound and Outbound Rules

Although you can configure firewall rules manually on an individual server, this becomes cumbersome when you need to deploy the same rule configuration to a large number of computers. You have the following options to apply the same set of rules to multiple computers:

- Apply rules using Group Policy. You can configure a Group Policy Object (GPO) to include firewall rules, and then apply that GPO to the computers on which you want to have the firewall rules applied. This is suitable when the target servers are members of an Active Directory domain.
- Export and import rules from one computer to another. You can configure a set of inbound and outbound rules on one computer, and then export those rules. You can then import the exported rules on other computers, thereby overwriting the current set of rules with the imported set of rules. This is a suitable strategy when the target computers are not members of an Active Directory domain.

IPsec Benefits and Usage

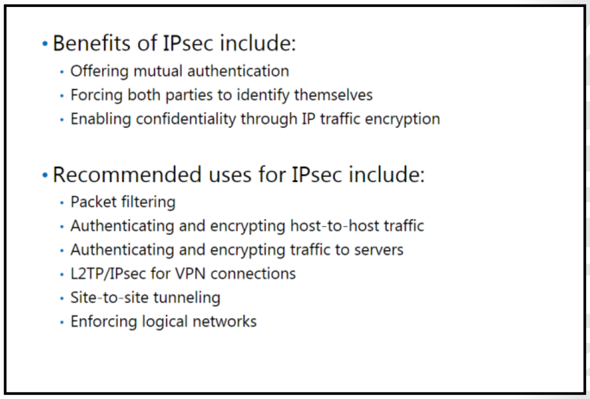

- Benefits of IPsec include:
  - Offering mutual authentication
  - Forcing both parties to identify themselves
  - Enabling confidentiality through IP traffic encryption
- Recommended uses for IPsec include:
  - Packet filtering
  - Authenticating and encrypting host-to-host traffic
  - Authenticating and encrypting traffic to servers
  - L2TP/IPsec for VPN connections
  - Site-to-site tunneling
  - Enforcing logical networks

IPsec is a framework of open standards for protecting communications over IP networks through cryptographic security services. IPsec supports:

- Network-level peer authentication
- Data-origin authentication
- Data integrity
- Data confidentiality through encryption
- Replay protection

IPsec is also a suite of protocols that can help protect data in transit through a network by using security services, and optionally by using digital certificates with public and private keys. Because of its design, IPsec helps provide much better security than previous protection methods. Network administrators who use IPsec may not need to configure security for individual programs.

IPsec Benefits

IPsec provides a private channel for sending and exchanging potentially sensitive or vulnerable data, such as partner and supply-chain data, personnel or medical records, or any other type of TCP/IP-based data. You can use IPsec to help attain confidentiality, integrity, and authentication while transporting data across insecure channels such as email, File Transfer Protocol (FTP) traffic, or news feeds. Though its original purpose was to secure traffic across public networks, many organizations have chosen to implement IPsec to address perceived weaknesses in their own private networks that might be susceptible to exploitation.

IPsec has these benefits:

- Offers mutual authentication before and during communications.
- Forces both parties to identify themselves during the communication process.
- Enables confidentiality through IP traffic encryption and digital packet authentication.

Recommended IPsec Usage

Some network environments are well suited to IPsec as a security solution, but others are not. We recommend IPsec for the following uses:

- Packet filtering. IPsec provides limited firewall capabilities for end systems. You can permit or block inbound or outbound traffic by using IPsec with the network address translation (NAT)/Basic Firewall component of the Routing and Remote Access Service.
- Securing host-to-host traffic on specific paths. You can use IPsec to provide protection for traffic between servers, other static IP addresses, or subnets. For example, IPsec can secure traffic between domain controllers in different sites, or between web servers and database servers.
- Securing traffic to servers. You can require IPsec protection for all client computers that access a server. Additionally, you can set restrictions on which computers can connect to a server that is running Windows Server 2012.
- Layer Two Tunneling Protocol (L2TP)/IPsec for virtual private network (VPN) connections. You can use the combination of the L2TP and IPsec (L2TP/IPsec) for all VPN scenarios. This does not require that you configure and deploy IPsec policies.
- Site-to-site (gateway-to-gateway) tunneling. You can use IPsec in tunnel mode for site-to-site (gateway-to-gateway) tunnels when you need interoperability with routers, gateways, or end systems that do not support L2TP/IPsec or Point-to-Point Tunneling Protocol (PPTP) connections.
- Enforcing logical networks (server/domain isolation). In a Windows Server–based network, you can logically isolate server and domain resources to limit access to authenticated and authorized computers. For example, you can create a logical network inside the existing physical network, where computers share common requirements for secure communications. To establish connectivity, each computer in this logically isolated network must provide authentication credentials to other computers.

  This isolation prevents unauthorized computers and programs from gaining inappropriate access to resources. Requests from computers that are not part of the isolated network are ignored. Server and domain isolation can help protect specific high-value servers and data, and protect managed computers from unmanaged or unauthorized computers and users.
  You can protect a network with two types of isolation:

  - Server isolation. To isolate a server, you configure specific servers to require IPsec policy when accepting authenticated communications from other computers. For example, you might configure the database server to accept connections only from the web application server.
  - Domain isolation. To isolate a domain, you use Active Directory domain membership to ensure that computers belonging to a domain accept only authenticated and secured communications from other computers within that same domain. The isolated network consists only of that domain's member computers. Domain isolation uses IPsec policy to protect traffic sent between domain members, including all client and server computers.

Note: Because IPsec depends on IP addresses for establishing secure connections, you cannot specify dynamic IP addresses. It is often necessary for a server to have a static IP address in IPsec policy filters. In large network deployments, and in some mobile user cases, using dynamic IP addresses at both ends of the connection can increase the complexity of the IPsec policy design.

Non-Recommended IPsec Usage

IPsec can reduce processing performance and increase network bandwidth consumption. Additionally, IPsec policies can be quite complex to configure and manage, and using IPsec can introduce application compatibility issues. For these reasons, avoid using IPsec in the following situations:

- Securing communication between domain members and their domain controllers. Using IPsec in this situation reduces network performance. In addition, we do not recommend using IPsec for secure communication between domain members and their domain controllers, because the required IPsec policy configuration and management is complex.
- Securing all network traffic. Using IPsec to secure all network traffic reduces network performance and introduces the following possible issues:
  - IPsec cannot negotiate security for multicast and broadcast traffic.
  - Traffic from real-time communications, from applications that require Internet Control Message Protocol (ICMP), and from peer-to-peer applications might be incompatible with IPsec.
  - Network management functions that must inspect the TCP, UDP, and protocol headers either are less effective or cannot function at all, due to IPsec encapsulation or IP payload encryption.

Additionally, the IPsec protocol and implementation have characteristics that require special consideration when you:

- Protect traffic over wireless 802.11 networks. You can use IPsec transport mode to protect traffic sent over 802.11 networks. However, do not use IPsec for providing security for corporate 802.11 wireless local area networks (LANs). Instead, use 802.11 Wi-Fi Protected Access (WPA) or WPA2 encryption and Institute of Electrical and Electronics Engineers, Inc. (IEEE) 802.1X authentication.

Support for IPsec, configuration management, and trust are required on client computers and servers. Because many computers on a network do not support IPsec, or they are not managed, it is not appropriate to use IPsec by itself to protect all 802.11 corporate wireless LAN traffic.

- Consider that IPsec tunnel mode policies are not optimized for mobile clients with dynamic IP addresses, nor does IPsec tunnel mode support dynamic address assignment or user authentication, which is required for remote-access VPN scenarios. Use L2TP/IPsec VPN connections to secure remote access traffic to organizational networks when that traffic is sent over public wireless networks that are connected to the Internet.
- Use IPsec in tunnel mode for remote access VPN connections. For Windows-based VPN clients and servers, we do not recommend using IPsec in tunnel mode for remote access VPN scenarios. Instead, use L2TP/IPsec or PPTP.

Connection Security Rules

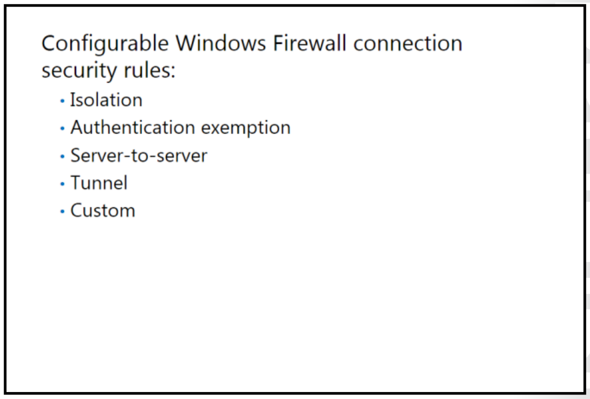

Configurable Windows Firewall connection security rules:

- Isolation
- Authentication exemption
- Server-to-server
- Tunnel
- Custom

Firewall rules allow or disallow traffic through the firewall, but do not secure that traffic. To secure traffic, you can create connection security rules. Windows Firewall with Advanced Security uses these rules to evaluate network traffic, and then blocks or allows messages based on the criteria that you establish in the rule.

In some circumstances, Windows Firewall with Advanced Security will block communication. If you configure settings that require security for a connection (in either direction), and if the two computers cannot authenticate each other, the connection is blocked. Windows Firewall with Advanced Security uses IPsec to enforce these rules.

The configurable Windows Firewall rules are as follows:

- Isolation rule. An *isolation rule* isolates computers by restricting connections based on credentials such as domain membership or health status. Isolation rules allow you to implement an isolation strategy for servers or domains.
- Authentication exemption. You can use an authentication exemption rule to designate connections that do not require authentication. You can designate computers by specific IP address, an IP address range, a subnet, or a predefined group such as a gateway.
- Server-to-server rule. A *server-to-server rule* protects connections between specific computers. This type of rule usually protects connections between servers. When you create the rule, you specify the network endpoints between which communications are protected. You then designate requirements and the authentication that you want to use.
- Tunnel rule. A tunnel rule enables you to protect connections between gateway computers. You typically use tunnel rules when connecting across the Internet between two security gateways. You must specify the tunnel endpoints by IP address, and then specify the authentication method.
- Custom. Use a custom rule to authenticate connections between two endpoints when you cannot set up the authentication rules that you need by using the other rules that are available in the new Connection Security Rule Wizard.

Authentication Options and Methods

Authentication Options

While using the Connection Security Rule Wizard to create a new rule, you can use the Authentication Requirements Wizard to specify how authentication is applied to inbound and outbound connections. If you request authentication, communications are enabled even if authentication fails. If you require authentication, the connection drops if authentication fails.

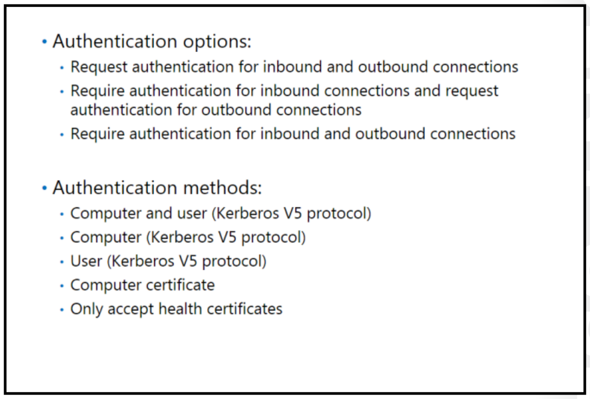

- Authentication options:
  - Request authentication for inbound and outbound connections
  - Require authentication for inbound connections and request authentication for outbound connections
  - Require authentication for inbound and outbound connections
- Authentication methods:
  - Computer and user (Kerberos V5 protocol)
  - Computer (Kerberos V5 protocol)
  - User (Kerberos V5 protocol)
  - Computer certificate
  - Only accept health certificates

The following list are options available for configuring authentication:

- Request Authentication for Inbound and Outbound Connections. Select this option to specify that all inbound and outbound traffic is authenticated, yet allow the connection even if authentication fails. If authentication succeeds, then traffic is protected. You typically use this option either in low-security environments, or in an environment where computers must connect but cannot perform the types of authentication that is available with Windows Firewall with Advanced Security.
- Require Authentication for Inbound Connections, and Request Authentication for Outbound Connections. Select this option if you require that all inbound traffic is authenticated, or else it is blocked. Outbound traffic can be authenticated, but is allowed even if authentication fails. If authentication succeeds for outbound traffic, that traffic is authenticated. You typically use this option in most IT environments in which computers that must connect can perform the authentication types that are available with Windows Firewall with Advanced Security.
- Require Authentication for Inbound and Outbound Connections. Select this option if you want to require that all inbound and outbound traffic either is authenticated, or it is blocked. You typically

use this option in higher-security IT environments where you must protect and control traffic flow, and in which the computers that must connect can perform the authentication types that are available with Windows Firewall with Advanced Security.

Authentication Methods

The following authentication methods are available:

- Computer and User (Kerberos version 5 (V5) protocol). This method uses both computer and user authentication, which means that you can request or require both the user and the computer to authenticate before communications continue. You can use the Kerberos V5 authentication protocol only if both computers and users are domain members.
- Computer (Kerberos V5 protocol). This method requests or requires the computer to authenticate by using the Kerberos V5 authentication protocol. You can use the Kerberos V5 authentication protocol only if both computers are domain members.
- User (Kerberos V5 protocol). This method requests or requires the user to authenticate by using the Kerberos V5 authentication protocol. You can use the Kerberos V5 authentication protocol only if the user is a domain member.
- Computer Certificate. This method requests or requires a valid computer certificate to authenticate, and you must have at least one certification authority (CA) to do this. Use this method if the computers are not part of the same Active Directory domain.

Note: You can use Group Policy to distribute certificates throughout your organization. This simplifies certificate administration for connection security rules.

- Accept only health certificates. The method requests or requires a valid health certificate to authenticate. Health certificates declare that a computer has met system health requirements (such as all software and other updates that network access requires), as determined by a NAP health policy server. These certificates are distributed during the NAP health evaluation process. Use this method only for supporting NAP.

Note: You can also specify advanced authentication methods, in which you determine which methods to use, and the order in which they should be attempted.

Best Practices for Designing Connection Security Rules

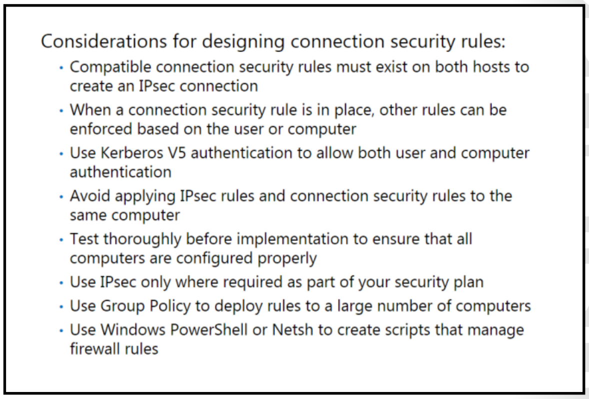

Considerations for designing connection security rules:

- Compatible connection security rules must exist on both hosts to create an IPsec connection
- When a connection security rule is in place, other rules can be enforced based on the user or computer
- Use Kerberos V5 authentication to allow both user and computer authentication
- Avoid applying IPsec rules and connection security rules to the same computer
- Test thoroughly before implementation to ensure that all computers are configured properly
- Use IPsec only where required as part of your security plan
- Use Group Policy to deploy rules to a large number of computers
- Use Windows PowerShell or Netsh to create scripts that manage firewall rules

Some of the best practices for designing connection security rules are as follows:

- Compatible connection security rules must exist on both hosts to create an IPsec connection. For example, authentication must be configured in the same way on both hosts.
- When a connection security rule is in place, other rules can be enforced based on the user or computer. This enables increased flexibility to restrict access to some applications by user or computer rather than by IP address. This avoids problems with changing IP addresses due to dynamic IP addressing.

- Use Kerberos V5 authentication to allow both user and computer authentication. The Kerberos V5 protocol is based on domain authentication, and requires no additional configuration. However, this authentication method is only suitable for computers that are members of the domain.

Note: Using Kerberos V5 authentication for connection security rules that involve domain controllers can result in the domain becoming inoperable.

- Avoid applying IPsec policies and connection security rules to the same computer. You can apply IPsec policies and connection security rules simultaneously, but we do not recommend this because the two can conflict. When a conflict occurs, it is difficult to determine where the problem is occurring.

Note: IPsec policies that are applied through connection security rules override IPsec policies that are applied through Group Policy. The IPsec policies in Group Policy are only provided for backwards compatibility for clients that do not support connection security rules, such as the Windows XP or Windows Server 2003 operating systems.

- Test the security rules thoroughly before implementation to ensure that all computers are configured properly. The best practice is to request IPsec authentication, and then verify functionality before requiring IPsec authentication.
- Use IPsec only where required as part of your security plan. Using IPsec increases network complexity, and you should not implement it without a defined purpose.
- Use Group Policy to deploy rules to a large number of computers. This is an automated process, and therefore, is less prone to error. Any new computers added to an OU have the rules applied automatically. Rules that you deploy using Group Policy override conflicting rules that are created on a local server.
- Use the Windows PowerShell® command-line interface or the Netsh command-line tool to create scripts that manage firewall rules. Scripts enable you to configure individual computers in a repeatable way that eliminates errors that are potentially introduced when using Windows Firewall with Advanced Security. In most cases, you use scripts only when it is difficult to configure an appropriate GPO.

Demonstration: Configuring Connection Security Rules

This demonstration shows how to:

- Enable ICMP traffic on LON-SVR1.
- Create a server-to-server rule on connecting servers.
- Create a server–to-server rule on LON-CL1.
- Test the rule.

Demonstration Steps

Enable ICMP traffic on LON-SVR1

- On LON-SVR1, open Windows Firewall with Advanced Security, and create a new inbound firewall rule to allow all ICMPv4 traffic, if it is secure.

Create a server-to-server rule on connecting servers

1. On LON-SVR1, in Windows Firewall with Advanced Security, create a new connection security rule with the following settings:

   a. Type: **Server-to-server**

   b. **Require authentication for inbound and outbound connections**

   c. Authentication: **Preshared key**

   d. Name: **Adatum-Server-to-Server**

Create a server-to-server rule on LON-CL1

1. Switch to LON-CL1 and open Windows Firewall with Advanced Security.

2. Create a new connection security rule with the following settings:

   a. Type: **Server-to-server**

   b. **Require authentication for inbound and outbound connections**

   c. Authentication: **Preshared key**

   d. Name: **Adatum-Server-to-Server**

Test the Rule

1. On LON-CL1, open a Windows PowerShell prompt and use the ping command to verify communications with LON-SVR1.

2. In Windows Firewall with Advanced Security, use the Monitoring node to verify secure communications.

Lesson 3
Designing and Implementing a NAP Infrastructure

First introduced in Windows Server 2008, NAP is a Windows Server operating system feature that prevents non-compliant computers from accessing a network. When you design NAP, you must ensure that the design meets your organization's requirements. As for other technologies, you must gather business requirements from the various departments.

NAP policies are used to define compliant and non-compliant clients. The available system health agents and system health validators control the options available for defining compliant and non-compliant clients.

By designing your NAP infrastructure carefully, and in such a way as to address both your organization's business requirements and your network security requirements, the implementation becomes transparent to users with computers that are in compliance.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe system health agents and system health validators.
- Define the properties of a NAP policy.
- Explain how to select appropriate system health validator settings.
- Determine how to manage unsupported client operating systems.
- Define a network policy to support your NAP implementation.
- Select a NAP enforcement method.
- Select a suitable remediation model.
- Describe how to implement NAP.

Systems Health Agents and System Health Validators

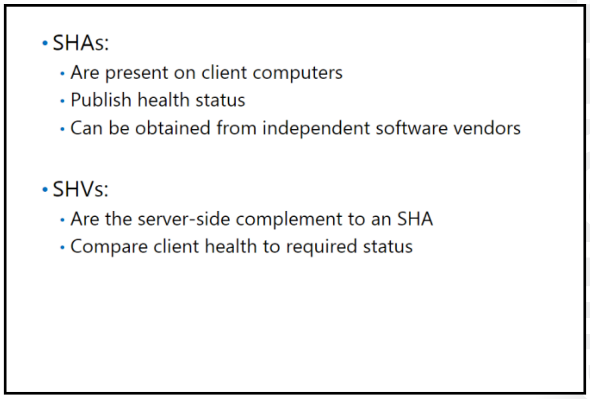

- SHAs:
  - Are present on client computers
  - Publish health status
  - Can be obtained from independent software vendors
- SHVs:
  - Are the server-side complement to an SHA
  - Compare client health to required status

A system health agent (SHA) is responsible for publishing the health status to an enforcement point. An SHA, which is present on NAP clients, is included with the following Windows operating systems:

- Windows XP Service Pack 3 (SP3)
- Windows Vista®
- Windows 7
- Windows 8
- Windows 8.1
- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2012 R2

Independent software providers also may provide an SHA for monitoring the status of their products.

A system health validator (SHV), which is present on a server that is running Network Policy Server (NPS), is responsible for comparing client health to the required health status. Each SHA on the client side must have a corresponding SHV on the server side.

An SHV is included with Windows Server 2012. It is capable of monitoring health for the following Windows operating systems:

- Windows XP SP3
- Windows Vista
- Windows 7
- Windows 8
- Windows 8.1
- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2012 R2

The settings for Windows 8 and Windows 8.1 also apply for Windows Server 2012 and Windows Server 2012 R2. The following are the settings that you can monitor for Windows 8 and Windows 8.1:

- Firewall is enabled.
- Antivirus application is on and up to date.
- Antispyware application is on and up to date.
- Automatic updating is enabled.
- All available security updates are installed.
- Locations are listed from where security updates can be downloaded.

The settings monitored by the SHV in Windows 8 and Windows 8.1 are based on the settings that are monitored by Windows Security Center on the client computer. For software to be monitored, it must be compatible with the Windows Security Center.

You can extend NAP to monitor additional settings and software. You can do this by deploying additional SHAs on NAP clients, and additional SHVs on servers that are running NPS.

Considerations for Defining a NAP Policy

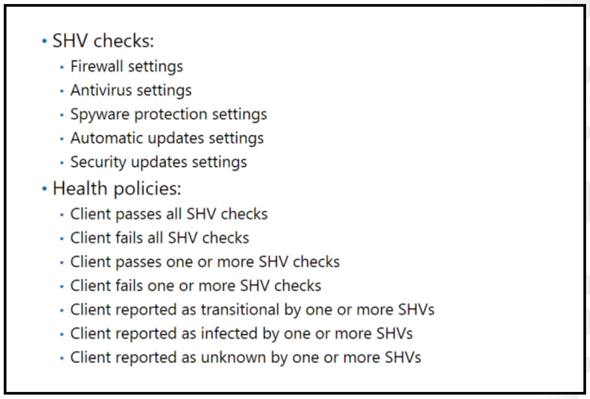

- SHV checks:
  - Firewall settings
  - Antivirus settings
  - Spyware protection settings
  - Automatic updates settings
  - Security updates settings
- Health policies:
  - Client passes all SHV checks
  - Client fails all SHV checks
  - Client passes one or more SHV checks
  - Client fails one or more SHV checks
  - Client reported as transitional by one or more SHVs
  - Client reported as infected by one or more SHVs
  - Client reported as unknown by one or more SHVs

A NAP policy consists of several elements:

- SHV checks. SHV checks determine how you track the health status of client computers. SHV checks consist of the following settings:
  - Firewall
  - Antivirus
  - Spyware protection
  - Automatic updates
  - Security updates

  Your organization's requirements determine which of these checks you select for the SHV. For example, you may be concerned only with the presence of antivirus and antimalware software on the client computers.

- Health policies. Health policies determine how your organization defines the health status of client computers, by evaluating the client SHV checks. When a client computer attempts to connect to the network, SHV checks are performed, and the results are evaluated against the defined health policies. The client computer must be classified under one of the following defined conditions:
  - Client passes all SHV checks.
  - Client fails all SHV checks.
  - Client passes one or more SHV checks.
  - Client fails one or more SHV checks.
  - Client reported as transitional by one or more SHVs.
  - Client reported as infected by one or more SHVs.
  - Client reported as unknown by one or more SHVs.

Typically, you define at least two health policies: one for compliant computers in which all SHVs checks are passed, and one for non-compliant computers in which the client fails one or more SHV checks. When planning your health policies, you may decide to define more than two health policies to evaluate clients as compliant, non-compliant, or unsupported.

Considerations When Defining Your SHV Settings

When defining your SHV settings, you must consider how your organization wants to control client configuration. Careful use of SHV settings can help to protect your computers against network security threats, and can form part of your defense-in-depth approach to network protection. The following are aspects of client configuration.

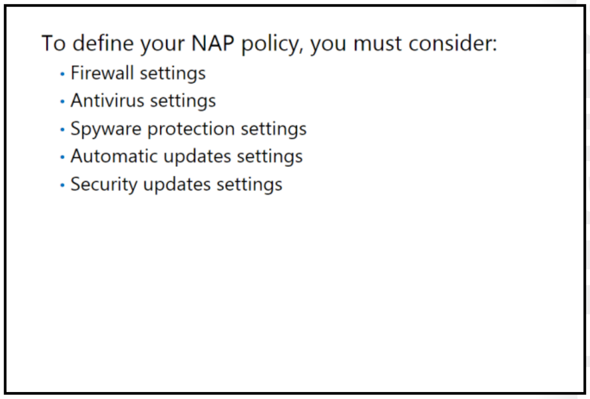

To define your NAP policy, you must consider:

- Firewall settings
- Antivirus settings
- Spyware protection settings
- Automatic updates settings
- Security updates settings

Firewall Settings

Windows SHV supports checking client computers for the presence and enabled status of a host-based firewall product. Specifically, you can select the **A firewall is enabled for all network connections** SHV option. By selecting this SHV option, and by incorporating it into a health policy, you can ensure that all client computers have a firewall enabled.

Antivirus Settings

The requirements regarding antivirus protection can vary widely from one organization to another. While one organization may require only the presence of a running antivirus package on client computers, another organization may insist on an up-to-date antivirus definition, and require that a recent scan have been performed on the client computer.

When determining how best to enforce antivirus software health requirements, consider the following factors:

- Manufacturer. Do you require a particular product from a specified software vendor, or is the presence of any antivirus product sufficient?
- Product version. If your organization requires a specific product, do you require compliant computers to have a specific version of the product, or is the latest version sufficient?
- Updates. Do you require the antivirus patterns to be up-to-date? If so, how do you define up-to-date? Within the last day, week, month?
- Scan configuration. Do you want compliant computers to have a full scan schedule configured? If so, do you additionally want to require that the full scan be performed within a specified time? For example, should the scan have occurred within the last week?

Note: The Windows operating system SHA is capable of determining whether an antivirus application is enabled, and whether it is up-to-date. Other vendor SHAs may determine other characteristics.

Spyware Protection Settings

As with antivirus settings, you can incorporate spyware protection requirements into your health policies as well. The Windows SHV can determine if a spyware protection application is enabled, and whether it is up to date. Other vendor SHAs may be able to determine additional characteristics.

Automatic Updates Settings

The Windows SHA is capable of determining whether automatic updates are enabled on a client computer. Keeping computers up to date is an important aspect of maintaining security for your network.

Security Update Settings

In addition to automatic updates, the Windows SHV also supports the configuration of security update settings. When designing your Security Updates Settings policy, consider the following questions:

- Is it important that Windows operating system updates that can improve security are installed on network computers? If not, do not enable the Windows SHV option to Restrict access for clients that do not have all available updates installed.
- If it is important that Windows operating system security updates are installed, with what frequency should the clients check for these security updates?
- How do you want your clients to obtain these updates? Should clients use Windows Update, WSUS, or both?

Unsupported Platforms

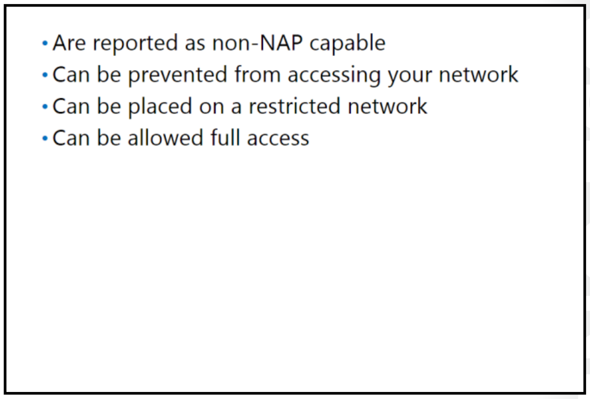

- Are reported as non-NAP capable
- Can be prevented from accessing your network
- Can be placed on a restricted network
- Can be allowed full access

When you design NAP, consider how to accommodate platforms that are not NAP-capable. For example, the Linux platform is not NAP-capable. If your organization has some client computers running Linux, then you must determine how to accommodate those clients.

Unsupported platforms are reported as not NAP-capable, rather than noncompliant. This enables you to differentiate between the two policies. However, a client computer that is NAP-capable will be reported as non-NAP capable if the NAP agent on the computer is disabled.

With unsupported platforms, you can take one of the following actions:

- Block the unsupported platform from accessing the network. You can prevent a client computer that is not NAP-capable from connecting at all. This is probably overly restrictive, but more secure than allowing a NAP-incapable client computer from accessing network resources.
- Place the unsupported platform on a restricted network. This restricts unsupported platforms to using only specific resources. However, it may be difficult to organize your network so that unsupported platforms have access to the resources they need, and still maintain its security with NAP.
- Allow the unsupported platform full access. Allowing unsupported platforms to have full access can be a temporary solution during your migration to using supported client computers. However, this creates a risk because unsupported client computers that are running older Windows operating systems are more likely to be sources of malware and viruses.

Considerations for Defining Your Network Policies

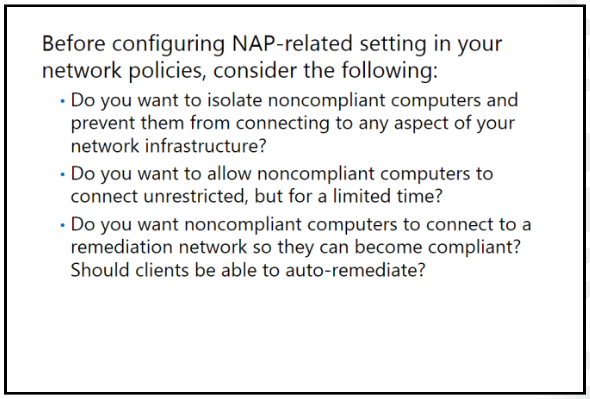

Before configuring NAP-related setting in your network policies, consider the following:

- Do you want to isolate noncompliant computers and prevent them from connecting to any aspect of your network infrastructure?
- Do you want to allow noncompliant computers to connect unrestricted, but for a limited time?
- Do you want noncompliant computers to connect to a remediation network so they can become compliant? Should clients be able to auto-remediate?

To implement NAP, in addition to configuring SHV settings and creating the required health policies, you must also configure appropriate network policies. These network policies require that the defined health policies be defined as single conditions. Typically, you create a network policy for each defined health policy, with that health policy as the defined condition.

When defining a network policy, after you define the conditions, you must then define whether you grant access to computers that meet the policy's conditions. You must select the Access granted option even for network policies that relate to noncompliant computers. This is because you configure NAP enforcement settings on the Network Policy Settings page.

When considering how to configure the NAP-related settings in your network policies, think carefully about what you want to achieve:

- Do you want to isolate noncompliant computers and prevent them from connecting to any aspect of your network infrastructure?

- Do you want to allow noncompliant computers to connect unrestricted but only for a limited time? For example, you may choose this configuration to test NAP rather than to enforce NAP.
- Do you want noncompliant computers to connect to a remediation network so they can become compliant? If so, should clients be able to auto-remediate?

Considerations for Selecting an Enforcement Method

NAP supports four enforcement methods: IPsec, 802.1X, VPN, and Dynamic Host Configuration Protocol (DHCP). When designing NAP enforcement, consider the benefits of each of these methods. You must also understand how your users work, and how their computers connect to your network. For example, there is little point in configuring VPN enforcement if your users do not attempt to connect using VPNs.

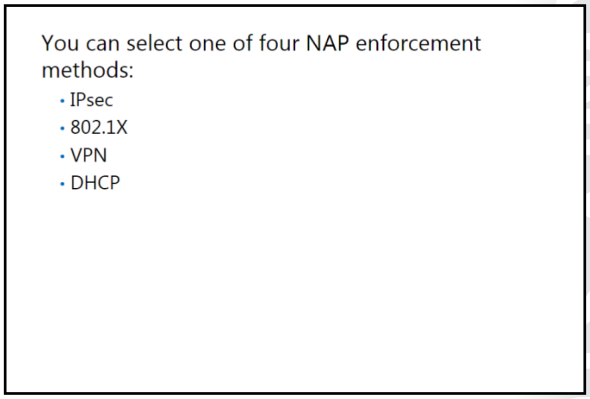

You can select one of four NAP enforcement methods:

- IPsec
- 802.1X
- VPN
- DHCP

IPsec Enforcement Considerations

With IPsec enforcement, a computer must first be compliant to initiate communications with other compliant computers. IPsec enforcement confines communication to compliant computers after they have connected successfully and obtained a valid IP address configuration. IPsec enforcement is the strongest form of limited network access or communication in NAP.

To implement IPsec enforcement, you must install additional software components on the network. You must have a Health Registration Authority (HRA) to act as an enforcement point, and a CA to generate health certificates. The HRA verifies the reported health status of a computer, and then issues a health certificate to the computer. IPsec then uses the health certificate for IPsec authentication.

No specific hardware components are required to implement IPsec. Therefore, you can implement IPsec enforcement in any environment.

IPsec enforcement divides a physical network into three logical networks: secure network, boundary network, and restricted network. A computer can be a member of only one of these logical networks at any time. The logical networks are defined in terms of which computers have health certificates, and which require IPsec authentication with health certificates for incoming communication attempts. The logical networks allow for limited network access and remediation, and provide compliant computers with a level of protection from noncompliant computers. When a computer is noncompliant, the computer is unable to complete IPsec authentication, and subsequently is limited to a restricted network. Ideally, the restricted network should have remediation servers on it.

Considerations for IPsec enforcement are as follows:

- IPsec enforcement is more complex to implement than other enforcement methods, because it requires an HRA and a CA.
- No additional hardware is required to implement IPsec enforcement. IPsec does not require upgrade switches or Wireless Application Protocols (WAPs), which might be required if you select 802.1X enforcement. You can implement IPsec enforcement in any environment.
- IPsec enforcement is very secure and difficult to circumvent.
- You can configure IPsec to encrypt communication for additional security.
- IPsec enforcement applies to IPv4 and IPv6 communication.

802.1X Enforcement Considerations

When you configure IEEE 802.1X enforcement, a computer must be in a state of compliance to obtain unlimited network access through an IEEE 802.1X–authenticated network connection. This can be an authenticating Ethernet switch, or an authenticating wireless access point (AP).

Authenticating switches or wireless APs act as enforcement points for NAP client computers. The client computer's health status forwarded to the NPS server as part of the authentication process.

When a computer is noncompliant, the switch places the computer on a separate virtual LAN (VLAN), or uses packet filters to restrict access only to remediation servers.

Considerations for 802.1X enforcement are as follows:

- Noncompliant computer isolation is enforced by the switch or wireless AP that connects with the client computer. This makes it very difficult to circumvent, and therefore very secure.
- You should use 802.1X enforcement for internal computers. This type of enforcement is appropriate for LAN computers with both wired and wireless connections.
- You cannot use 802.1X enforcement if your switches and wireless APs do not support the use of 802.1X for authentication.

VPN Enforcement Considerations

With VPN enforcement, a computer must be compliant to obtain unlimited network access through a remote access VPN connection. For noncompliant computers, network access is limited through a set of IP packet filters that the VPN server applies to the VPN connection.

VPN enforcement requires the use of a NAP-integrated VPN server. The Routing and Remote Access service (RRAS) server included with Windows Server 2012 is NAP-integrated. The client computer health status is sent as part of the authentication process.

When a computer is noncompliant, the VPN connection is still authenticated. However, IP filters restrict access to only remediation servers.

Considerations for VPN enforcement include the following:

- VPN enforcement is best suited to situations in which a VPN is already in use. It is unlikely that you would implement VPN connections on an internal network to use VPN enforcement.
- Use VPN enforcement to ensure that employees who connect from home computers are not introducing malware to your network. Home computers are often not well maintained by users and represent a higher risk. Many users do not have antivirus software, or do not apply Windows operating system updates regularly.
- Use VPN enforcement to ensure that roaming laptops are not introducing malware to your network. Roaming laptops are more susceptible to malware than computers that are connected directly to the corporate network. This is because they may be unable to download antivirus updates, Windows application updates, and operating system updates from outside the corporate network. Roaming laptops are also more likely to be in environments where malware is present.

DHCP Enforcement Considerations

With DHCP enforcement, a computer must be compliant to obtain an unlimited access IPv4 address configuration from a DHCP server. For noncompliant computers, network access is restricted with an IPv4 address configuration that limits access to the restricted network.

DHCP enforcement requires the use of a NAP-integrated DHCP server. The DHCP server included with Windows Server 2012 is NAP-integrated for IPv4 addressing, but not for IPv6. The health status of the client computer is sent with the DHCP lease request.

If the client computer is noncompliant, a lease is given with the following settings:

- A default gateway of 0.0.0.0
- A subnet mask of 255.255.255.255
- Static routes to remediation servers

Considerations for DHCP enforcement include the following:

- DHCP enforcement is easy to implement, and can apply to any computer with a dynamic IP address.
- DHCP enforcement is easy to circumvent. A user can circumvent DHCP enforcement by using a static IP address. In addition, a noncompliant computer user could add static host routes to reach servers that are not remediation servers.
- DHCP enforcement for IPv6 clients is not possible. If computers on your network use IPv6 addresses to communicate, DHCP enforcement is ineffective.

Planning IPsec Enforcement for NAP

IPsec enforcement divides a physical network into three logical networks: secure network, boundary network, and restricted network. A computer is a member of only one of these logical network at any time. The logical networks are defined by which computers have health certificates and which computers require IPsec authentication with health certificates for incoming communication attempts. The logical networks allow for limited network access and remediation, and provide compliant computers with protection from noncompliant computers.

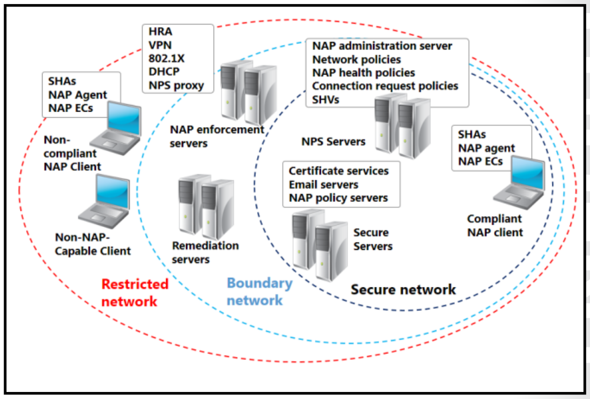

IPsec enforcement defines the three logical networks as follows:

- Secure network. The *secure network* is a set of computers that have health certificates, and which require incoming communication attempts to use health certificates for IPsec authentication. On a managed network, most server and client computers that are members of the Active Directory domain would be in the secure network.
- Boundary network. The *boundary network* is a set of computers that have health certificates but do not require incoming communication attempts to use health certificates for IPsec authentication. Computers in the boundary network must be accessible to computers on the entire network.
- Restricted network. The *restricted network* is a set of computers that do not have health certificates. This includes noncompliant NAP client computers, guests on the network, or computers that are not NAP-capable, such as computers that are running versions of Windows operating systems that do not support NAP, or Apple Macintosh and UNIX-based computers.

Based on the three logical networks, the types of initiated communications possible are secure, boundary, and restricted.

Secure Network

Computers in the secure network can initiate communications with computers in all three logical networks. Communications initiated to computers in the secure network or boundary network are authenticated with IPsec and health certificates. Communications initiated to computers in the restricted network are not authenticated with IPsec.

Computers in the secure network accept communications initiated from computers in the secure and boundary networks that IPsec authenticates, but do not accept communications initiated from computers in the restricted network. For example, a client computer in the secure network can request a webpage from a web server in the secure network. However, a client computer in the restricted network cannot.

You can configure the requirements for initiated communication on a TCP or UDP port basis to limit specific traffic. For example, you can require IPsec authentication with health certificates for remote procedure call (RPC) traffic, but not for web traffic. In this case, a client computer in the restricted network could request a webpage from a web server in the secure network, but not be able to use RPC to connect to that same server.

Boundary Network

Computers in the boundary network can initiate communications with computers in either the secure or boundary networks, providing they are authenticated with IPsec and health certificates. Computers in the boundary network also can initiate communication with computers in the restricted network even though IPsec does not authenticate them.

Computers in the boundary network will accept communications that are initiated from computers in both the secure and boundary networks, providing they are authenticated with IPsec and health certificates. Computers in the boundary network will also accept communications from computers in the restricted network that IPsec does not authenticate.

Boundary network members typically consist only of the HRA and NAP remediation servers. Servers in the boundary network must be accessible from noncompliant NAP clients in the restricted network to perform initial remediation functions and to obtain health certificates. Additionally, these servers must be accessible from compliant computers in the secure network to perform ongoing remediation functions, renew health certificates, and manage computers in the boundary network.

A computer is a member of the secure or boundary network for the time specified in the health certificate's validity period. Before the health certificate expires, the IPsec-protected NAP client contacts the HRA to obtain a new health certificate. You can configure the validity time for health certificates on the HRA. Validity time typically spans hours rather than years, in the case of computer or user certificates.

Restricted Network

Computers in the restricted network can initiate communications with computers in the restricted and boundary networks. Computers in the restricted network cannot initiate communications with computers in the secure network, unless they are allowed specifically through the IPsec policy settings of the secure network's computers. However, computers in the restricted network can accept communications initiated from computers in all three logical networks.

Configuring HRA Server

To support IPsec NAP enforcement, you must configure an HRA server. This process involves the following steps:

1. Configure authentication requirements. When you install the HRA, you are prompted to configure the HRA to either issue certificates only when users are authenticated to the domain, or to optionally provide health certificates to anonymous users. If you select to allow only domain-authenticated users, a single website, DomainHRA, is created to support this configuration. If you choose to allow anonymous users to obtain health certificates, an additional website, NonDomainHRA, is created to support that configuration.

2. Configure CAs. The HRA must be associated (either during installation or immediately following), with either a stand-alone or enterprise CA. This is discussed further in the next topic.

3. Configure the request policy. The security settings that the HRA uses to communicate with clients are known as *request policy settings*. You can use the HRA snap-in to specify these security mechanisms and to determine which asymmetric key algorithm, hash algorithm, and cryptographic service provider (CSP) the HRA server uses to encrypt communication with client computers.

Note: Configuring request policy settings on your HRA server is not mandatory. By default, a NAP-capable client computer initiates a negotiation process with an HRA server by using a mutually acceptable default security mechanism for encrypting communication.

To obtain and issues certificates, the HRA must be associated with a CA. This process involves choosing a CA type, verifying CA security settings, and configuring additional settings.

Choosing a CA Type

When configuring the HRA to use a CA, you can select one of the following:

- Stand-alone CA. A *stand-alone CA* issues certificates that are not based on templates. Consequently, you do not need to configure a certificate template. However, you must still configure CA security settings and certificate issuance requirements so that the HRA can request and issue health certificates automatically to client computers that are compliant with health policies.
- Enterprise CA. An *enterprise CA* certificate is based on a specific template. Therefore, if you select an enterprise CA, you must configure the required certificate template as part of the CA preparation process.

Note: If your enterprise CA is installed on a computer that is running Windows Server 2008 or Windows Server 2012, the required HRA template already exists. If your enterprise CA is running Windows Server 2003, then you must create the required template manually.

Complete the following tasks on your enterprise CA to ensure that it is ready to support the requirements of your HRA:

1. Verify certificate availability. Use the Certificate Templates snap-in to check for the presence of the System Health Authentication template.
2. Verify certificate enrollment permissions for the HRA. To check that the HRA has the required permissions to obtain and issue health certificates, open the Certificate Templates snap-in and view the properties of your System Health Authentication template. Check the security settings to ensure that both the Enroll and Autoenroll permissions have been granted to the DNS name of your HRA server.

Verify CA Security Settings

After selecting the CA type, you must now verify the CA security settings. For NAP client computers to obtain health certificates automatically once they have been determined to be compliant with network health requirements, you must configure your NAP CAs to issue health certificates automatically. Use the following process to ensure that certificates are issued automatically.

1. Open the Certification Authority management console snap-in.
2. Verify that the Policy Module for your CA is configured with the following value: **Follow the settings in the certificate template, if applicable. Otherwise, automatically issue the certificate.**

Note: This process applies to both enterprise and stand-alone CA servers.

Configure Additional Settings

You can add the CA to the HRA during installation of the HRA role or at any time thereafter by using the HRA console. After you add the CA to the HRA, you can use the HRA console to complete these additional tasks:

1. Configure the CA wait time. The HRA attempts to obtain health certificates only from the CA that is configured first in the processing order, unless that CA has been marked as unavailable. You can change the number of minutes to wait before identifying a CA as unavailable.

2. Configure health certificate validity period. Client computers attempt to renew their health certificate 15 minutes prior to expiration, or when a change in client health status occurs. You can configure a custom validity period for health certificates. The default validity period is 4 hours.

IPsec enforcement policies

After configuring the HRA, the next step prior to configuring NAP with IPsec enforcement is to configure IPsec Enforcement policies. These policies create the three logical networks in the enforcement topology: the secure network, the boundary network, and the restricted network.

The process of configuring the IPsec policies consists of the following steps:

1. Create the required organizational units (OUs). You must create two OUs for the application of GPOs:

- You must create a boundary OU for computers with NAP exemption certificates that request, but do not require, that incoming communications authenticate with a health certificate.

- You must also create a secure OU for computers running the Windows Vista, Windows 7, Windows 8, Windows Server 2008, or Windows Server 2012 operating systems.

2. Create GPOs. You must create and link a GPO to each of the OUs you created:

- Boundary GPO. Within the GPO, configure Windows Firewall with Advanced Security. Create an isolation rule with these settings:

  - Requirements: Request authentication for inbound and outbound connections

  - Authentication method: Only accept health certificates

- Secure GPO. Within the GPO, configure Windows Firewall with Advanced Security. Create an isolation rule with these settings:

  - Requirements: Require authentication for inbound connections and request authentication for outbound connections

  - Authentication method: Only accept health certificates

3. Enforce IPsec rules on the client computers. After completing the previous steps, you must now move the relevant computers to the appropriate OUs so that the correct GPOs apply.

Remediation Server Groups

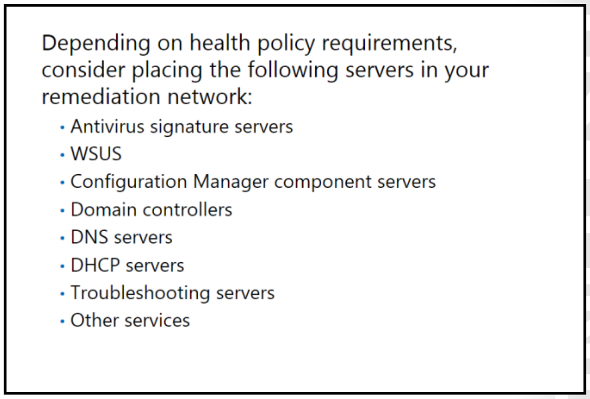

Depending on health policy requirements, consider placing the following servers in your remediation network:

- Antivirus signature servers
- WSUS
- Configuration Manager component servers
- Domain controllers
- DNS servers
- DHCP servers
- Troubleshooting servers
- Other services

A *remediation server group* is a list of restricted network servers that provide resources to bring noncompliant NAP-capable clients into compliance with defined client health policies.

A remediation server hosts the updates that a NAP agent can use to bring noncompliant client computers into compliance with health policy, as defined by NPS. For example, a remediation server can host antivirus signatures. If a health policy requires that client computers have the latest antivirus definitions, then the antivirus SHA, the antivirus SHV, the antivirus policy server, and the remediation server work together to update noncompliant computers.

Depending on your health policy requirements, consider placing the following servers in your remediation network:

- Antivirus signature servers. These servers enable noncompliant computers to access antivirus updates.
- WSUS. This service enables noncompliant computers to obtain the required software updates.
- Microsoft® System Center Configuration Manager component servers. Configuration Manager management points, software update points, and distribution points host the software updates required to bring computers into compliance.
- Domain controllers. Noncompliant computers might require access to domain services on the noncompliant network for authentication purposes, to download policies from Group Policy, or to maintain domain profile settings.
- DNS servers. Noncompliant computers must have access to DNS to resolve host names.
- DHCP servers. Noncompliant computers must have access to a DHCP server if the client’s IP profile changes on the noncompliant network, or if the DHCP lease expires.
- Troubleshooting servers. When you configure a remediation server group, you can provide a troubleshooting URL with instructions about how to bring computers back into compliance with your health policies.
- Other services. Consider providing access to the Internet on your remediation network so that noncompliant computers can reach remediation services, such as Windows Update and other Internet resources.

Demonstration: Implementing NAP

This demonstration shows how to:

- Install the NPS server role.
- Configure NPS as a NAP health policy server.
- Configure health policies.
- Configure network policies for compliant computers.
- Configure network policies for noncompliant computers.

- Configure the DHCP server role for NAP.
- Configure client NAP settings.
- Test NAP.

Demonstration Steps

Install the NPS server role

1. Switch to LON-DC1, and sign in as a domain administrator.
2. Open Server Manager, and install the Network Policy and Access Services role.

Configure NPS as a NAP health policy server

1. Open the Network Policy Server console.
2. Configure the Windows Security Health Validator to require that all Windows 8 computers are running a firewall.

Configure health policies

1. Create a health policy named **Compliant** in which the condition is that **Client passes all SHV checks**.
2. Create another health policy named **Noncompliant** in which the condition is that **Client fails one or more SHV checks**.

Configure network policies for compliant computers

1. Disable the two existing network policies. Otherwise, these policies would interfere with the processing of the policies you are about to create.
2. Create a new network policy named **Compliant-Full-Access** that has a condition of the **Compliant** health policy. For this policy, grant computers unrestricted access.

Configure network policies for noncompliant computers

- Create a new network policy named **Noncompliant-Restricted** that has a condition of the **Noncompliant health policy**. Grant computers **restricted access**.

Configure the DHCP server role for NAP

1. Open the DHCP console.
2. Modify the properties of the IPv4 scope to support NAP.
3. Create a new DHCP policy that allocates appropriate DHCP scope options to noncompliant computers. These options should assign a DNS suffix of **restricted.Adatum.com**.

Configure client NAP settings

1. On LON-CL1, enable the **DHCP Quarantine Enforcement Client**.
2. Start the **Network Access Protection Agent** service.
3. Use the local Group Policy Management Console to enable the Security Center.
4. Reconfigure LON-CL1 to obtain an IP address from a DHCP server.

Test NAP

1. Verify the obtained configuration by using the Ipconfig tool.
2. Disable and stop the Windows Firewall service.
3. In the System Tray area, click the **Network Access Protection** pop-up warning. Review the information in the **Network Access Protection** dialog box, and then click **Close**.

Note: Proceed with the following steps even if the pop-up window does not display.

4. Verify the obtained configuration by using the Ipconfig tool.
5. Notice that the computer has a subnet mask of 255.255.255.255 and a DNS Suffix of restricted.Adatum.com.
6. Leave all windows open.

Lab: Designing and Implementing Network Protection

Scenario

A. Datum Corporation has recently experienced problems with malware and the interception of network traffic. The malware was introduced from computers that were not compliant with corporate security and maintenance policies. The network interception was caused by an unknown person connecting a packet-sniffing device physically to the network.

A. Datum wants to improve its security by ensuring that sensitive servers are protected by appropriate firewall rules, and that sensitive servers only accept network traffic if that traffic is authenticated and encrypted. Any clients interacting with sensitive servers need to have met a minimum client health benchmark.

Objectives

At the end of this lab, you should be able to:

- Design a Windows Firewall solution.
- Implement Windows Firewall.
- Design a NAP solution.
- Implement a NAP solution with IPsec enforcement.

Lab Setup

Estimated Time: 75 minutes

Virtual machines	20413C-LON-DC1 20413C-LON-SVR1 20413C-LON-SVR2 20413C-LON-CL1
User Name	Adatum\Administrator
Password	Pa$$w0rd

For this lab, use the available virtual machine environment. Before you begin the lab, 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 **20413C-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in using the following credentials:

- User name: **Administrator**
- Password: **Pa$$w0rd**
- Domain: **Adatum**

5. Repeat steps 2 through 4 for 20413C-LON-SVR1 and 20413C-LON-CL1.

Note: Important: Do not start 20413C-LON-SVR2 until instructed to do so.

Exercise 1: Designing a Windows Firewall Solution

Scenario

A. Datum IT staff are concerned that the current structure for security is not efficient for allocating resources. A team of security consultants has recently completed an analysis of security needs and risks, and you must now implement some of the proposed changes within the network infrastructure.

After analyzing network security by using the defense-in-depth model, the security consultant team advises that you could improve internal security by implementing Windows Firewall. To maximize security, they propose implementing outbound rules on workstations and servers.

To secure network communications further, the consultant team has proposed that you secure communication between all users in the A. Datum organization's research group. This prevents non-research users from accessing research data or applications.

To support the implementation of Windows Firewall rules, the IT staff has created a complete list of network applications currently in use at A. Datum. The applications and their ports are listed in the following tables.

Client application	Destination port	Executable
File and printer sharing server message block (SMB)	445	n/a
Research custom application	10101	Research.exe
Internet Explorer	80, 443, 8080	Iexplore.exe
Email client	Random RPC	Outook.exe

Server application	Incoming port	Executable
File and printer sharing (SMB)	445	n/a
Research custom application	10101	ResearchSrv.exe
Customer service web application	8080	n/a
Email server	Random RPC	Store.exe

Windows Firewall Deployment Strategy	
Document Reference Number: BS0929/1	
Document Author Date	Brad Sutton 29th September
• **Requirements Overview** • Design Windows Firewall to help to address the following issues: • In addition to the listed server applications, you must also account for network services, such as domain logons and DNS lookups. • There are no more than two servers running any specific application.	

Windows Firewall Deployment Strategy

Proposals

1. What inbound rules should you implement on servers?
2. What outbound rules should you implement on servers?
3. What inbound rules should you implement on Windows 8 workstations?
4. What outbound rules should you implement on Windows 8 workstations?
5. How will you deploy Windows Firewall on servers running Windows Server?
6. How will you deploy Windows Firewall on workstations?

The main tasks for this exercise are as follows:

1. Read the supporting documentation
2. Update the proposal document with your planned course of action
3. Examine the suggested proposals in the Lab Answer Key
4. Discuss your proposed solution with the class, as guided by your instructor

▶ Task 1: Read the supporting documentation

- Read the documentation provided.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the Windows Firewall Deployment Strategy document.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones in the Lab Answer Key.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you should have completed the Windows Firewall design for A. Datum.

Exercise 2: Implementing a Windows Firewall solution

Scenario

You must now implement a trial of your proposed Windows Firewall solution within the A. Datum network.

The main tasks for this exercise are as follows:

1. Move the computers to the Research OU
2. Create a Group Policy Object (GPO) and link it to the Research OU
3. Create the required connection security rules
4. Create the firewall rules

5. Refresh the Group Policy settings

6. Attempt to connect to the web server

7. Verify settings

8. To prepare for the next exercise

▶ Task 1: Move the computers to the Research OU

1. Switch to LON-DC1.
2. Open Active Directory Users and Computers.
3. Move LON-CL1 and LON-SVR1 from the Computers container to the Research OU.

▶ Task 2: Create a Group Policy Object (GPO) and link it to the Research OU

1. Open Group Policy Management.
2. Create a new GPO named **Research Department Application Security Policy** and link it to the Research OU.

▶ Task 3: Create the required connection security rules

1. Open the Research Department Application Security Policy for editing.
2. In the Group Policy Management Editor, under Computer Configuration, browse to **Policies\Windows Settings\Security Settings\Windows Firewall with Advanced Security\Windows Firewall with Advanced Security – LDAP://CN={GUID}**, and then click **Connection Security Rules**.
3. Create a new Connection Security Rule with the following properties:
   - Rule type: **Custom**
   - Endpoints: Default
   - Requirements: **Require authentication for inbound connections and request authentication for outbound connections**
   - Authentication method: **Computer and user (Kerberos V5)**
   - Protocol and Ports:
   - Endpoint 1: **TCP** port **80**
   - Endpoint 2: default
   - Profile: Domain
   - Name: **Research Department Application Security rule**

▶ Task 4: Create the firewall rules

1. Create a new Inbound firewall rule with the following properties:
   - Rule type: **Custom**
   - Program: Default
   - Protocol and Ports: Local port: **TCP** port **80**
   - Scope: Default
   - Action:

- **Allow the connection if it is secure**

- **Allow the connection if it is authenticated and integrity-protected**
    - Users: Default
    - Computers: **Only allow connections from these computers**:
- **LON-CL1**
- **LON-SVR1**
    - Profile: Domain
    - Name: **Research Department Application Firewall rule**

▶ Task 5: Refresh the Group Policy settings

1. Switch to LON-CL1.
2. Open a command prompt, type **gpupdate /force**, and then press Enter.
3. Restart the computer.
4. Switch to LON-SVR1.
5. Open a command prompt, type **gpupdate /force**, and then press Enter.
6. Restart the computer. Wait until LON-SVR1 has restarted before proceeding.

▶ Task 6: Attempt to connect to the web server

1. Switch to LON-CL1. You must wait until LON-SVR1 has restarted before proceeding.
2. Sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
3. From the desktop, start Internet Explorer, and attempt to open the **http://LON-svr1** webpage. The default Internet Information Services (IIS) 8 webpage should load successfully.

▶ Task 7: Verify settings

1. Open Windows Firewall with Advanced Security.
2. Open Monitoring, click **Security Associations**, and then click **Main Mode**.
3. Double-click any associations listed.
4. Make a note of the First authentication method.
5. Expand **Quick Mode**.
6. In the right pane, double-click the item listed.
7. Make a note of the Remote port.

▶ Task 8: To prepare for the next exercise

When you are finished with this exercise, revert two of the virtual machines to their initial state, and start two additional virtual machines. To do this perform the following steps:

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-CL1**, and then click **Revert**.
3. In the **Revert Virtual Machines** dialog box, click **Revert**.
4. Repeat step 2 and 3 for 20413C-LON-SVR1 and 20413C-LON-DC1.
5. Click **20413C-LON-DC1**, and in the Actions pane, click **Start**.
6. In the Actions pane, click **Connect**. Wait until the virtual machine starts.

7. Sign in using the following credentials:

- User name: **Administrator**
- Password: **Pa$$w0rd**
- Domain: **Adatum**

8. Repeat steps 5 through 7 for 20413C-LON-SVR2, and 20413C-LON-CL1.

Results: After completing this exercise, you should have configured the required firewall rules.

Exercise 3: Designing a Network Access Protection (NAP) Solution

Scenario

You are required to design and implement a NAP solution to provide additional protection in the A. Datum network.

The Paris site provides services for all branch offices in the European area. However, for now, NAP is being implemented in London as a trial for the rest of A. Datum. You must consider various scenarios, and then test them.

NAP Deployment Strategy

Document Reference Number: BS1001/1

Document Author	Brad Sutton
Date	1st October

Requirements Overview

Design a NAP strategy to support the following objectives:

- Help to prevent the spread of malware in the A Datum European offices.
- Ensure that only client computers that have a firewall enabled, have recently checked the local WSUS server for updates, and have an active, up-to-date antimalware solution are able to connect to computers in the Research department.
- Ensure that computers that are not up-to-date with anti-malware definitions or WSUS server updates are able to obtain those update.

Additional Information

- All client computers have been upgraded to Windows 8.1.
- All client computers use dynamic IP addresses.
- A WSUS server is located on the internal network.
- All clients and servers in the Research department are members of the A. Datum Active Directory domain.

Proposals

1. Which NAP enforcement method is appropriate for the given scenario?

NAP Deployment Strategy

2. What is the simplest way to apply the necessary client computer configuration to several computers simultaneously?

3. How will you ensure that the configuration applies to only client computers, and not to servers?

4. What determines the options available for verifying client computer status? How can you expand these options?

5. How do noncompliant computers access remediation servers?

6. Which servers should you configure as remediation servers?

7. In addition to the existing servers hosting the Domain Controller, DNS, DHCP, and WSUS server roles, what additional roles should you plan to deploy to implement this solution?

The main tasks for this exercise are as follows:

1. Read the supporting documentation

2. Update the proposal document with your planned course of action

3. Examine the suggested proposals in the Lab Answer Key

4. Discuss your proposed solution with the class, as guided by your instructor

▶ Task 1: Read the supporting documentation

- Read the documentation provided.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the NAP Deployment Strategy document.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones in the Lab Answer Key.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you should have completed the NAP design.

Exercise 4: Implementing NAP with IPsec Enforcement

Scenario

You will now implement NAP with IPsec enforcement.

The main tasks for this exercise are as follows:

1. Configure NAP health compliance certificates

2. Install the Network Policy and Access Services role on LON-SVR2

3. Configure the HRA

4. Configure health policies

5. Configure the health validator

6. Create the NAP client configuration GPO

7. Test a client on the internal network

8. Test compliance

9. To prepare for the next module

▶ Task 1: Configure NAP health compliance certificates

1. On LON-DC1, open **certsrv.msc**.

2. Duplicate the Workstation Authentication certificate template. Use the following properties:

- **General** tab, **Template display** name: **NAP Health Certificate**

- **Subject Name** tab: **Supply in the request**

- **Extensions** tab, **Application Policies**: add **System Health Authentication**

3. Save the new template, and then issue it.

4. Modify the properties of the AdatumCA CA:

- Security tab: allow LON-SVR2 all permissions on the CA.

5. Open an elevated Windows PowerShell window, and run the following three commands, pressing Enter at the end of each line:

```
Certutil -setreg policy\EditFlags +EDITF_ATTRIBUTEENDDATE
net stop certsvc
net start certsvc
```

▶ Task 2: Install the Network Policy and Access Services role on LON-SVR2

1. On LON-SVR2, request a Computer certificate from the **AdatumCA** certification authority.
2. Using Server Manager, install the Network Policy and Access Services role with the following role services:
   - **Network Policy Server**
   - **Health Registration Authority**
3. Configure the settings as follows:
   a. Use an existing remote CA: **AdatumCA**
   b. Click **Yes, require requestors to be authenticated as members of a domain**.
   c. Choose an existing certificate for SSL encryption: **LON-SVR2.Adatum.com**

▶ Task 3: Configure the HRA

1. On LON-SVR2, in Server Manager, open the Health Registration Authority.
2. Verify that **LON-DC1.Adatum.com\AdatumCA** is listed in the **Certification Authority** list.
3. Open the **Certification Authority** object's properties, and make the following changes:
   - **Use enterprise certification authority**:
     a. Authenticated compliant certificate template: **NAP Health Certificate**
     b. Anonymous compliant certificate template: **NAP Health Certificate**
     c. Verify that the validity period for certificates approved by this Health Registration Authority is set to **4** hours.
4. Close the Health Registration Authority window.

▶ Task 4: Configure health policies

1. On LON-SVR2, open the Network Policy Server console.
2. Run the Configure NAP Wizard:
   a. On the **Select Network Connection Method For Use with NAP** page, in **Network connection method**, select the **IPsec with Health Registration Authority (HRA)** check box.
   b. On the **Define NAP Health Policy** page, clear the **Enable auto-remediation of client computers** check box.

▶ Task 5: Configure the health validator

1. On LON-SVR2, in the Network Policy Server console, under Network Access Protection, open the Default Configuration for the Windows Security Health Validator.
2. On the **Windows 8/Windows 7/Windows Vista** tab, clear all check boxes, and then select the **A firewall is enabled for all network connections** check box.

▶ Task 6: Create the NAP client configuration GPO

1. On LON-DC1, use the Active Directory Users and Computers console to create the **NAP_Clients OU**.
2. Move the LON-CL1 computer account into this OU.
3. On LON-DC1, open the Group Policy Management Console.
4. Create a GPO named **NAP Client Configuration**, and then link it to the **NAP_Clients OU**.

5. Open the new GPO for editing, and configure the following settings:

    a. Browse to **Computer Configuration\Policies\Windows Settings\Security Settings\Network Access Protection\NAP Client Configuration\Enforcement Clients**:

    - Enable **IPsec Relying Party**.

    b. Browse to **Computer Configuration\Policies\Windows Settings\Security Settings\Network Access Protection\NAP Client Configuration\Health Registration Settings\Trusted Server Groups**, and then create a new group named **Internal**:

    - On the **Add Servers** page, type **https://lon-svr2.adatum.com/domainhra/hcsrvext.dll**, click **Add**, and then click **Finish**.

    c. Browse to **Computer Configuration\Policies\Windows Settings\Security Settings\System Services**:

    i. Double-click **Network Access Protection Agent**, and in the **Network Access Protection Agent Properties** dialog box, select the **Define this policy setting** check box.

    ii. Under **Select services startup mode**, click **Automatic**, and then click **OK**.

    d. Browse to **Computer Configuration\Policies\Administrative Templates: Policy definitions (ADMX files) retrieved from the local computer\Windows Components\Security Center**:

    i. In the details pane, double-click **Turn on Security Center**.

    ii. In the **Turn on Security Center** dialog box, click **Enabled**, and then click **OK**.

6. Close the Group Policy Management Editor.

▶ Task 7: Test a client on the internal network

1. On LON-CL1, sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.

2. In Windows PowerShell, run the following commands, pressing Enter at the end of each line:

```
gpupdate /force
get-service napagent
netsh nap client show grouppolicy
```

Note: These commands verify that the correct policies are applying to the NAP client.

▶ Task 8: Test compliance

1. Open Windows PowerShell.

2. At the command prompt, type the following command, and then press Enter:

```
napstat
```

3. Notice that the client is fully compliant.

4. Disable and stop the Windows Firewall service.

5. Restart LON-CL1, and sign in as **Adatum\administrator** with the password **Pa$$w0rd**.

6. Repeat steps 1 through 3.

    You should see that the client is no longer compliant.

▶ Task 9: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. In the Virtual Machines list, right-click **20413C-LON-CL1**, and then click **Revert**.
3. In the **Revert Virtual Machines** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for 20413C-LON-SVR2 and 20413C-LON-DC1.

Results: After completing this exercise, you should have implemented NAP with IPsec enforcement.

Question: What was your approach to the firewall design exercise?

Question: What was your approach to the NAP design exercise?

Question: How does the network access design compare with network access implementation in your organization?

Module Review and Takeaways

Review Questions

Verify the correctness of the statement by placing a mark in the column to the right.

Statement	Answer
The Windows SHV can determine both the state of the firewall (enabled or disabled), and whether it is up to date.	

Question: What are some common forms of network attacks?

Question: What server role (or roles) must you deploy to support NAP?

Question: What are the recommended uses for IPsec?

Real-world Issues and Scenarios

Scenario: Tailspin Toys is planning to implement NAP as part of its overall security infrastructure. They want an enforcement method that is applicable to all network clients, regardless of how they connect. Currently, a PKI is in place. What enforcement method or methods would you recommend?

Answer: IPsec enforcement would be suitable, as would 802.1X enforcement, depending on whether the switches and access points support 802.1X authentication. DHCP enforcement would be unsuitable, because clients with a manually assigned IP configuration can bypass NAP. VPN enforcement would also be unsuitable, because not all clients are connecting by VPN.

Scenario: Wingtip Toys wants to implement IPsec NAP enforcement. What infrastructure components must be in place to support this method?

Answer: Aside from the general requirements for NAP, IPsec also requires that you deploy a Health Registration Authority (HRA) and a Public Key Infrastructure (PKI) for health certificates.

Module 12

Designing and Implementing Remote Access Services

Contents:

Module Overview

When you are designing remote network access, you must carefully plan access for users who attempt to connect to network resources from outside your organization. Without proper consideration of how best to enable remote network access, you risk leaving your network resources open to security breaches and possible misuse.

You can implement several technologies to provide remote access, depending on the data needing to be accessed, the type of client computers being used, and existing security requirements. The remote access technologies available in Windows Server® 2012 R2 that you will learn about in this module are virtual private network (VPN), DirectAccess, and Web Application Proxy.

VPNs allow almost any client computer to access resources on a corporate network by using a public Internet connection. Client computers do not need to run a Windows® operating system to access a VPN, which makes this technology widely implemented. However, VPN connections are only available after a user signs in to their workstation and initiates the connection.

DirectAccess connections can be made automatically when a computer starts. It does not require a user to sign in and initiate the connection. However, DirectAccess is only available to Windows 8.1, Windows 8, and Windows 7 clients.

Not all remote users need access to all corporate network resources while working remotely. Some users might only need access to certain web applications, such as Microsoft® SharePoint® sites, and web mail based on Microsoft Exchange Server®. In such scenarios, you can use a Web Application Proxy to provide access to these web applications to remote users. You can also use Web Application Proxy to publish applications to users who are not part of your company's Active Directory® infrastructure.

Objectives

After completing this module, you will be able to:

- Plan and implement DirectAccess.
- Plan and implement virtual private networking.
- Plan and implement Web Application Proxy.
- Plan a complex remote access infrastructure.

Lesson 1
Planning and Implementing DirectAccess

Organizations often rely on VPN connections to provide remote users with secure access to data and to resources on the corporate network. VPN connections are easy to configure, and are supported by different clients. However, the user is responsible for initiating the VPN connection, and connection settings could require additional configuration on the corporate firewall. In addition, VPN connections usually enable remote access to the entire corporate network, and organizations cannot effectively manage remote computers unless they are connected.

To overcome such limitations in VPN connections, organizations can implement DirectAccess to provide a seamless connection between the internal network and the remote computer on the Internet. With DirectAccess, organizations can manage remote computers because they are always connected.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the components that are required to implement DirectAccess.
- Describe the DirectAccess options in Windows Server 2012 R2.
- Configure DirectAccess by using the Getting Started Wizard.
- Describe the configuration changes that the Getting Started Wizard makes.
- Describe how a simple DirectAccess deployment works.
- Describe the limitations of the DirectAccess Getting Started Wizard.
- Plan for DirectAccess tunnelling protocols.
- Plan for a public key infrastructure (PKI) integration with DirectAccess.
- Plan network configuration options.
- Plan for network location servers.

DirectAccess Components

To deploy and configure DirectAccess, your organization must support the following infrastructure components:

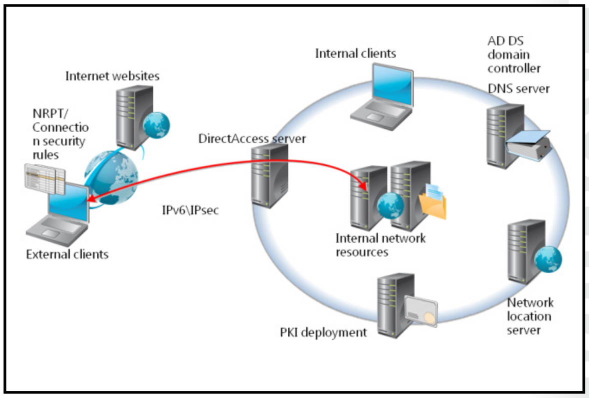

- DirectAccess server
- DirectAccess clients
- Network location server
- Internal resources
- Active Directory domain
- Group Policy
- PKI (optional for the internal network)
- Domain Name System (DNS) server
- Network Access Protection (NAP) server

DirectAccess Server

The DirectAccess server can be any Windows Server 2012 R2 or Windows Server 2012 server that you join to a domain, that accepts connections from DirectAccess clients, and that establishes communication with intranet resources. This server provides authentication services for DirectAccess clients, and acts as an Internet Protocol security (IPsec) tunnel mode endpoint for external traffic. The new remote access server role allows centralized administration, configuration, and monitoring for both DirectAccess and VPN connectivity.

Compared with the previous implementation in Windows Server 2008 R2, the new DirectAccess Getting Started Wizard simplifies DirectAccess management for small and medium organizations. The wizard does this by removing the need for a full PKI deployment, and by removing the requirement for two consecutive public IPv4 addresses for the physical adapter that is connected to the Internet. In Windows Server 2012, the DirectAccess Getting Started Wizard detects the actual implementation state of the DirectAccess server, and selects the best deployment automatically. This hides from the administrator the complexity of manually configuring IPv6 transition technologies.

DirectAccess Clients

DirectAccess clients can be any domain-joined computer that is running Windows 8.1 Enterprise, Windows 8 Enterprise, Windows 7 Enterprise, or Windows 7 Ultimate.

Note: With off-premise provisioning, you can join the client computer in a domain without connecting the client computer in your internal premises.

The DirectAccess client computer connects to the DirectAccess server by using IPv6 and IPsec. If a native IPv6 network is not available, the client establishes an IPv6-over-IPv4 tunnel by using either 6to4 or Teredo. Note that the user does not have to be logged on to the computer for this step to complete.

If a firewall or proxy server prevents the client computer that is using 6to4 or Teredo from connecting to the DirectAccess server, the client computer attempts to connect automatically by using the IP over HTTPS (IP-HTTPS) protocol, which uses a Secure Sockets Layer (SSL) connection to ensure connectivity. The client has access to the Name Resolution Policy Table (NRPT) rules and the connection security tunnel rules.

Network Location Server

DirectAccess clients use the network location server to determine their location. If the client computer can connect with HTTPS, then the client computer assumes it is on the intranet and disables DirectAccess components. If the network location server is not contactable, then the client assumes it is on the Internet. The network location server is installed with the web server role.

Note: You distribute the URL for the network location server via Group Policy Object (GPO).

Internal Resources

You can configure any IPv6–capable application that is running on internal servers or client computers to be available for DirectAccess clients. For older applications and servers that are not based on Windows operating systems and have no IPv6 support, Windows Server 2012 now includes native support for protocol translation (NAT64) and name resolution (DNS64) gateways to convert IPv6 communication from DirectAccess client to IPv4 for the internal servers.

Note: As in the past, you also can achieve this functionality with Microsoft Forefront® Unified Access Gateway. Likewise, as in past versions, these translation services do not support sessions initiated by internal devices, only requests originating from IPv6 DirectAccess clients.

Active Directory Domains

You must deploy at least one Active Directory domain, running a minimum of Windows Server 2008 R2 domain-functional level. Windows Server 2012 DirectAccess provides integrated multiple domain support, which allows client computers from different domains to access resources that may be located in different trusted domains.

Group Policy

Group Policy is required for the centralized administration and deployment of DirectAccess settings. The DirectAccess Getting Started Wizard creates a set of GPOs in addition to the settings for DirectAccess clients, the DirectAccess server, and selected servers.

PKI

PKI deployment is optional for simplified configuration and management. DirectAccess in Windows Server 2012 R2 and Windows Server 2012 enables client authentication requests to be sent over a HTTPS–based Kerberos authentication proxy service that is running on the DirectAccess server. This eliminates the need for establishing a second IPsec tunnel between clients and domain controllers. The Kerberos authentication proxy will send Kerberos authentication requests to domain controllers on behalf of the client.

However, for a full DirectAccess configuration that allows NAP integration, two-factor authentication, and force tunneling, you still will need to implement certificates for authentication for every client that will participate in DirectAccess communication.

DNS Server

When using ISATAP, you must use at least Windows Server 2008 R2, Windows Server 2008 Service Pack 2 (SP2) or newer, or a non-Microsoft DNS server that supports DNS message exchanges over ISATAP.

NAP Servers

NAP is an optional component of the DirectAccess solution that you can use to provide compliance checking and enforcing security policy for DirectAccess clients over the Internet. DirectAccess in Windows Server 2012 R2 and Windows Server 2012 provides the ability to configure NAP health checks directly from the setup user interface, instead of editing the GPO manually as required with DirectAccess in Windows Server 2008 R2.

Windows Server 2012 R2 DirectAccess Options

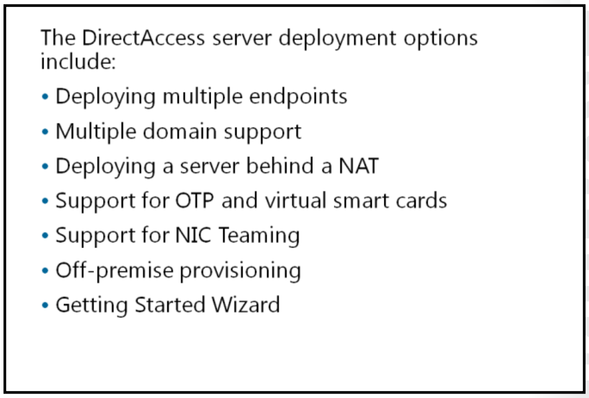

The DirectAccess server deployment options include:

- Deploying multiple endpoints
- Multiple domain support
- Deploying a server behind a NAT
- Support for OTP and virtual smart cards
- Support for NIC Teaming
- Off-premise provisioning
- Getting Started Wizard

DirectAccess server deployment options in Windows Server 2012 R2 include:

- Deploying multiple endpoints. When you implement DirectAccess on multiple servers in different network locations, the DirectAccess client computer that is running the Windows 8 operating system selects the closest endpoint automatically. However, for DirectAccess client computers running Windows 7, you must specify the endpoint manually.

Note: The closest point is determined based on the roundtrip time for the connection to the server.

- Multiple domains and multiple forest support. Organizations that have complex multidomain or multiforest infrastructure can deploy DirectAccess servers in multiple domains or forests. In this scenario, DirectAccess client computers can connect to DirectAccess servers that are located in different domains or forests.
- Deploying a DirectAccess server behind a network address translation (NAT). You can deploy a DirectAccess server behind a NAT device, by using a single or multiple interfaces, which removes the prerequisite for a public address. In this configuration, only IP-HTTPS deploys, which allows a secure IP tunnel to be established by using a secure HTTP connection.
- Support for one-time password (OTP) and virtual smart cards. DirectAccess supports OTP authentication, where users are authenticated by providing a combination of user name, password, and an OTP. This feature requires a PKI deployment. In addition, DirectAccess can use the Trusted Platform Module (TPM)–based virtual smart card, which uses the TPM of a client computer to act as a virtual smart card for two-factor authentication.
- Offload network adapters with support for network interface card (NIC) Teaming. NIC Teaming in Windows Server 2012 is fully supported, and does not require non-Microsoft drivers. DirectAccess servers support NIC Teaming. This capability allows DirectAccess client computers to benefit from bandwidth aggregation on the network cards, and failover capability in case one of the network cards is not working.
- Off-premise provisioning. With the new Djoin.exe tool, you can easily provision a non-domain computer with an Active Directory binary large object so that the computer can join a domain without being connected to the internal network. After you join the computer to the domain, it can access the intranet resources by using DirectAccess.
- DirectAccess Getting Started Wizard. The DirectAccess Getting Started Wizard provides a quick way to configure DirectAccess for basic remote access scenarios.

Demonstration: Running the Getting Started Wizard

This demonstration shows how to:

- Create a security group in Active Directory Domain Services (AD DS) for DirectAccess client computers.
- Configure DirectAccess.

Demonstration Steps

Create a security group in AD DS for DirectAccess client computers

1. On LON-DC1, open the Active Directory Users and Computers console, and create an organizational unit (OU) with the name **DA_Clients OU**.
2. In that OU, create a global security group with the name **DA_Clients**.
3. Add LON-CL1 to the **DA_Clients** security group.

Configure DirectAccess

1. Switch to LON-RTR.
2. On LON-RTR, change the IPv4 address for Ethernet 2 as follows:
   - IP address: **131.107.0.21**
   - Subnet mask: **255.255.0.0**

3. Restart LON-RTR, and sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
4. On LON-RTR, in the Server Manager console, select **Remote Access Management**.
5. In the Remote Access Management console, start the **Run the Getting Started Wizard**.
6. Complete the wizard using the following settings:
   - Configure Remote Access: **Deploy DirectAccess only**
   - Verify that Edge is selected
   - Type the public name or IPv4 address used by clients to connect to Remote Access server: **131.107.0.21**
   - Remote Access Review page: next to **Remote Clients**, click **Change**
   - Domain Computers: **DA_Clients**
   - Clear **Enable DirectAccess for mobile computers only**
   - Network Connectivity Assistant: click **Finish**
   - Remote Access review page: click **OK**
   - Configure Remote Access: click **Finish**
7. In the **Applying Getting Started Wizard Settings** dialog box, click **Close**.

When you finish the demonstration, revert all virtual machines to their initial state. To do this, perform the following steps:

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machines** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for the following machines: **20413C-LON-RTR**.

Configuration Changes in the Getting Started Wizard

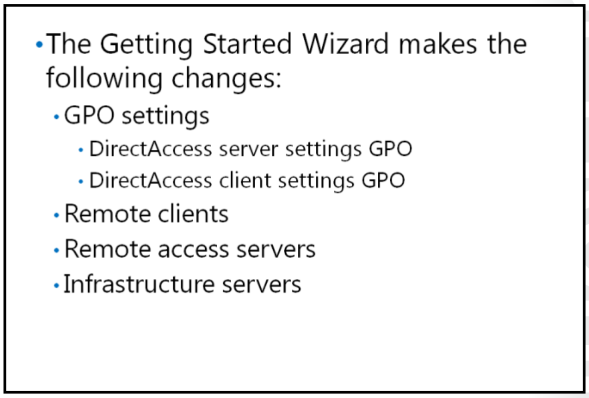

- The Getting Started Wizard makes the following changes:
  - GPO settings
    - DirectAccess server settings GPO
    - DirectAccess client settings GPO
  - Remote clients
  - Remote access servers
  - Infrastructure servers

The Getting Started Wizard makes multiple configuration changes so that DirectAccess clients can connect to the intranet. These changes include:

- GPO settings. The wizard creates two GPOs to define which computers will be allowed to connect to the corporate network by using DirectAccess:
  - DirectAccess server settings GPO. Defines settings that will apply to DirectAccess servers. This GPO contains Windows Firewall with Advanced Security connection security rules that are used to establish connections between DirectAccess clients and the server.

  - DirectAccess client settings GPO. Defines settings that will apply to DirectAccess clients. This GPO contains Windows Firewall with Advanced Security connection security rules that are used to establish connections between DirectAccess clients and the server. This GPO configures all clients to trust the self-signed certificate created by the DirectAccess server. It also contains NRPT entries that the client uses to determine:
    - Which DNS server to use to access hosts while connected to either the intranet or the Internet.
    - The IPv6 transition technology settings required for the connection.

Note: The IPv6 transition technologies that DirectAccess uses will be discussed in more detail later in this lesson.

- Remote clients. In the Getting Started Wizard, you can configure the following client computer settings for DirectAccess:
  - Select groups. You can select which groups of client computers to configure for DirectAccess. By default, the Domain Computers group will be configured for DirectAccess. In the wizard, you can edit this setting and replace the Domain Computers group with a custom security group.
  - Enable DirectAccess for mobile computers only. This setting is enabled by default, and it can be disabled in the wizard.
  - Network Connectivity Assistant. The Network Connectivity Assistant runs on every client computer and provides DirectAccess connectivity information, diagnostics, and remediation support.
  - Resources that validate connectivity to an internal network. DirectAccess client computers need information that will help them decide whether they are located on an intranet or on the Internet. Therefore, they will contact the resources you provide in the Getting Started Wizard. You can provide a URL that an HTTP request will access, or you can provide or fully qualified domain name (FQDN) that will be contacted by a **ping** command. By default, this setting is not configured.
  - Helpdesk email address. By default, this setting is not configured.
  - DirectAccess connection name. The default name for this setting is Workplace Connection.
  - Allow DirectAccess clients to use local name resolution. This setting is disabled by default.
- Remote access servers. In the Getting Started Wizard, you define the network topology where the DirectAccess server is located:
  - On an edge of the internal corporate network, where the edge server has two network adapters
  - On a server located behind an edge device, where the server has two network adapters
  - On a server located behind an edge device, where the server has one network adapter
- Infrastructure servers. In the Getting Started Wizard, you define infrastructure servers. DirectAccess clients connect to these servers before they connect to internal corporate resources. By default, two entries are configured: the domain name suffix, and the DirectAccess-NLS name followed by the domain name suffix. For example, if the domain name is contoso.com, then the following entries are configured by default: contoso.com and DirectAccess-NLS.contoso.com.

How a Simple DirectAccess Deployment Works

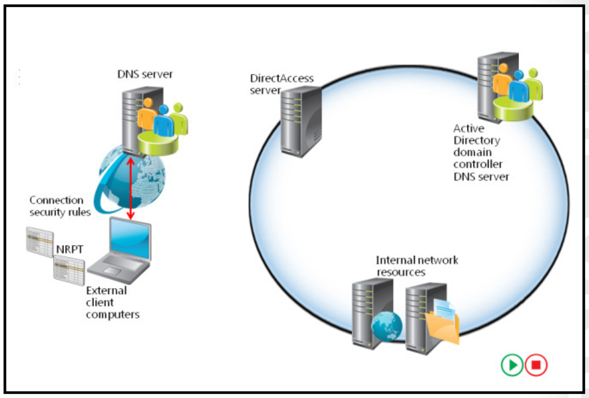
DNS server
DirectAccess server
Active Directory domain controller DNS server
Connection security rules
NRPT
External client computers
Internal network resources

If a DirectAccess client cannot reach the URL address specified for the network location server, the DirectAccess client assumes that it is not connected to the intranet, and therefore it must be located on the Internet. When the client computer cannot communicate with the network location server, it begins using NRPT and connection security rules. The NRPT has DirectAccess-based rules for name resolution, and connection security rules define DirectAccess IPsec tunnels for communication with intranet resources.

Internet-connected DirectAccess clients use the following process to connect to intranet resources.

1. The DirectAccess client attempts to access the network location server.
2. The client attempts to locate a domain controller.
3. The client attempts to access intranet resources first, and then attempts to access Internet resources.

How do DirectAccess Clients Attempt to Access the Network Location Server?

The process for the DirectAccess client attempting to access the network location server is as follows:

1. The client tries to resolve the FQDN of the network location server URL. Because the FQDN of the network location server URL corresponds to an exemption rule in the NRPT, the DirectAccess client does not send the DNS query to a locally configured DNS server, such as an Internet-based DNS server. An external Internet-based DNS server would not be able to resolve the name.
2. The DirectAccess client processes the name resolution request as defined in the DirectAccess exemption rules in the NRPT.
3. Because the network location server is not found on the same network where the DirectAccess client is currently located, the DirectAccess client applies a public or private firewall network profile to the attached network.
4. The Connection Security tunnel rules for DirectAccess from the firewall network profile are applied.

The DirectAccess client uses a combination of NRPT rules and connection security rules to locate and access intranet resources across the Internet through the DirectAccess server.

How do DirectAccess Clients Attempt to Locate a Domain Controller?

After starting up and determining its network location, the DirectAccess client attempts to locate and log on to a domain controller. This process creates an IPsec tunnel (or infrastructure tunnel) by using the IPsec tunnel mode and encapsulating security payload (ESP), to the DirectAccess server. The process is as follows:

1. The DNS name for the domain controller matches the intranet namespace rule in the NRPT, which specifies the IPv6 address of the intranet DNS server. The DNS client service constructs the DNS name query that is addressed to the IPv6 address of the intranet DNS server, and forwards it to the DirectAccess client's TCP/IP stack for sending.
2. Before sending the packet, the TCP/IP stack determines if there are Windows Firewall outgoing rules or connection security rules for the packet.
3. Because the destination IPv6 address in the DNS name query matches a connection security rule that corresponds with the infrastructure tunnel, the DirectAccess client uses Authenticated Internet

Protocol (AuthIP) and IPsec to negotiate and authenticate an encrypted IPsec tunnel to the DirectAccess server. The DirectAccess client, both the computer and the user, authenticates itself with its installed computer certificate and its NTLM credentials, respectively.

Note: AuthIP enhances authentication in IPsec by adding support for user-based authentication with Kerberos version 5 (V5) or SSL certificates. AuthIP also supports efficient protocol negotiation and the use of multiple sets of credentials for authentication.

4. The DirectAccess client sends the DNS name query through the IPsec infrastructure tunnel to the DirectAccess server.
5. The DirectAccess server forwards the DNS name query to the intranet DNS server. The DNS name query response is sent back to the DirectAccess server and back through the IPsec infrastructure tunnel to the DirectAccess client.

Subsequent domain logon traffic goes through the IPsec infrastructure tunnel. When the user on the DirectAccess client signs in, the domain logon traffic goes through the IPsec infrastructure tunnel.

How do DirectAccess Clients Attempt to Access Intranet Resources?

The first time that the DirectAccess client sends traffic to an intranet location (such as an email server) that is not on the list of destinations for the infrastructure tunnel, the following process occurs:

1. The application or process that attempts to communicate constructs a message or payload, and hands it off to the TCP/IP stack for sending.
2. Prior to sending the packet, the TCP/IP stack checks to determine if there are Windows Firewall outgoing rules or connection security rules for the packet.
3. Because the destination IPv6 address matches the connection security rule that corresponds with the intranet tunnel that specifies the IPv6 address space of the entire intranet, the DirectAccess client uses AuthIP and IPsec to negotiate and authenticate an additional IPsec tunnel to the DirectAccess server. The DirectAccess client authenticates itself with its installed computer certificate and the user account's Kerberos authentication credentials.
4. The DirectAccess client sends the packet through the intranet tunnel to the DirectAccess server.
5. The DirectAccess server forwards the packet to the intranet resources. The response is sent back to the DirectAccess server and back through the intranet tunnel to the DirectAccess client.

Any subsequent intranet traffic that does not match an intranet destination in the infrastructure tunnel connection security rule goes through the intranet tunnel.

How do DirectAccess Clients Attempt to Access Internet Resources?

When the user or a process on the DirectAccess client attempts to access an Internet resource (such as an Internet web server), the following process occurs:

1. The DNS client service passes the DNS name for the Internet resource through the NRPT. There will be no matches. The DNS client service constructs the DNS name query that is addressed to the IP address of an interface-configured Internet DNS server, and hands it off to the TCP/IP stack for sending.
2. Prior to sending the packet, the TCP/IP stack checks to determine if there are Windows Firewall outgoing rules or connection security rules for the packet.
3. Because the destination IP address in the DNS name query does not match the connection security rules for the tunnels to the DirectAccess server, the DirectAccess client sends the DNS name query normally.

4. The Internet DNS server responds with the IP address of the Internet resource.
5. The user application or process creates the first packet to send to the Internet resource. Prior to sending the packet, the TCP/IP stack checks to determine if there are Windows Firewall outgoing rules or connection security rules for the packet.
6. Because the destination IP address in the DNS name query does not match the connection security rules for the tunnels to the DirectAccess server, the DirectAccess client sends the packet normally.

Any subsequent Internet resource traffic that does not match a destination either in the infrastructure intranet tunnel or in the connection security rules, sends and receives normally.

The process of accessing the domain controller and intranet resources is very similar to the connection process, because both of these processes use NRPT tables to locate an appropriate DNS server to resolve the name queries. However, the main difference is in the IPsec tunnel that is established between the client and DirectAccess server. When accessing the domain controller, all the DNS queries are sent through the IPsec infrastructure tunnel, and when accessing intranet resources, a second IPsec tunnel is established to access intranet resources.

Limitations of the Getting Started Wizard

The Getting Started Wizard is easy to implement, but it is not suitable for every large deployment. Specifically, it is not suitable for large deployments that:

- Need to support multisite access.
- Require a highly available infrastructure.
- Require support for Windows 7 computers in a DirectAccess scenario.

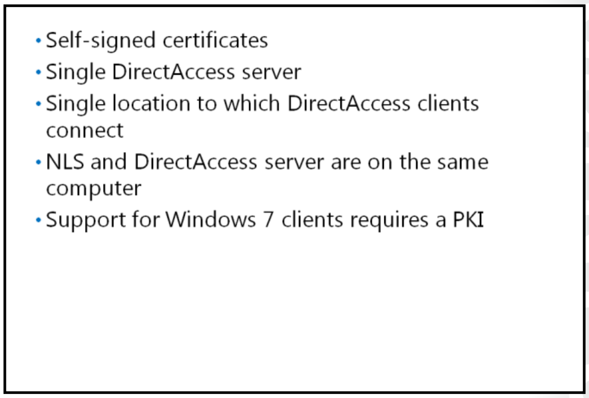

- Self-signed certificates
- Single DirectAccess server
- Single location to which DirectAccess clients connect
- NLS and DirectAccess server are on the same computer
- Support for Windows 7 clients requires a PKI

Self-Signed Certificates

The Getting Started Wizard creates a self-signed certificate to enable SSL connections to the DirectAccess and network location servers. For DirectAccess to function, you need to ensure that the certificate revocation list (CRL) distribution point for both certificates is available externally. In addition, you cannot use the self-signed certificate in multisite deployments.

Note: The CRL contains all revoked certificates and the reasons for revocation.

Because of these limitations, most organizations will either configure a public certificate for DirectAccess and network location server, or provide certificates generated by an internal certification authority (CA).

Those organizations that have implemented an internal CA can use the Web Server certificate template to issue a certificate to the DirectAccess and network location servers. The organization must also ensure that CRL distribution points are accessible from the Internet.

Network Location Server Design

The network location server is a critical part of a DirectAccess deployment. The Getting Started Wizard deploys the network location server on the same server as the DirectAccess server. If DirectAccess client computers on the intranet cannot successfully locate and access the secure webpage on the network location server, they might not be able to access intranet resources.

When DirectAccess clients obtain a physical connection to the intranet or experience a network status change on the intranet (such as an address change when roaming between subnets), they attempt an HTTPS connection to the network location server URL. If the client can establish an HTTPS connection to network location server and then verify the revocation status for the Web server's certificate, the client determines that it is on the intranet. As a result, the NRPT will be disabled on the client, and Windows Firewall will be configured to use the Domain profile with no IPsec tunnels.

You must deploy the network location server on a highly available, high-capacity intranet web server. Larger companies should consider implementing the network location server on a Network Load Balancing (NLB) cluster or by using external hardware balancer.

Support for Windows 7

The Getting Started Wizard configures the remote access server to act as a Kerberos authentication proxy to perform IPsec authentication without requiring certificates. Client authentication requests are sent to a Kerberos proxy service that is running on the DirectAccess server. The Kerberos proxy then sends Kerberos requests to domain controllers on behalf of the client. This configuration is only applicable for clients running the following Windows operating systems: Windows 8, Windows 8.1, Windows Server 2012, or Windows Server 2012 R2. If you need to support DirectAccess on Windows 7 clients, you must deploy a PKI to issue computer certificates for backward compatibility.

Planning the DirectAccess Tunnelling Protocol

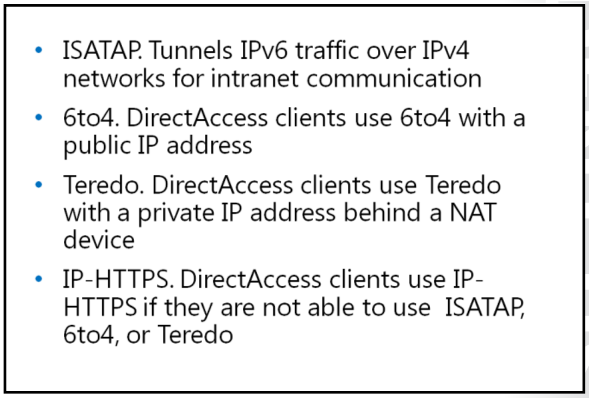

- ISATAP. Tunnels IPv6 traffic over IPv4 networks for intranet communication
- 6to4. DirectAccess clients use 6to4 with a public IP address
- Teredo. DirectAccess clients use Teredo with a private IP address behind a NAT device
- IP-HTTPS. DirectAccess clients use IP-HTTPS if they are not able to use ISATAP, 6to4, or Teredo

DirectAccess uses IPv6 and IPsec when clients connect to internal resources. However, many organizations do not have native IPv6 infrastructure. Therefore, DirectAccess uses transitioning tunneling technologies to connect IPv6 clients with IPv4 internal resources, and to communicate through IPv4-based Internet.

DirectAccess tunneling protocols include:

- ISATAP. You use ISATAP when the DirectAccess client requires access to IPv6 hosts in the internal network, and you are using an ISATAP router for IPv6 connectivity. ISATAP enables DirectAccess clients to connect to the DirectAccess server over IPv4 networks for intranet communication. By using ISATAP, an IPv4 network emulates a logical IPv6 subnet to other ISATAP hosts, where ISATAP hosts automatically tunnel to each other for IPv6 connectivity. Windows Vista®, Windows Server 2008, and newer Windows client and server operating systems can act as ISATAP hosts. ISATAP does not require any changes on IPv4 routers because IPv6 packets are tunneled within an IPv4 header. To use ISATAP, you must configure DNS servers to answer ISATAP queries, and you must enable IPv6 on network hosts.
- 6to4. 6to4 enables DirectAccess clients to connect to the DirectAccess server over the IPv4-based Internet. You can use 6to4 when clients have a public IPv4 Internet address. IPv6 packets are encapsulated in an IPv4 header, and then sent over the 6to4 tunnel adapter to the DirectAccess server. You can configure the 6to4 tunnel adapter for DirectAccess clients and the DirectAccess server by using a GPO. 6to4 does not work if clients are located behind an IPv4 NAT device. 6to4 uses the IP protocol on port 41 for communication.
- Teredo. Teredo enables DirectAccess clients to connect to the DirectAccess server across the IPv4 Internet, when clients are located behind an IPv4 NAT device. In this scenario, you should configure the firewall to allow outbound traffic on User Datagram Protocol (UDP) port 3544. Clients that have a

private IPv4 address use Teredo to encapsulate IPv6 packets in an IPv4 header, and send them over the IPv4-based Internet. You can configure Teredo for both DirectAccess clients and the DirectAccess server by using a GPO.

- IP-HTTPS. IP-HTTPS enables DirectAccess clients to connect to the DirectAccess server over the IPv4-based Internet. -Clients use IP-HTTPS when they are unable to connect to the DirectAccess server by using ISATAP, 6to4, or Teredo. You can configure IP-HTTPS for DirectAccess clients and the DirectAccess server by using Group Policy.

Note: You can enable and disable the DirectAccess tunneling protocol by using either a GPO or the **netsh** interface command. To use GPOs, navigate to Computer Configuration/Policies/Administrative Templates/Network/TCPIP Settings/IPv6 Transition Technologies, and then configure the setting. To use the **netsh** command, type the following command: **netsh interface <*protocol*> set state <*enabled or disabled*>**. For example, to enable ISATAP, you would use the following command: **netsh interface isatap set state enabled**.

Planning PKI Integration with DirectAccess

While planning for DirectAccess implementation, organizations can choose to use a private CA or a public CA. If an organization has already deployed an internal PKI infrastructure that they are using for different purposes (such as user or server authentication), the organization can customize the current PKI infrastructure further to enhance the DirectAccess deployment. You will require an SSL certificate for both the location server and the DirectAccess server. Additionally, you may need certificates for DirectAccess clients if you choose to use certificate-based authentication for clients or if you have Windows 7 clients. Using a public CA for client certificates can be expensive. Therefore, in this scenario, you should maintain your own PKI.

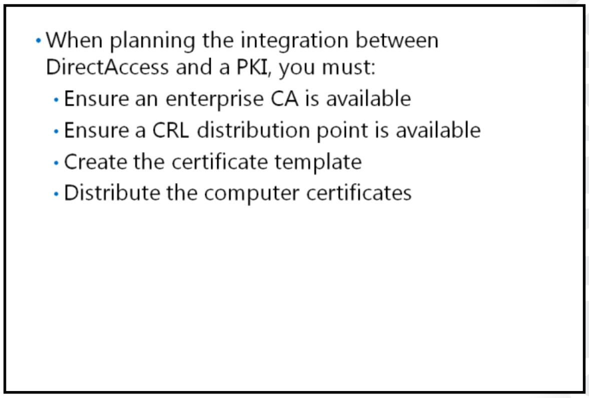

- When planning the integration between DirectAccess and a PKI, you must:
  - Ensure an enterprise CA is available
  - Ensure a CRL distribution point is available
  - Create the certificate template
  - Distribute the computer certificates

When planning the integration between DirectAccess and a PKI, you must:

- Ensure an enterprise CA is present. You can use enterprise CAs to create certificate templates that you will use later for issuing certificates manually or by GPOs.
- Ensure a CRL distribution point is available. Clients check CRL distribution points to ensure that the certificate provided by a DirectAccess server is still valid. The CRL must be available both in the intranet and on the Internet.
- Create the certificate template. DirectAccess uses a web server certificate for the network location server and for IP-HTTPS connections on the DirectAccess server. DirectAccess uses computer certificates for DirectAccess clients and servers to be used for IPsec.
- Distribute the computer certificates. The most efficient way to distribute computer certificates is by using Group Policy. Certificate distribution is extremely important for clients that use certificate-based authentication. If you do not use a GPO for autoenrollment, you would need to manually distribute and install the certificates on each computer.

Planning Network Configuration Options

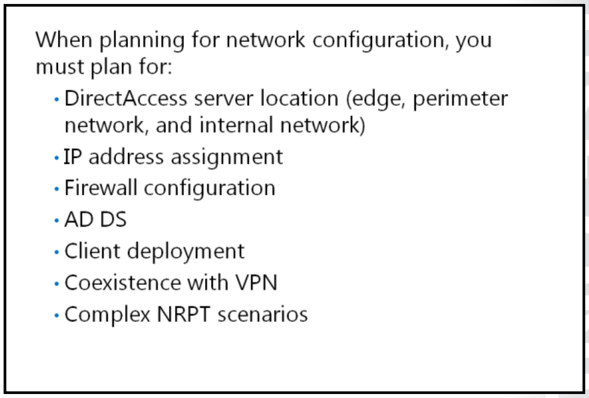

When planning for network configuration, you must plan for:

- DirectAccess server location (edge, perimeter network, and internal network)
- IP address assignment
- Firewall configuration
- AD DS
- Client deployment
- Coexistence with VPN
- Complex NRPT scenarios

Depending on your organization's business requirements, you can configure multiple network topologies when you deploy an advanced DirectAccess infrastructure.

Consider the following aspects when planning for internal network configuration:

- Plan for DirectAccess server location. You can install the DirectAccess server in different network configurations:
  - Edge. For this configuration, you install the DirectAccess server role service on a computer that acts as an edge server. An edge server also acts as a firewall. The edge server has two network adapters: one that connects to the Internet, and the other that connects to the internal network.
  - Behind an edge device with two network adapters. In this configuration, you install the DirectAccess role service on a computer that is located on a perimeter network behind an edge device. The DirectAccess server has two network adapters, one that connects to the perimeter network, and the other that connects to the internal network.
  - Behind an edge device with one network adapter. This configuration assumes that you installed the DirectAccess role service on a computer that is located on the internal network.
- Plan the IP address assignment. You should plan your IP addressing based on whether your organization has deployed native IPv6 addressing, both IPv6 and IPv4 addressing, or only IPv4 addressing. In a scenario where both the Internet and intranet IP addressing is IPv4, you must configure the external network adapter of the DirectAccess server with two consecutive public IPv4 addresses. The Teredo tunneling protocol requires this configuration, because the DirectAccess server will act as a Teredo server.
- Plan the firewall configuration. The DirectAccess server requires specific ports to be open on the corporate firewall so that the DirectAccess client computers can connect from the Internet to the internal network. The following firewall rules are required for DirectAccess on an IPv4 network:
  - Teredo traffic. UDP destination port 3544 inbound and UDP source port 3544 outbound.
  - 6to4 traffic. IP Protocol 41 inbound and outbound.
  - IP-HTTPS. TCP destination port 443 and TCP source port 443 outbound.
  - For scenarios where DirectAccess and a network location server are installed on the same server with a single adapter, TCP port 62000 on the server should be open.
- Plan for AD DS. DirectAccess requires that you install at least one domain controller on a server running Windows Server 2008 R2 or a newer Windows Server operating system. The computer where you install the DirectAccess role service should be a domain member. The DirectAccess client computers also have to be domain members. DirectAccess clients can establish a connection from the Internet with any domain in the same forest as the DirectAccess server, and with any domain that has a two-way trust with the DirectAccess server forest.

- Plan for client deployment. Before deploying clients, you should configure the following settings:
  - Create a security group for DirectAccess client computers, and configure the group membership.
  - Configure DirectAccess to be available either for all computers in the domain, or only for mobile computers.
  - Configure the Network Connectivity Assistant.
- Plan for DirectAccess and VPN server coexistence. If you use a single server to host DirectAccess and VPN connections, you must ensure that the network policy rules that you create for VPN connections will not block DirectAccess connections.
- Plan NRPT changes. Consider existing NRPT settings before deploying DirectAccess. You need to ensure that existing NRPT rules do not conflict with the NRPT rules used for DirectAccess clients.

Planning Network Location Servers

The network location server which hosts the network location server website resides on the DirectAccess server or on another server in your organization. If the network location server website resides on the DirectAccess server, the website is created automatically when you deploy DirectAccess. If the network location server website resides on another computer that is running the Windows Server operating system, you must manually install Internet Information Services (IIS) on that computer, and configure the network location server website.

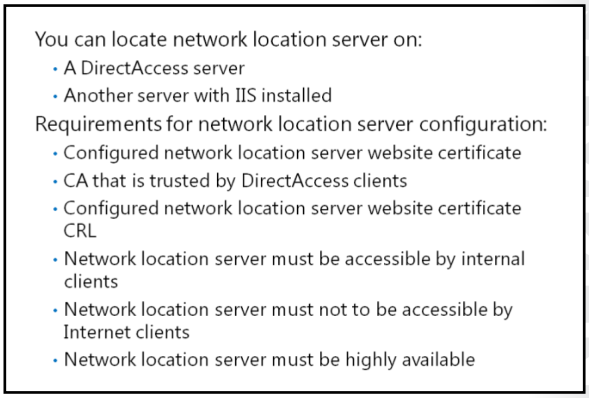

You can locate network location server on:

- A DirectAccess server
- Another server with IIS installed

Requirements for network location server configuration:

- Configured network location server website certificate
- CA that is trusted by DirectAccess clients
- Configured network location server website certificate CRL
- Network location server must be accessible by internal clients
- Network location server must not to be accessible by Internet clients
- Network location server must be highly available

You should configure the network location server to meet the following requirements:

- An HTTPS server certificate is configured for the network location server website.
- A CA that issues the HTTPS certificate for the network location server website should be trusted by the DirectAccess client computers.
- The server certificate of the network location server website must be checked against a CRL.
- The DirectAccess client computers on the internal network must be able to resolve the name of the network location server.
- The network location server should not be accessible to DirectAccess client computers on the Internet.
- If DirectAccess is business-critical for the organization, you should configure the network location server with high availability for computers located on the internal network.

Lesson 2
Planning and Implementing VPN

When your users require remote access to your organization's network resources and applications, you must determine how to provide this access safely. Although DirectAccess can provide the necessary access, only clients that are running Windows 8.1, Windows 8, or Windows 7 can use it. If you have other client operating systems, you should have to implement VPN access. As part of your VPN infrastructure design, you must consider which VPN tunneling protocol to implement, then select the appropriate hardware, and then determine the suitable location to place the VPN.

Lesson Objectives

After completing this lesson, you will be able to:

- Determine what factors are important in selecting a network access design.
- Select a suitable VPN tunneling protocol.
- Select a suitable authentication and encryption method.
- Plan a VPN services strategy.
- Plan and configure network policies.
- Design a suitable remote access services strategy.
- Implement a VPN, and configure network access policies.
- Plan client connectivity to VPNs.

Considerations for Implementing VPNs

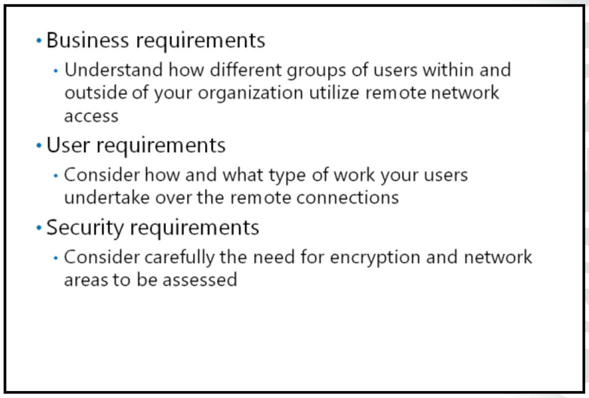

- Business requirements
  - Understand how different groups of users within and outside of your organization utilize remote network access
- User requirements
  - Consider how and what type of work your users undertake over the remote connections
- Security requirements
  - Consider carefully the need for encryption and network areas to be assessed

Most organizations are more likely to use DirectAccess as a solution for providing remote access to corporate resources. However, other scenarios still exist where a traditional VPN solution is more beneficial.

Typical VPN Scenarios

The main limitation of DirectAccess is that it is a client connection technology, and even as a client technology, it is restricted to Windows 8.1, Windows 8, and Windows 7 clients. Because of this, you may still require VPN for the following scenarios:

- Site-to-site connectivity. You still need to configure VPN connections to connect two sites over the Internet.
- Legacy clients. Client computers that do not run Windows 8.1, Windows 8, or Windows 7, and that require remote access to corporate resources will need to connect by using standard VPN technologies.
- Integration with public cloud services. Public cloud services such as Windows Azure™ are commonly used as an extension to the corporate environment. To access the virtual machines that are hosted in

this environment, you need a secure communication tunnel. You can use a site-to-site VPN to establish communication between the corporate network and the public cloud.

When designing network access, you must first gather information about your business needs and your users' requirements. You must then determine your security needs based on your business and user needs assessment.

Business Requirements

Business requirements for your organization include understanding how different groups of users (both within and outside of your organization) utilize remote network access. Requirements that you should consider include:

- Internal staff. Various members of your organization may need access to data from non-traditional locations. Examples of internal staff that you need to consider include traveling executives, sales people visiting clients, users with mobile devices, and remote users with wireless connections.
- External users. You may need to provide access for external users to access data or applications in your organization. External users can include partners, suppliers, or vendors.
- The data being accessed. The type of access that you provide should be determined in part by the type of data that users need to access. If users need to access an application remotely, Remote Desktop Services (RDS) may be an effective solution. If users need to access a small volume of data, a secure website may be the most appropriate solution.

User Requirements

When determining user requirements, you must consider how and what type of work your users undertake over remote connections. You should consider the following:

- Tasks performed by internal staff. The type of access that you provide may vary depending on the tasks being performed by internal staff. For example, applications requiring fast data connectivity may not be appropriate for VPN connections.
- Tasks performed by external users. The type of access that you decide to provide may vary depending on the tasks being performed by external users. For example, if external users perform data entry into a web-based application, you may consider using SSL to provide sufficient security.
- The duration of the connection. The duration of a user connection is important because connection duration is limited for some network access methods. You must ensure that your network access method allows connections of the required duration to prevent any problems for users. For example, if external users connect to a VPN for up to 12 hours at a time, you should ensure that the VPN server allows connections for at least 12 hours.
- The number of concurrent users. You must design your network access solution to support the maximum number of concurrent users that you expect, and to allocate for growth. For example, if you expect 150 concurrent VPN connections, you must ensure that your Internet connection and VPN server can manage that capacity, and more.

Security Requirements

After determining and planning your network access, you must now consider carefully how to provide this access in a secure manner. When determining security requirements, you need to consider the following:

- Types of client computers. Different operating systems provide different security capabilities. By knowing the types of client computers in use within your organization, you can determine if you require a non-Microsoft solution for additional security.
- Need for encryption. In general, data that travels over a public network should be encrypted. Depending on how sensitive the data is, you may also need to encrypt it over your internal networks.

- Network areas to be accessed. As part of a security design, you must segment your network into security zones. You can restrict specific types of clients to specific security zones to mitigate the risk of an unauthorized user gaining access to data. For example, you might give VPN users access to only a subset of servers in your organization, and make sensitive data unavailable when attempting access through a VPN connection.

Best Practices

When you perform your initial needs analysis and security assessment for your remote network access infrastructure design, consider the following guidelines:

- Determine and document any existing network access infrastructure. Gather information about the number and types of users, the types of connections used, and the duration of connections. In addition, gather information about the types of servers and client computers used, and the connection protocols used.
- Interview the following groups for design input:
  - Administrators
  - Users
  - Application owners
  - Security groups
  - Management
  - Any other groups that will be part of the network design

  Failure to gather information from a key group might cause your design to fail.
- Collect data on the specific security needs for your organization. Your remote access design should balance the business needs for accessing data with the security needs of restricting access to data. Failure to gather security information might lead to a design that has security flaws and is susceptible to unauthorized access. Review existing compliance regulations to ensure that your implementation follows company guidelines.

Gather information about your organization's future needs. Most network access designs take into consideration anticipated growth over the next three to five years. This will ensure that your design will remain useful over time, and will avoid an expensive redesign in the near future.

Selecting a Suitable Tunnelling Protocol

If you decide to implement a VPN solution, you must choose the VPN tunneling protocol that you will use. Windows Server 2012 R2 and Windows Server 2012 support four VPN tunneling protocols:

- Point-to-Point Tunneling Protocol (PPTP)
- Layer Two Tunneling Protocol (L2TP) over IPsec
- Secure Socket Tunneling Protocol (SSTP)
- Internet Key Exchange version 2 (IKEv2)

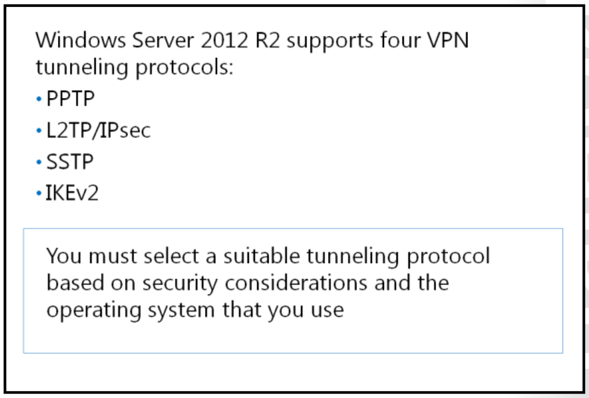

Windows Server 2012 R2 supports four VPN tunneling protocols:

- PPTP
- L2TP/IPsec
- SSTP
- IKEv2

You must select a suitable tunneling protocol based on security considerations and the operating system that you use

PPTP

PPTP enables you to encrypt and encapsulate multiprotocol traffic that is then sent across either a private IP network, or across a public IP network such as the Internet. You can use PPTP for site-to-site and VPN connections. When you use the Internet as the VPN public network, the PPTP server is a VPN server that is PPTP-enabled and that has two interfaces: one on the Internet, and the other on the intranet.

PPTP uses Microsoft Point-to-Point Encryption to encrypt PPP frames. For this encryption, PPTP uses encryption keys that are generated from the authentication process of Microsoft Challenge Handshake Authentication Protocol version 2 (MSCHAPv2) or Extensible Authentication Protocol-Transport Layer Security (EAP-TLS). VPN clients must use the MSCHAPv2 or EAP-TLS authentication protocol so that the payloads of frames can be encrypted.

To implement PPTP, you must configure your firewall to pass both TCP Port 1723 and IP Protocol ID 47.

L2TP

L2TP is a combination of PPTP and Layer 2 Forwarding, and contains the best features of them both. You can use L2TP to encrypt multiprotocol traffic that you want to send over any medium supporting point-to-point datagram delivery, such as IP or asynchronous transfer mode (ATM).

For encryption, the Microsoft implementation of L2TP does not use Microsoft Point-to-Point Encryption; instead, it uses IPsec in Transport Mode. This method is called L2TP/IPsec. The encryption keys generated from the IKE negotiation process encrypt the L2TP message by using one of the following protocols: Advanced Encryption Standard (AES) 256, AES 192, AES 128, and Triple Data Encryption Standard (3DES) encryption algorithms.

To implement L2TP, you must configure your firewall to pass UDP Port 500, UDP Port 1701, UDP Port 4500, and IP Protocol ID 50.

SSTP

You can use SSTP to allow PPTP and L2TP/IPsec traffic through firewalls and web proxies. SSTP is a tunneling protocol that uses the HTTPS protocol over Transmission Control Protocol (TCP) port 443. SSTP includes a method to encapsulate PPP traffic over the SSL channel of the HTTPS protocol. Because SSTP uses Point-to-Point Protocol (PPP), strong authentication methods such as EAP-TLS are possible. Transport-level security with enhanced key negotiation, encryption, and integrity checking also is possible with SSTP.

Note: Because of SSTP’s reliance on HTTPS, you can be reasonably sure that you can initiate a VPN connection by using SSTP where other tunnel types are not permitted. For example, at a hotel where the firewall only allows HTTP and HTTPS traffic to pass, you could still use an SSTP-based tunnel.

Note: PPTP, L2TP, and SSTP depend heavily on the features originally specified for PPP. PPP was designed originally to send data across dial-up or dedicated point-to-point connections. For IP, PPP encapsulates IP packets within PPP frames, and then transmits the encapsulated PPP packets across a point-to-point link..

IKEv2

For computers running Windows 8.1, Windows 8, Windows 7, Windows Server 2012 R2, Windows Server 2012, or Windows Server 2008 R2, you can use IKEv2. IKEv2 uses the IPsec tunnel mode protocol over UDP port 500. Because of its support for mobility, IKEv2 is much more resilient to changing network connectivity. This makes it a good choice for mobile users who move between access points and even switch between wired and wireless connections.

This ability is a requirement of *VPN Reconnect*, a feature of Windows 8 and Windows 7.

Note: VPN Reconnect uses IKEv2 technology to provide seamless and consistent VPN connectivity. VPN Reconnect reestablishes a VPN connection automatically when Internet connectivity becomes available again. Users who connect with a wireless mobile broadband benefit most from this capability.

Using IKEv2 and IPsec enables support for strong authentication and encryption methods.

Note: You can configure the maximum allowed network outage time for an IKEv2 VPN connection. By default this value is set to 30 minutes, but can range from five minutes to eight hours.

Comparing the Tunneling Protocols

When choosing between PPTP, L2TP/IPsec, SSTP, and IKEv2 remote access VPN solutions, consider the following:

- You can use PPTP with a variety of Microsoft clients, including:
  - Windows 8.1
  - Windows 8
  - Windows 7
  - Windows Vista
  - Windows XP
  - Windows Server 2012 R2
  - Windows Server 2012
  - Windows Server 2008 R2
  - Windows Server 2008
  - Windows Sever 2003 R2
  - Windows Server 2003
  - Windows 2000 Server

  Unlike L2TP/IPsec, PPTP does not require a PKI, and instead uses encryption keys. PPTP encrypts data, and any captured packets require the encryption key to be deciphered. Thus, PPTP-based VPN connections provide data confidentiality. However, PPTP-based VPN connections do not provide:

  - Data integrity, which confirms that the data was not modified in transit.
  - Data-origin authentication, which confirms that the data was sent by the authorized user.
- You can use L2TP only with client computers that are running the following Windows operating systems:
  - Windows 8.1
  - Windows 8
  - Windows 7
  - Windows Vista
  - Windows XP
  - Windows 2000 Server

L2TP supports computer certificates or a preshared key as the IPsec authentication method. Authentication by using computer certificates is the recommended authentication method. In this method, a PKI issues computer certificates to the VPN server and all VPN client computers. L2TP/IPsec VPN connections use IPsec to provide data confidentiality, integrity, and authentication.

Unlike both PPTP and SSTP, L2TP/IPsec enables computer authentication at the IPsec layer, and user-level authentication at the PPP layer.

- Some older clients have problems using IPsec from behind a NAT device; this problem can preclude the use of L2TP/IPsec in these situations.
- You can use SSTP only with client computers that are running the following Windows operating systems:
  - Windows 8.1
  - Windows 8
  - Windows 7
  - Windows Vista SP1
  - Windows Server 2012 R2
  - Windows Server 2012
  - Windows Server 2008 R2
  - Windows Server 2008

  Using SSL, SSTP VPN connections provide data confidentiality, integrity, and authentication. Unlike both L2TP and PPTP, you do not need to reconfigure your firewall to support SSTP connections. However, unlike PPTP and L2TP, SSTP does not support site-to-site connections. Although SSL requires a certificate, the certificate only needs to be installed on the VPN server. However, VPN clients require that the CA certificate be in their trust list.
- IKEv2 supports VPN Reconnect, which the other protocols do not support. However, remember that IKEv2 is only supported by the following Windows operating systems:
  - Windows 8.1
  - Windows 8
  - Windows 7
  - Windows Server 2012 R2
  - Windows Server 2012
  - Windows Server 2008 R2

Selecting an Authentication and Encryption Method

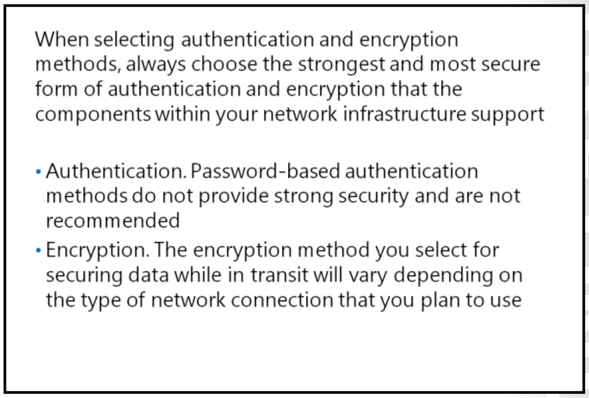

When selecting authentication and encryption methods, always choose the strongest and most secure form of authentication and encryption that the components within your network infrastructure support

- Authentication. Password-based authentication methods do not provide strong security and are not recommended
- Encryption. The encryption method you select for securing data while in transit will vary depending on the type of network connection that you plan to use

Before enabling remote network access to resources and applications within your network, you must consider how you will implement authentication and encryption to meets your security design needs.

Authentication

You can choose from the following authentication methods:

- Password Authentication Protocol (PAP). PAP uses plaintext passwords and is the least secure authentication protocol. It typically is negotiated if the remote access client and remote access server cannot negotiate a more secure form of validation. PAP is included in Windows Server 2012 to support older client operating systems that support no other authentication method.

- CHAP. CHAP is a challenge-response authentication protocol that uses the industry-standard MD5 hashing scheme to encrypt the response. Various vendors of network access servers and clients use CHAP. Because CHAP requires the use of a reversibly encrypted password, you should consider using another authentication protocol, such as MSCHAP v2.

- MSCHAP v2. MSCHAP v2 is a one-way, encrypted password, mutual-authentication process that works as follows:

  1. The authenticator (the remote access server or the computer that is running Network Policy Server (NPS)) sends a challenge to the remote access client. The challenge consists of a session identifier and an arbitrary challenge string.

  2. The remote access client sends a response that contains a one-way encryption of the received challenge string, the peer challenge string, the session identifier, and the user password.

  3. The authenticator checks the response from the client and sends back a response. This response contains an indication of the success or failure of the connection attempt, and an authenticated response based on the sent challenge string, the peer challenge string, the client's encrypted response, and the user password.

  4. The remote access client verifies the authentication response, and if correct, uses the connection. If the authentication response is not correct, the remote access client terminates the connection.

- Extensible Authentication Protocol (EAP). With EAP, an arbitrary authentication mechanism authenticates a remote access connection. The remote access client and the authenticator (either the remote access server or the Remote Authentication Dial-In User Service (RADIUS) server) negotiate the exact authentication scheme to be used. Routing and Remote Access includes support for EAP-TLS by default. You can plug in other EAP modules to the server that is running Routing and Remote Access, to provide other EAP methods.

- Other options. In addition to the previously mentioned authentication methods, two other options that you can enable when selecting an authentication method include:

- Unauthenticated Access. Unauthenticated Access allows remote systems to connect without authentication. As a best practice, you should never enable this option in a production environment, because it puts your network at risk. Nonetheless, this option is sometimes useful for troubleshooting authentication issues in a test environment.

- Machine Certificate for IKEv2. Select this option if you want to use VPN Reconnect.

Encryption

The encryption method that you select for securing data while in transit varies depending on the type of network connection that you plan to use. Windows Server 2012 supports the use of the following:

- Microsoft Point-to-Point Encryption. Microsoft Point-to-Point Encryption uses the RSA public key cipher for encryption and decryption with an RC4 stream cipher to encrypt data for PPP or PPTP connections. PPTP connections use Microsoft Point-to-Point Encryption with MSCHAP, MSCHAP v2, EAP-MD5 Challenge, or EAP-TLS authentication. (Note that we do not recommend the use of EAP-MD5 due to security weaknesses.)
- IPsec. IPsec is used for encryption by L2TP VPN connections. IPsec can perform authentication based on a preshared key, Kerberos authentication, or certificates. (As a best practice, you should base authentication on certificates.) In addition, L2TP will perform user-based authentication with CHAP, MSCHAP, MSCHAP v2, EAP-MD5 Challenge, or EAP-TLS authentication. IPsec also supports IKEv2 VPNs.
- SSL. SSL is used for encryption by SSTP VPN connections. SSL requires a certificate to be installed on the server, but not on the client computers. SSL is firewall-friendly because it is a web protocol.

IPsec and SSL are more secure than Microsoft Point-to-Point Encryption. IPsec provides additional authentication security by requiring computer-based authentication. However, this computer-based authentication requires additional administrative effort.

Best Practices

When selecting authentication and encryption methods and protocols, always select the strongest and most secure form of authentication and encryption that the components within your network infrastructure support. Where security is especially important, consider implementing multifactor authentication. Multifactor authentication uses a combination of components in addition to user names and passwords. This provides more secure access to network resources. For example, deploying a smart card solution is a common way to implement multifactor authentication.

Planning a VPN Services Strategy

When planning a VPN services strategy, you must consider the hardware to use, the placement of VPN servers, and the user environment.

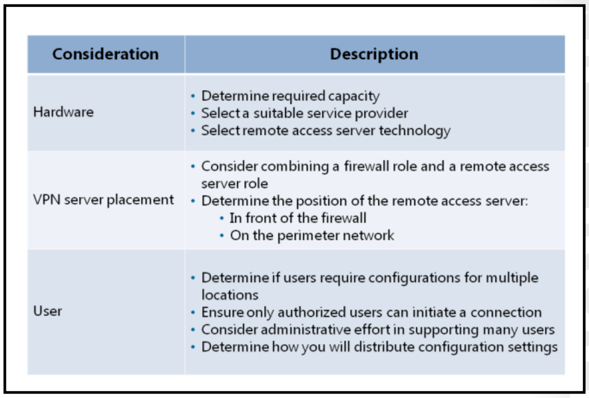

Consideration	Description
Hardware	• Determine required capacity • Select a suitable service provider • Select remote access server technology
VPN server placement	• Consider combining a firewall role and a remote access server role • Determine the position of the remote access server: • In front of the firewall • On the perimeter network
User	• Determine if users require configurations for multiple locations • Ensure only authorized users can initiate a connection • Consider administrative effort in supporting many users • Determine how you will distribute configuration settings

Hardware Considerations

When determining hardware for a remote access solution, you should consider the following factors:

- Your capacity requirement is determined by many factors, including the number of users, what tasks your users perform, from where they are connecting, and what level of security you require. If you estimate a capacity that is inadequate, your remote access infrastructure could slow down user productivity. If you estimate a capacity that is too large, you might end up paying for capacity that you are not using. Provisioning communication links requires you to determine an appropriate connection capacity and installation time. The link capacity must support the maximum number of users expected, and support future growth.

- Select a service provider that can meet your needs for a service level agreement (SLA). Uptime and recovery requirements for outages are a significant part of an SLA. Costs will vary between SLA providers, depending on conditions contained within the SLA.
- Remote access servers can be either dedicated hardware servers, or a software-based solution such as Windows Server 2012 Routing and Remote Access Services (RRAS). Generally, a server-based software solution is more flexible than a dedicated remote access device, because the server on which the remote access software is installed can be used for additional functions in addition to providing remote access.

Considerations for Placing VPN Servers

Consider the following strategies for placing VPN servers:

- Many firewall solutions can also be configured as VPN servers. You can consider this option if your firewall has both the required capabilities and sufficient capacity. It is worth noting though, that the cost of this solution may be higher than using RRAS.
- A VPN server that is positioned in front of your firewall is easier to implement. It also enables you to apply firewall rules to clients that access the internal network. The major drawback to this strategy is that your VPN server will not be protected.
- A VPN server on your perimeter network enables firewall rules to apply to incoming clients, and protects your VPN server. However, configuring firewall rules is more complex than other options, because you must configure the two firewalls that create the perimeter network appropriately. In addition, when you host the VPN server on a perimeter network, it can use an internal IP address with ports redirected from the external firewall.

Considerations for Configuring the User Environment

When determining how you will configure remote access client computers, consider the following:

- Travel. Do the remote access users require access configurations that cover multiple locations? For example, if you have users who travel frequently or need to access your private network from home and other locations, you will need to create multiple connection environments on the client.
- Security. Because the remote access client is the starting point of the remote access request, it is imperative that only authorized users can initiate a request. For example, you should configure a salesperson’s laptop in such a way that only the salesperson can initiate a remote access request.
- Number of remote users. If you have hundreds of users that require remote access configurations, you must consider the amount of administrative overhead that will be involved in setting up each user’s computer for remote access.
- Distribution. How will you design a distribution process to provide each remote access user with the configuration settings that they need? For example, if the user’s computer is at home, is there a website where the user can access the connection configuration? Alternatively, will the user require removable media to configure the settings?

Planning Network Policies

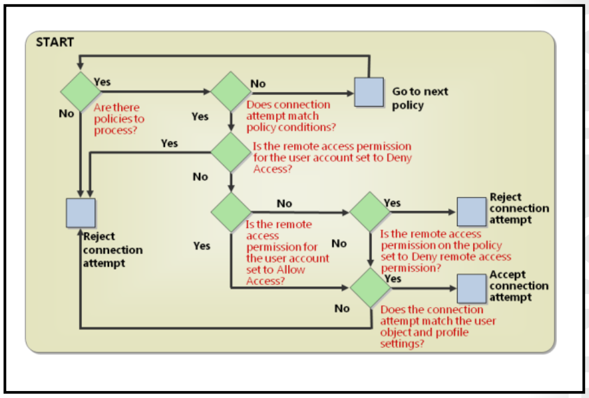
START
Yes
No
Are there policies to process?
No
Does connection attempt match policy conditions?
Go to next policy
Yes
Yes
Is the remote access permission for the user account set to Deny Access?
No
Reject connection attempt
No
Yes
Reject connection attempt
Is the remote access permission for the user account set to Allow Access?
Yes
No
Is the remote access permission on the policy set to Deny remote access permission?
Yes
Accept connection attempt
No
Does the connection attempt match the user object and profile settings?

Network policies are sets of conditions, constraints, and settings that enable you to designate who is authorized to connect to the network, and the circumstances under which they can or cannot connect. As such, network policies determine whether a connection attempt is successful. If the connection attempt is successful, the network policy then defines connection characteristics—such as day and time restrictions, and session idle-disconnect times.

When you deploy NAP, a health policy is added to the network policy configuration so that your NPS performs client health checks during the authorization process.

You can view network policies as rules, with each rule having a set of conditions and settings. NPS compares the rule's conditions to the properties of connection requests. If a match occurs between the rule and the connection request, the settings that you define in the rule are applied to the connection.

When you configure multiple network policies in NPS, they are an ordered set of rules. NPS checks each connection request against the list's first rule, then the second, and so on, until it finds a match.

Note: After NPS discovers a matching rule, it disregards further rules. Therefore, it is important that you order your network policies appropriately.

Each network policy has a Policy State setting that allows you to enable or disable the policy. When you disable a network policy, NPS does not evaluate that policy when authorizing connection requests.

Controlling Network Access with Network Policies

When NPS performs authorization of a connection request, it compares the request with each network policy in the ordered list of policies, starting with the first policy and moving down the list.

If NPS finds a policy in which the conditions match the connection request, NPS uses the matching policy and the dial-in properties of the user account to perform authorization.

If you configure the dial-in properties of the user account to grant or control access through network policy, and if NPS authorizes the connection request, NPS applies the settings that you configure in the network policy to the connection in the following way:

- If NPS does not find a network policy that matches the connection request, NPS rejects the connection unless the dial-in properties on the user account are set to grant access.
- If the dial-in properties of the user account are set to deny access, NPS rejects the connection request.
- The default network policies deny access to all users. This ensures that only users to which you have specifically granted access are allowed access. To allow users access, you create additional network policies with conditions that match authorized users and change the order of the policies to ensure your new policy is listed before the default network policy.

When planning your network policies, consider how you want the constraints and conditions to control the connection from particular groups of users, and then choose appropriate conditions to effect these settings on those users.

For example, suppose you have the following two objectives:

- You want to allow members of the Domain Admins group to connect at any time of the week but insist on an L2TP tunnel-type connection.
- You want all other users to connect with any tunnel type, but only on weekends.

You must consider how to implement two network policies to achieve this objective. If you configure a condition of Any time of the week, in addition to administrators all other user connection attempts will match this condition and subsequently the settings within the policy. Therefore, you might consider creating a condition that looks for membership of the Domain Admins group, and then a constraint of an L2TP tunnel-type. You will then require a second policy to address the needs of all other users.

Discussion: Designing Remote Access

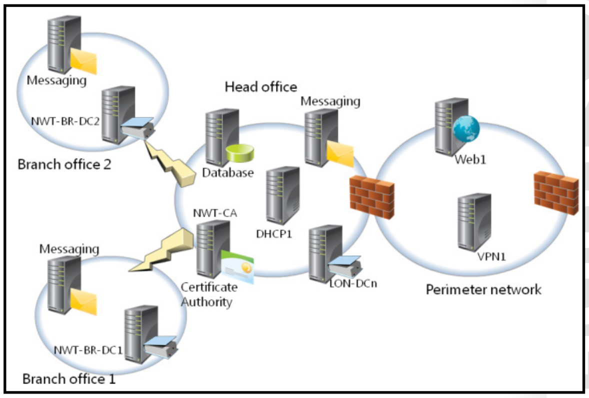

Northwind Traders has three locations, as shown on the slide. You have been given the following information about the Northwind Traders network:

- The perimeter network contains a web server and a VPN server that are running NPS and RRAS.
- The Dynamic Host Configuration Protocol (DHCP) server provides IPv4 configurations to all network clients.
- Exchange Server 2010 provides email via multiple role servers that are deployed to each physical location. Users' mailboxes are stored on their local messaging servers.
- A line-of-business (LOB) application that is built on Microsoft SQL Server® resides on the database server in the head office.
- Domain controllers reside in each physical location, and additionally provide DNS name resolution.
- A CA resides in the head office, and issues private certificates within the Northwind Traders organization.

The Northwind Traders sales force requires access to the database application from anywhere in the world. The sales manager has requested that the IT department modify the network design to support this requirement as soon as possible. In addition, the sales users have requested access to their email mailboxes while roaming.

Use the following requirements to assist in determining an appropriate network access design:

- All sales users have a laptop installed with either Windows 7 or Windows 8.
- Sales users require access to the database application from their customers' sites, their own homes, and from their hotels while traveling.
- The database contains sensitive information, and any network traffic to and from the database must be encrypted.
- It is important that only sales personnel can access the database remotely.
- Sales users need access to their email while traveling.

Consider the current configuration, and then using the information contained within the preceding topics, determine how you would design remote access to support the sales users' needs. To help, consider the following discussion questions.

> **Question:** How would you propose to support the sales users' needs to access their email?
>
> **Question:** What additional network components do you require to support your design, if any?
>
> **Question:** To facilitate database access, which VPN tunnel type would you recommend?

Demonstration: Implementing a VPN

This demonstration shows how to:

- Configure a VPN server.
- Configure a VPN client.
- Create a VPN policy based on the Windows Groups condition.
- Test the VPN.

Demonstration Steps

Configure a VPN server

1. Sign in to LON-RTR as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. In Server Manager, add the **Network Policy and Access Services** role.
3. Register the server in AD DS.
4. Open Routing and Remote Access, and disable the existing configuration.
5. Reconfigure LON-RTR as a VPN server by using the following settings:
   - **Ethernet 2** is the public interface
   - The VPN server allocates addresses from the pool: **172.16.0.100** - **172.16.0.111**
   - The server is configured with the option **No, use Routing and Remote Access to authenticate connection requests**.
6. Start the VPN service.

Configure a VPN client

1. On LON-CL2, create a new VPN connection with the following properties:
   - Internet address to connect to: **10.10.0.1**
   - Destination name: **Adatum VPN**
   - Allow other people to use this connection: Enabled
2. After you create the VPN, review the properties of the connection.
3. From the **Security** tab, reconfigure the VPN by using the following settings:
   - Type of VPN: **Point to Point Protocol (PPTP)**
   - Authentication: **Allow these protocols =Microsoft CHAP Version 2 (MS-CHAP v2)**

4. Test the VPN connection.
5. Wait for the VPN connection to connect.

Note: Your connection will be unsuccessful, and you will receive an error relating to authentication issues.

Create a VPN policy based on the Windows Groups condition

1. On LON-RTR, switch to the Network Policy Server console.
2. Disable the two existing network policies. These would interfere with the processing of the policy you are about to create.
3. Create a new network policy by using the following properties:
   - Policy name: **Adatum VPN Policy**
   - Type of network access server: **Remote Access Server(VPN-Dial up)**
   - Condition: **Windows Groups = Domain Admins**
   - Permission: **Access granted**
   - Authentication methods: Default settings
   - Constraints: Default settings
   - Settings: Default settings

Test the VPN

1. Switch to LON-CL2.
2. Test the Adatum VPN connection. This connection is now successful.

When you finish the demonstration, revert all virtual machines to their initial state. To do this, perform the following steps:

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machines** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for the following machines: 20413C-LON-RTR and 20413C-LON-CL2.

Planning Client Connectivity to VPNs

When planning for client connectivity to VPNs, consider the following factors:

- Client operating system. Depending on the client operating system, you can use different VPN technologies. For instance, VPN Reconnect with IKEv2 is only available to client computers running Windows 7 or newer. DirectAccess also is only available to client computers running Windows 7 or newer, but requires an Enterprise edition of the Windows client operating system.

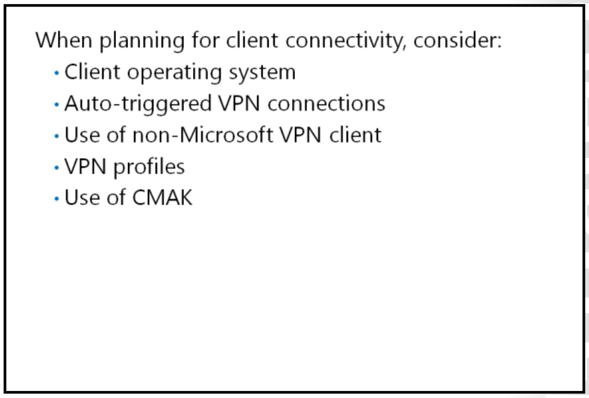

When planning for client connectivity, consider:

- Client operating system
- Auto-triggered VPN connections
- Use of non-Microsoft VPN client
- VPN profiles
- Use of CMAK

- Auto-triggered VPN connections. One of the main concerns of users when connecting remotely is the need to start a VPN connection. Organizations have been asking for a way to start a VPN connection automatically from a client computer without user intervention. Although you can achieve this by adding a VPN connection to the list of processes to start when the user signs in, the connection still requires the VPN client to start, and sometimes prompt the user for credentials. DirectAccess solves this problem by using a client that checks automatically for connectivity to the corporate network, and starts a remote connection automatically if it is unable to find the corporate network locally.
- Other VPN clients. If you need to provide clients that are using other, non-Microsoft operating systems with access to VPN, you must ensure that you configure the remote access server to accept the tunneling protocols and authentication types supported by these non-Microsoft clients.
- VPN profiles. Windows 8.1 includes the following inbox VPN clients:
  - Check Point VPN
  - F5 VPN
  - Juniper Networks Juno Pulse
  - SonicWALL Mobile Connect
- Client settings. You need to identify how to set up the VPN connection for your client computers. You can use the Connection Manager Administration Kit (CMAK) in Windows operating system to create connection profiles and then distribute these profiles by using software distribution. This process makes the connection profiles available on client images for new computers, or though network shares.

Additional Reading: For more information on CMAK, visit http://go.microsoft.com/fwlink/?LinkID=391894.

Lesson 3
Planning and Implementing Web Application Proxy

Many organizations need to provide access to web applications that are on the corporate network, to users who are not on the corporate network but who connect to the corporate network through the Internet. The process of configuring an application so that it is accessible from the Internet is called *publishing*. Windows Server 2012 R2 introduces the Web Application Proxy role service, which you can use for publishing applications. Web Application Proxy is deployed as a component of the Remote Access role in Windows Server 2012 R2.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the use of Web Application Proxy.
- Plan for authentication options.
- Plan for single sign-on (SSO).
- Plan for application publishing.
- Publish a secure web site.
- Plan for Workplace Join with Web Application Proxy.

Web Application Proxy Overview

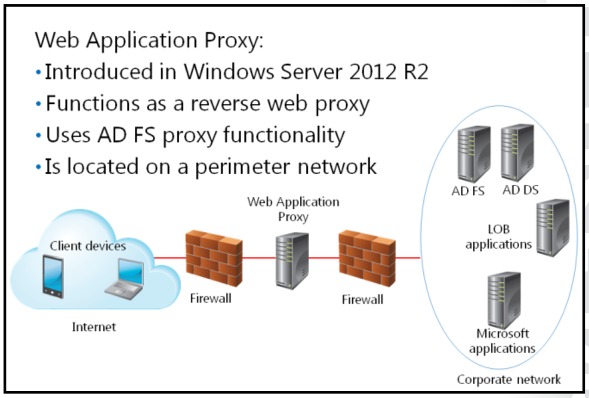

Web Application Proxy:

- Introduced in Windows Server 2012 R2
- Functions as a reverse web proxy
- Uses AD FS proxy functionality
- Is located on a perimeter network

For users who connect remotely to the corporate network, Web Application Proxy (introduced in Windows Server 2012 R2) functions as a reverse web proxy to provide access to internal corporate web applications. Web Application Proxy uses the Active Directory Federation Services (AD FS) technology to preauthenticate Internet users, and acts as an AD FS proxy for publishing claims-aware applications.

AD FS is required during the Web Application Proxy configuration process. After Web Application Proxy configuration is complete, you can publish both claims-aware applications that use AD FS preauthentication and web applications that use pass-through preauthentication.

You typically place the Web Application Proxy servers on the perimeter network between two firewall devices. The AD FS server and applications that are published are located on the corporate network, together with domain controllers and other internal servers. The AD FS server and applications are protected by the second firewall. This scenario provides secure access to corporate applications for users whose computers are connected to the Internet, and, at the same time, protects the corporate IT infrastructure from security threats from the Internet.

Planning Authentication Options

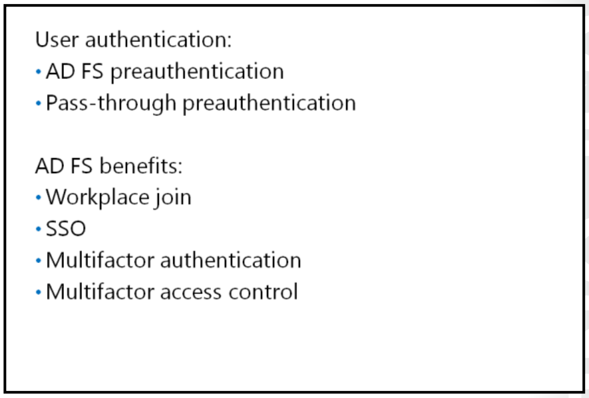

User authentication:

- AD FS preauthentication
- Pass-through preauthentication

AD FS benefits:

- Workplace join
- SSO
- Multifactor authentication
- Multifactor access control

Web Application Proxy in Windows Server 2012 R2 supports two types of preauthentication:

- AD FS preauthentication. AD FS preauthentication uses AD FS for web applications that use claims-based authentication. When a user initiates a connection to the corporate web application, the first entry point to which the user will connect is Web Application Proxy. Next, Web Application Proxy will preauthenticate the user in the AD FS server, and if the authentication succeeds, Web Application Proxy will establish a connection to the corporate network web server that hosts the application.
- Pass-through preauthentication. Pass-through preauthentication does not use AD FS for authentication, and Web Application Proxy does not preauthenticate the user. Instead, the user connects to the web application through Web Application Proxy. If the web application is configured for authentication, then Web Application Proxy authenticates the user. The main advantage of pass-through authentication is that you can publish applications that are not claims-aware to external users.

AD FS preauthentication provides the following benefits when compared to pass-through preauthentication:

- Workplace Join. Workplace Join is a new feature in AD FS in Windows Server 2012 R2, which allows you to add devices to the workplace that are not members of the Active Directory domain—such as smart phones, tablets, or non-company laptops. After you add these non-domain devices to the workplace, you can configure them for AD FS preauthentication.
- SSO. SSO allows users that are preauthenticated by AD FS to enter their credentials only once to sign onto a corporate network. If users subsequently access other applications that use AD FS for authentication, they will not be prompted again for their credentials.
- Multifactor authentication. Multifactor authentication allows you to configure multiple types of credentials to strengthen security. For example, you can configure authentication so that users must enter user names and passwords together with an activated smart card.
- Multifactor access control. Multifactor access control is used by organizations that want to implementing authorization claim rules to strengthen their security in publishing web applications. You configure authorization claim rules so that they issue a permit or deny claim, which will in turn determine whether a user or a group will be allowed or denied access to the web application that is using AD FS preauthentication.

Planning for SSO

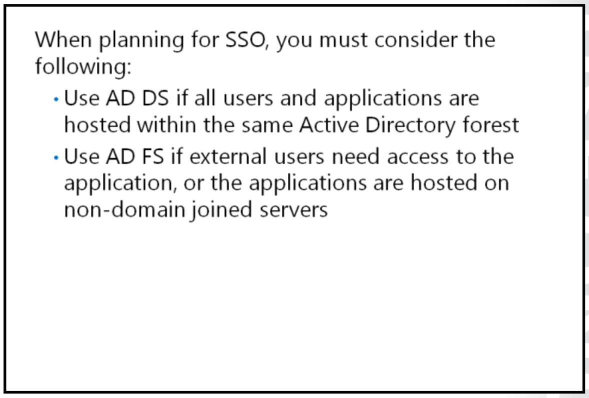

When planning for SSO, you must consider the following:

- Use AD DS if all users and applications are hosted within the same Active Directory forest
- Use AD FS if external users need access to the application, or the applications are hosted on non-domain joined servers

For many organizations, configuring access to applications and services by using SSO can be as simple as using AD DS. If all users are members of the same Active Directory forest, and if all applications run on servers that are members of the same Active Directory forest, you typically can use Active Directory authentication to provide access to the application.

However, there are situations where you must use AD FS to enable SSO to optimize the user experience. These situations include the following circumstances:

- The applications might be running on:
  - Servers that do not have Windows Server installed.
  - Servers that do not support Active Directory authentication.
  - Servers that have Windows Server installed, but that are not domain-joined.
  - Servers that are running a non-Microsoft web server service.
- The applications might require Security Assertion Markup Language or Web services for authentication and authorization.
- Large organizations frequently have multiple domains and forests that might result from mergers and acquisitions, or from 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. You can achieve this by using Workplace-Joined devices.

Organizations can use AD FS to enable SSO in these scenarios. If the organization has a single Active Directory 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: 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.

Note: SSO requires that you install AD FS on a computer that is running Windows Server 2012 R2.

Note: For detailed information on how to use AD FS to implement and manage single sign-on, visit http://go.microsoft.com/fwlink/?LinkID=391895

Planning for Application Publishing

The most important factor to consider when planning for application publishing by using Web Application Proxy is the type of application that you want to publish. You can publish applications that require claims-based authentication, integrated Windows authentication, Microsoft Office Forms Based Authentication, or Windows Store apps.

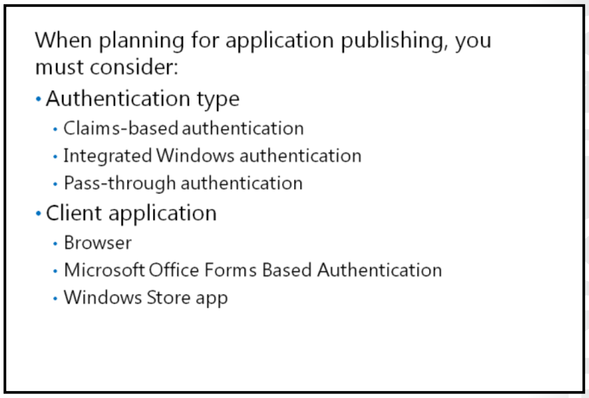

When planning for application publishing, you must consider:

- Authentication type
  - Claims-based authentication
  - Integrated Windows authentication
  - Pass-through authentication
- Client application
  - Browser
  - Microsoft Office Forms Based Authentication
  - Windows Store app

Claims-Based Authentication

For claims-based authentication, you must ensure that the following prerequisites are met:

- The AD FS server must have a claims-aware relying party trust for the application.
- The Web Application Proxy server does not need to be part of an Active Directory domain.
- Applications must be configured to use AD FS for SSO.

Additional Reading: For more information on how to plan and implement application publishing by using claims-based authentication, visit http://go.microsoft.com/fwlink/?LinkID=391896.

Integrated Windows Authentication

For applications that require integrated Windows authentication, the Web Application Proxy server must perform preauthentication of the user. In this scenario, you must ensure that the following prerequisites are met:

- The AD FS server must have a nonclaims-aware relying party trust for the application.
- The Web Application Proxy server must be part of an AD DS domain.
- The Web Application Proxy server must be able to provide delegation for the users who need access to the application.
- The application must be running on a computer that is running Windows Server 2012 R2 or Windows Server 2012.

Additional Reading: For more information on how to plan and implement application publishing by using integrated Windows authentication, visit http://go.microsoft.com/fwlink/?LinkID=391897.

Pass-through Authentication

Applications that use pass-through authentication do not require any extra planning. However, they cannot utilize any AD FS features, such as Workplace Join, multifactor authentication, and multifactor access control.

Microsoft Office Forms Based Authentication Clients

To publish an application for clients that use Microsoft Office Forms Based Authentication, you must add a relying party trust for the application to the AD FS server. The type of relying party trust must match the authentication requirements of the application, either claims-based or integrated Windows authentication.

Windows Store App Clients

To publish a web application to be accessed from a Windows Store app, you must ensure that the following prerequisites are met:

- The Windows Store app must support Open Authorization (OAuth) 2.0.
- The OAuth endpoint in AD FS must be proxy-enabled.
- You must use the Windows PowerShell® cmdlet **Set-WebApplicationProxyConfiguration** to configure the Web Application Proxy server with the URL for the AD FS server.

Additional Reading: For more information on planning application publishing, visit http://go.microsoft.com/fwlink/?LinkID=391898.

Additional Reading: For information on publishing SharePoint Server and Exchange Server through Web Application Proxy, visit http://go.microsoft.com/fwlink/?LinkID=391899.

Demonstration: Publishing a Secure Web Site

In this demonstration, you will learn how to:

- Install the AD FS role.
- Install the Web Application Proxy role service.
- Obtain a certificate.
- Obtain a certificate for a website.
- Configure Web Application Proxy.
- Publish the internal website.
- Configure internal website authentication.
- Disable DirectAccess on a client computer.
- Verify access to the internal website from the client computer.

Demonstration Steps

Install the AD FS role

1. Create a Microsoft Group Key Distribution Service (KDS) root key.
2. On LON-DC1, in Server Manager, add the **Active Directory Federation Services** server role.
3. On LON-DC1, run the AD FS Federation Server Configuration Wizard by using the following parameters:
   - Create a new Federation Service.
   - Create a stand-alone deployment.
   - Use the **LON-DC1.Adatum.com** certificate.
   - Choose a service name of **LON-DC1.Adatum.com**.
4. On LON-DC1, open Windows PowerShell, and then use the following command to enable modification of the assigned certificates:

```
set-ADFSProperties –AutoCertificateRollover $False
```

5. In the AD FS Management console, add the **LON-DC1.Adatum.com** certificate as a new token-signing certificate.

6. Verify that the certificate has a subject of **CN=LON-DC1.Adatum.com** and its purposes include **Proves your identity to a remote computer** and **Ensures the identity of a remote computer**.
7. Make the new certificate the primary certificate and remove the old certificate

Install the Web Application Proxy role service

1. Switch to LON-RTR.
2. Open Server Manager, and then on the **Manage** menu, click **Add roles and features**.
3. In the Add Roles and Features Wizard, on the **Select server roles** page, expand **Remote Access**, and then click **Web Application Proxy**.

Obtain a certificate

- Open a Microsoft Management Console (MMC), add the Certificates - Computer account snap-in, and then request a new certificate with the following settings:
  - Subject Name: **Common Name**: **lon-dc1.adatum.com**
  - Alternative name: **DNS**: **lon-dc1.adatum.com**, **enterpriseregistration.adatum.com**, **lon-svr1.adatum.com**

Obtain a certificate for a website

1. Switch to LON-SVR1.
2. Open an MMC, add the **Certificates - Computer account** snap-in, and then request a new certificate with the Subject Name **Common Name: lon-svr1.adatum.com**.
3. Open the Internet Information Services (IIS) Manager console,
4. navigate to **LON-SVR1/Sites**, and then click **Default Web site**.
5. Configure site bindings by entering **lon-svr1.adatum.com** as a host name, and by selecting **lon-svr1.adatum.com** as the **SSL Certificate**.
6. Close the Internet Information Services (IIS) Manager console.

Configure Web Application Proxy

1. From Server Manager, open the Remote Access Management console.
2. In the navigation pane, click **Web Application Proxy**, and run the Web Application Proxy Configuration Wizard.
3. In the Web Application Proxy Configuration Wizard, for **Federation service name**, enter **lon-dc1.adatum.com** as the FQDN of the federation service.
4. In the **User name** and **Password** text boxes, type **Administrator** and **Pa$$w0rd**.
5. On the **AD FS Proxy Certificate** page, in the list of certificates currently installed on the Web Application Proxy server, select **lon-dc1.adatum.com** as the certificate to be used by Web Application Proxy for AD FS proxy functionality.
6. On the **Results** page, verify that the configuration was successful, and then close the wizard.

Publish the internal website

1. On the Web Application Proxy server, in the Remote Access Management console, in the navigation pane, start the Publish New Application Wizard.
2. On the **Preauthentication** page, click **Pass-through**.
3. In the **Name** text box, type **LON-SVR1 Web** as a friendly name for the application.

4. In the **External URL** text box, type **https://lon-svr1.adatum.com** as the external URL for this application.

5. In the **External certificate** list, click the **lon-dc1.adatum.com** certificate.

6. In the **Backend server URL** text box, ensure that **https://lon-svr1.adatum.com** displays.

Note: This value is entered automatically when you enter the external URL.

7. On the **Confirmation** page, review the settings, and then click **Publish**.

8. On the **Results** page, ensure that the application publishes successfully, and then click **Close**.

Configure internal website authentication

1. Switch to LON-SVR1.

2. From Server Manager, open the Internet Information Services (IIS) Manager console.

3. In the Internet Information Services (IIS) Manager console, navigate to **Default Web Site**.

4. Configure Authentication for the Default Web Site with following settings:

   o Windows Authentication: Enabled

   o Anonymous Authentication: Disabled

5. Close the Internet Information Services (IIS) Manager console.

Disable DirectAccess on a client computer

1. Switch to LON-CL1.

2. Open Control Panel.

3. In Control Panel, remove LON-CL1 from the **Adatum.com** domain, add LON-CL1 to a workgroup named **WORKGROUP**, and then restart LON-CL1.

Verify access to the internal website from the client computer

1. Sign in with user name **Admin** and password **Pa$$w0rd**.

2. Open Control Panel.

3. In Control Panel, remove LON-CL1 from the **Adatum.com** domain, add LON-CL1 to a workgroup named **WORKGROUP**, and then restart LON-CL1.

4. Sign in with user name **Admin** and password **Pa$$w0rd**.

5. On LON-CL1, use Notepad to create a file named **Hosts**.

6. Add the following content to the hosts file: **131.107.0.2 lon-svr1.adatum.com**

7. Copy the hosts file to **C:\Windows\System32\drivers\etc**.

8. On LON-CL1, open Windows Internet Explorer®, and then type the following address: **https://lon-svr1.adatum.com**.

9. When prompted, in the **Internet Explorer** dialog box, type **Adatum\Bill** for the user name and **Pa$$w0rd** for the password, and then click **OK**.

10. Verify that the default IIS 8.0 web page for LON-SVR1 opens.

11. If you are unable to connect to https://lon-svr1.adatum.com, perform the following steps:

    a. On LON-CL1, on the **Start** screen, type **cmd**, and then press Enter.

    b. In the command prompt window, type **regedit**, and then press Enter.

c. In the **User Account Control** dialog box, click **Yes**.

d. In the Registry Editor window, in the navigation pane, navigate to **HKLM\Software\Policies\Microsoft\Windows NT\DNSClient\DNSPolicyConfig**.

e. Right-click each of the entries starting with **DA**, click **Delete**, and in the **Confirm Key Delete** dialog box, click **Yes**.

f. Close the Registry Editor window.

12. Restart LON-CL1 and perform steps 8 through 10 to verify connectivity to default IIS 8.0 web page on LON-SVR1.

When you finish the demonstration, revert all virtual machines to their initial state. To do this, perform the following steps:

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machines** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for the following machines: 20413C-LON-RTR. 20413C-LON-SVR1, and 20413C-LON-CL1.

Planning Workplace Join with Web Application Proxy

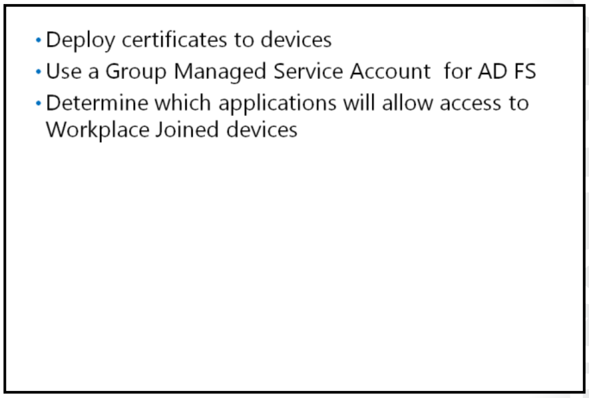

- Deploy certificates to devices
- Use a Group Managed Service Account for AD FS
- Determine which applications will allow access to Workplace Joined devices

You can use Workplace Join to connect non-domain joined devices to your organizations' network, and provide SSO for these devices. This allows users to connect to company resources by using their own devices. You can use Workplace Join on Windows 8.1, Windows Server 2012 R2, and Apple iOS mobile operating system devices.

When planning to use Workplace Join with Web Application Proxy, consider the following best practices:

- Deploy certificates to devices. Devices must trust the CA that issues the certificates used by the AD FS server that is providing Workplace Join. The certificate that AD FS uses must contain a subject name and a subject alternate name. Both the subject name and the subject alternate name must be the same name that the clients use to connect to the Web Application Proxy server.
- Use a Group Managed Service Account for the AD FS services. You can use Group Managed Service Accounts across multiple servers and have their password changed automatically by the servers. This makes them more secure than regular user accounts. We recommend you to use a Group Managed Service Account instead of a user account to facilitate password management.
- Determine which applications will allow access to Workplace-Joined devices. After a device is Workplace-Joined, you will be able to retrieve information about the device from Active Directory, and then use this information to generate device claims. Because the end user usually owns Workplace-Joined devices, you might not want to have sensitive information from your web applications cached on these devices. For sensitive applications, you can use claims to determine whether users are accessing the application from a Workplace-Joined device or from a domain-joined device, and then determine how to proceed with authorization.

Lesson 4
Planning a Complex Remote Access Infrastructure

Larger organizations that provide remote access to hundreds, or thousands of users from multiple locations must consider how to provide a fault-tolerant and highly available remote access solution. Administrators for these organizations should take into account the different available forms of authentication, access to multiple Active Directory forests, remote connection monitoring, secure access to resources, and capacity planning.

Lesson Objectives

After completing this lesson, you will be able to:

- Plan for a highly available remote access infrastructure.
- Plan for remote access capacity.
- Plan for remote access with multiple locations.
- Plan for additional authentication options.
- Plan for remote access with multiple forests.
- Plan for a RADIUS implementation.

Planning a Highly Available Remote Access Infrastructure

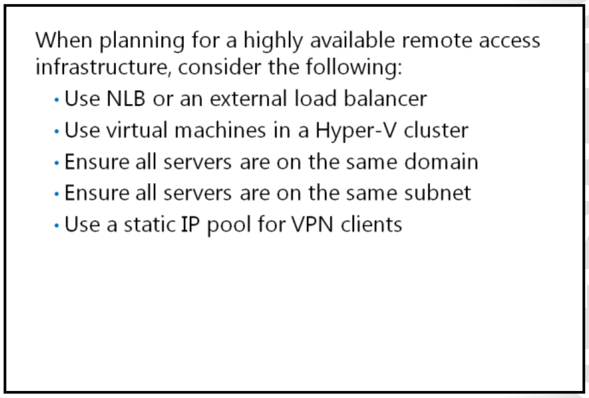

When planning for a highly available remote access infrastructure, consider the following:

- Use NLB or an external load balancer
- Use virtual machines in a Hyper-V cluster
- Ensure all servers are on the same domain
- Ensure all servers are on the same subnet
- Use a static IP pool for VPN clients

For larger implementations of remote access server, you may need to implement multiple servers to manage larger demands and provide high availability of services. You can cluster the remote access server role on an NLB cluster to provide high availability and load balancing.

When planning for a highly available remote access infrastructure, you must follow these guidelines:

- Decide whether to use NLB, or an external load balancer.
- If you use NLB, each server in the cluster must have the Remote Access role service and the NLB installed.
- Consider using virtual machines in a Hyper-V® cluster to provide high availability along with load balancing.
- Virtual machines must have media access control (MAC) address spoofing enabled for NLB.
- All servers must belong to the same Active Directory domain.
- All servers must be in the same subnet.
- All servers must have the same number of network adapters for DirectAccess.
- Use identical certificates for each server, and consider using a PKI.
- If your implementation uses only IPv6, you must use a IPv6 addressing prefix of /59.

- VPN client traffic must use a static IP pool for load balancing.
- Any changes made by using the RRAS management console must be replicated manually to all servers in the cluster.

Additional Reading: For more information on planning a load-balanced cluster deployment, visit http://go.microsoft.com/fwlink/?LinkID=391900.

Planning Remote Access Capacity

Microsoft has performed tests for remote access (DirectAccess) capacity planning on high-end and low-end servers. The tests were based on a mix of IP-HTTPS and Teredo clients One-third of the test subjects were IP-HTTPS clients, and the remaining two-thirds of the subjects were Teredo clients. For better performance, Receive Side Scaling was enabled, and hyper-threading was disabled. Low-end servers had 4 cores and 4 gigabytes (GB) of random access memory (RAM), whereas high-end servers had 8 cores and 8 GB of RAM. The following table lists the results.

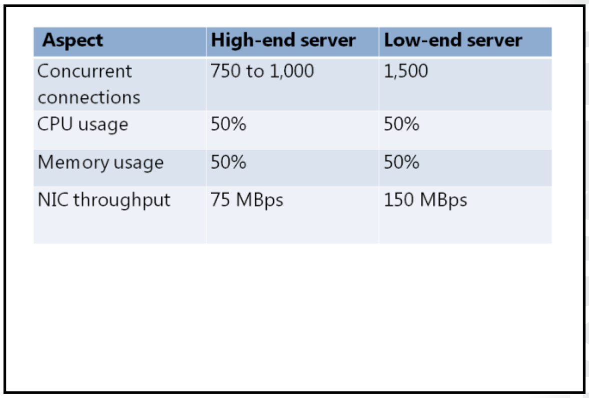

Aspect	High-end server	Low-end server
Concurrent connections	750 to 1,000	1,500
CPU usage	50%	50%
Memory usage	50%	50%
NIC throughput	75 MBps	150 MBps

Aspect	High-end server	Low-end server
Concurrent connections	750 to 1,000	1,500
CPU usage	50%	50%
Memory usage	50%	50%
NIC throughput	75 megabytes per second (MBps)	150 MBps

Additional Reading: For complete details on the DirectAccess capacity planning tests performed, visit http://go.microsoft.com/fwlink/?LinkID=391901.

When planning for remote access capacity, you can use the results in the table to determine the number of servers that you should use based on your user load.

Planning Remote Access With Multiple Locations

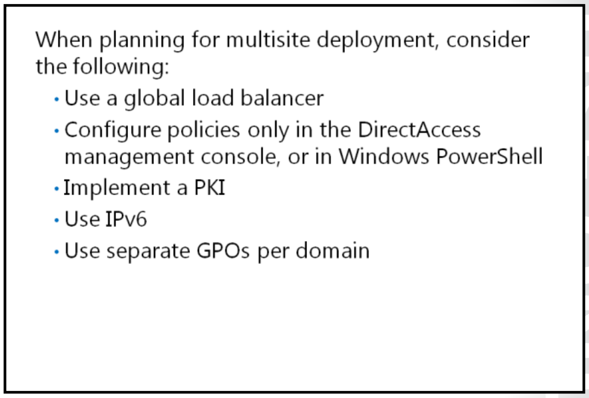

When planning for multisite deployment, consider the following:

- Use a global load balancer
- Configure policies only in the DirectAccess management console, or in Windows PowerShell
- Implement a PKI
- Use IPv6
- Use separate GPOs per domain

For larger organizations with offices spread across the globe, you might have to deploy remote access servers in more than one location. You can provide high availability on each location by using a clustered deployment. However, you still want users to be able to connect to the remote access server that is closest to their location.

When planning for a multisite deployment, you should take the following considerations into account:

- Windows 8 and Windows 8.1 clients can identify an entry point automatically.
- You can use a global load balancer to identify the best point of entry.
- Windows 7 clients cannot select their endpoints automatically. Therefore, you must enable Windows 7 support on each endpoint.
- You must change policies only by using the DirectAccess management console and Windows PowerShell.
- You must implement a PKI.
- The corporate network must use IPv6. ISATAP is not supported.
- Create a separate GPO per domain for easier management.

Additional Reading: For more information on planning to deploy multiple Remote Access servers in multiple locations, visit http://go.microsoft.com/fwlink/?LinkID=391902.

Planning Additional Authentication Options

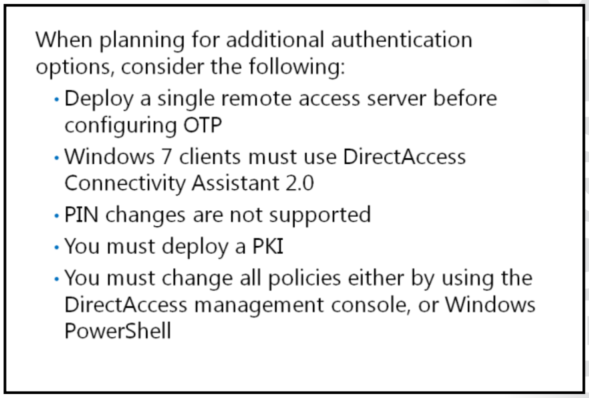

When planning for additional authentication options, consider the following:

- Deploy a single remote access server before configuring OTP
- Windows 7 clients must use DirectAccess Connectivity Assistant 2.0
- PIN changes are not supported
- You must deploy a PKI
- You must change all policies either by using the DirectAccess management console, or Windows PowerShell

DirectAccess clients can authenticate by using standard Active Directory credentials, or by using a two-factor authentication that supports OTPs. When planning for two-factor authentication, consider the following:

- You must deploy a single remote access server before configuring OTP.
- Windows 7 clients must use the Microsoft DirectAccess Connectivity Assistant 2.0.
- PIN changes are not supported.
- You must deploy a PKI.
- You must change all policies either in the DirectAccess management console, or by using Windows PowerShell.

Planning Remote Access with Multiple Forests

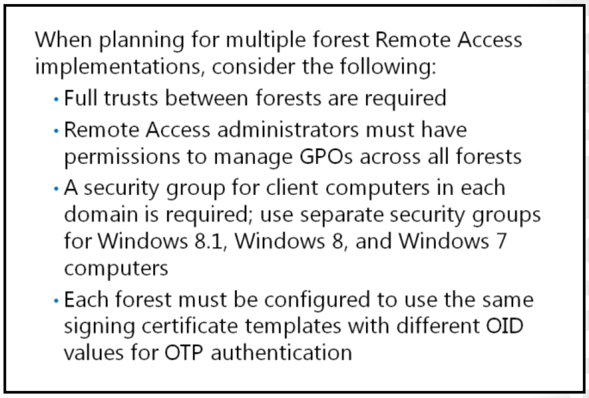

When planning for multiple forest Remote Access implementations, consider the following:

- Full trusts between forests are required
- Remote Access administrators must have permissions to manage GPOs across all forests
- A security group for client computers in each domain is required; use separate security groups for Windows 8.1, Windows 8, and Windows 7 computers
- Each forest must be configured to use the same signing certificate templates with different OID values for OTP authentication

In more complex environments, remote access servers may require connectivity to resources in multiple forests. When planning for multiple forest access, consider the following:

- A full trust is required between forests.
- Remote Access administrators must have permissions to manage GPOs in all domains for which access is required.
- A security group for remote client computers is required for each forest. We recommend you to have a security group for each domain and for each client type (Windows 7, Windows 8, or Windows 8.1).
- If you are using OTP, you must configure each forest to use the same signing certificate templates with different object identifier (OID) values.

Additional Reading: For more information on configuring OTP in a multiple forest environment, visit http://go.microsoft.com/fwlink/?LinkID=391903.

- If you use IPsec, all clients and servers must have certificates issued by the same authority.

Planning a RADIUS Implementation

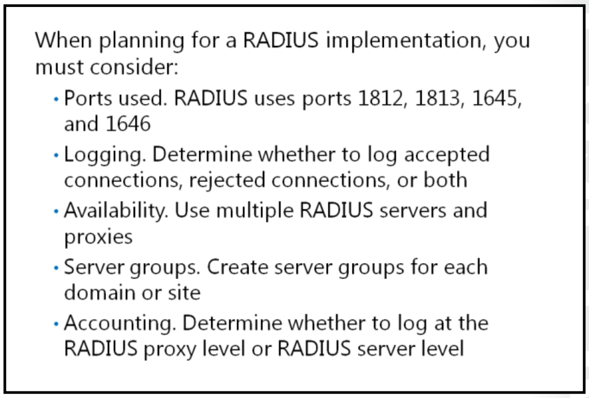

When planning for a RADIUS implementation, you must consider:

- Ports used. RADIUS uses ports 1812, 1813, 1645, and 1646
- Logging. Determine whether to log accepted connections, rejected connections, or both
- Availability. Use multiple RADIUS servers and proxies
- Server groups. Create server groups for each domain or site
- Accounting. Determine whether to log at the RADIUS proxy level or RADIUS server level

Remote access solutions can use RADIUS servers and proxies. When using RADIUS in a remote access solution, you must consider the following:

- Plan for RADIUS ports. RADIUS uses UDP traffic on ports 1812 and 1645 for authentication traffic, and ports 1813 and 1646 for accounting messages.
- Plan for logging. If there is too much logging, it becomes difficult to find the information you need. Conversely, if there is too little logging, you may not have the information you require. Therefore, determine ahead of time what to log. For example, determine whether you want to log accepted connections, rejected connections, or both.
- Plan for availability. You can use two RADIUS proxy servers as a primary RADIUS server, and a secondary RADIUS server for most RADIUS clients. For RADIUS clients that do not allow the use of a secondary RADIUS server, you must use NLB.
- Plan for RADIUS server groups. A RADIUS proxy can send connections to a RADIUS server group. If you have multiple domains, you usually need to create a RADIUS server group for each domain. However, if the domain has multiple sites that contain RADIUS servers, you need to create a RADIUS server group for each site.

- Plan for accounting. Accounting messages can be logged at the RADIUS proxy or RADIUS server level. Logging accounting at the RADIUS proxy level avoids traffic being sent over the intranet and provides centralized logging. However, if your solution provides access to RADIUS servers in different domains, and if there is a need for logging at the domain level, consider logging accounting data at the RADIUS server level. You should also determine whether to use SQL Server or flat files for logging data. For larger environments and added reporting flexibility, we recommend using SQL Server.

Additional Reading: To know more about planning NPS as a RADIUS server, visit http://go.microsoft.com/fwlink/?LinkID=391904.

Lab: Designing and Implementing Network Access Services

Scenario

A. Datum Corporation is evaluating the network access requirements for users who occasionally work from hotel rooms or from home. Many of the management staff in the London head office travel to remote locations and must access organizational data from hotel rooms. Executives also want to work from home, or while away on vacation. In addition, many workers in Trey Research and Contoso, Ltd work from home. They occasionally bring their computers into the corporate office, and often fail at their attempts to connect to all corporate network resources because their systems are in breach of the corporate network health policies.

Objectives

After completing the lab, you will be able to:

- Design a remote access strategy.
- Plan and implement a DirectAccess solution.
- Plan and implement a VPN solution.
- Plan and implement a Web Application Proxy solution.

Lab Setup

Estimated Time: 120 minutes

Virtual machines	20413C-LON-DC1 20413C-LON-SVR1 20413C-LON-RTR 20413C-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, start Hyper-V Manager.
2. In Hyper-V Manager, click **20413C-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Sign in using the following credentials:
   - User name: **Administrator**
   - Password: **Pa$$w0rd**
   - Domain: **Adatum**
5. Repeat steps 2 through 4 for **20413C-LON-RTR**.
6. Important: Do not start 20413C-LON-SVR1 and 20413C-LON-CL1 until instructed to do so.

Exercise 1: Designing a Remote Access Strategy

Scenario

The current VPN deployment at A. Datum Corporation consists of a single VPN server. Clients use L2TP connections, and have connectivity to the entire network when connected. Although a VPN server is currently in place in the London location, you must design a remote access solution based on user and business requirements.

Remote Access Services Strategy	
Document Reference Number: BS0905/1	
Document Author	Brad Sutton
Date	5th September

Requirements Overview

Design a remote access strategy to support the following objectives:

- Data security is very important.
- A. Datum has an infrastructure in place for deploying certificates and smart cards.
- Some executives have had problems with VPN connections being blocked by hotel firewalls.
- Users from non-European sites have complained about slow access to data over the VPN.
- There is only a single Internet connection for Contoso, Ltd, which is located in Paris.
- The current service provider for Internet access provides no guarantees for availability. Availability guarantees are required for disaster recovery planning.
- It is important that lost connections are reestablished automatically.
- Where possible, computers must be brought into the management scope of the A. Datum or Trey Research networks.

Additional Information

- Datum has Windows 8.1, Windows 8, and Windows 7 clients.
- Some users who use computers that are not members of the adatum.com domain require access to the internal network. Security for these devices is critical.
- Some Contoso users only require access to a secure website on the internal network at A. Datum.
- Automatic client connection is very important.

Proposals

1. What technology should you use to provide automatic client connectivity?

2. How can you provide secure access for computers that are not joined to any domain?

3. How can you provide Contoso users with access to the secure website on the A. Datum internal network?

The main tasks for this exercise are as follows:

1. Read the supporting documentation.
2. Update the proposal document with your planned course of action.
3. Examine the suggested proposals in the Lab Answer Key.
4. Discuss your proposed solution with the class, as guided by your instructor.

▶ Task 1: Read the supporting documentation

- Read the documentation provided.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the Remote Access Services Strategy document.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones in the Lab Answer Key.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you should have successfully designed a remote access strategy.

Exercise 2: Planning and Implementing a DirectAccess Solution

Scenario

You now must determine the enterprise components that must be in place to support DirectAccess deployment. You also must determine the server roles that you must deploy prior to deploying DirectAccess.

Supplemental Documentation

DirectAccess Strategy	
Document Reference Number: BS0907/1	
Document Author Date	Brad Sutton 7th September
Requirements Overview Create a DirectAccess implementation plan to support the following objectives: • Determine what enterprise components are required for DirectAccess. • Determine what server roles are required for DirectAccess.	
Proposals 1. What components must be in place to support DirectAccess? 2. Will you implement a PKI?	

DirectAccess Strategy
3. What must you configure on the DNS servers to support your planned deployment?

The main tasks for this exercise are as follows:

1. Read the supporting documentation.
2. Update the proposal document with your planned course of action.
3. Examine the suggested proposals in the Lab Answer Key.
4. Discuss your proposed solution with the class, as guided by your instructor.
5. Configure AD DS and DNS.
6. Configure required certificates.
7. Configure internal resources.
8. Configure a DirectAccess server.
9. Configure DirectAccess Group Policy Object (GPO) settings.
10. Verify certificate distribution.
11. Verify IP configuration.
12. Move LON-CL1, and verify connectivity to intranet resources.
13. Revert virtual machines.

▶ Task 1: Read the supporting documentation

- Read the documentation provided.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the DirectAccess Strategy document.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones in the Lab Answer Key.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

▶ Task 5: Configure AD DS and DNS

1. Create a security group for DirectAccess client computers by performing the following steps:
   a. Switch to LON-DC1.
   b. Open the **Active Directory Users and Computers** console, and create an OU named **DA_Clients OU**.
   c. Within that OU, create a **Global Security group** named **DA_Clients**.
   d. Modify the membership of the **DA_Clients** group to include **LON-CL1**.
   e. Close Active Directory Users and Computers.

2. Configure firewall rules for ICMPv6 traffic by performing the following steps:

   a. Open the Group Policy Management Console, and then open **Default Domain Policy**.

   b. In the Group Policy Management Editor, navigate to **Computer Configuration\Policies\Windows Settings\Security Settings\Windows Firewall with Advanced Security\Windows Firewall with Advanced Security**.

   c. Create a new inbound rule with the following settings:

      - Rule Type: **Custom**
      - Protocol type: **ICMPv6**
      - Specific ICMP types: **Echo Request**
      - Name: **Inbound ICMPv6 Echo Requests**

   d. Create a new outbound rule with the following settings:

      - Rule Type: **Custom**
      - Protocol type: **ICMPv6**
      - Specific ICMP types: **Echo Request**
      - Action: **Allow the connection**
      - Name: **Outbound ICMPv6 Echo Requests**

   e. Close both the Group Policy Management Editor and the Group Policy Management Console.

3. Create required DNS records by performing the following steps:

   a. Open the **DNS Manager** console, and then create a new host record with the following settings:

      - Name: **nls**
      - IP Address: **172.16.0.11**

   b. Close the DNS Manager console.

4. Remove ISATAP from the DNS global query block list by performing the following steps:

   a. Open a command prompt window, type the following command, and then press Enter:

      ```
      dnscmd /config /globalqueryblocklist wpad
      ```

      Ensure that the **Command completed successfully** message displays.

   b. Close the command prompt window.

5. Switch to LON-RTR and configure the DNS suffix by performing the following steps:

   a. In the **Ethernet Properties** dialog box, in the **Internet Protocol Version 4 (TCP/IPv4)** dialog box, add the **Adatum.com** DNS suffix.

   b. Close the **Ethernet Properties** dialog box.

6. Configure the Ethernet 2 properties as follows:

   - Change the **Ethernet 2\ Internet Protocol Version 4 (TCP/IPv4)** configuration using the following configuration settings:

     - IP address: **131.107.0.2**
     - Subnet mask: **255.255.0.0**

▶ Task 6: Configure required certificates

1. Configure the CRL distribution settings by performing the following steps:

   a. Switch to LON-DC1, and open the Certification Authority console.

   b. Configure the **AdatumCA** CA with the following extension settings:

      - Add Location: **http://crl.adatum.com/crld/**
      - Variable: **CAName**, **CRLNameSuffix**, **DeltaCRLAllowed**
      - Location: **.crl**
      - Select the following:
      - **Include in CRLs. Clients use this to find Delta CRL locations.**
      - **Include in the CDP extension of issued certificates.**
      - Do not restart Active Directory Certificate Services.
      - Add Location: **\\LON-RTR\crldist$\**
      - Variable: **CAName**, **CRLNameSuffix**, **DeltaCRLAllowed**
      - Location: **.crl**
      - Select the following:
        - **Include in CRLs. Clients use this to find Delta CRL locations.**
        - **Include in the CDP extension of issued certificates.**
      - Restart Active Directory Certificate Services.

2. Duplicate the web certificate template and configure appropriate permission by performing the following steps:

   a. In the Certificate Templates console, in the contents pane, duplicate the **Web Server** template by using the following options:

      - Template display name: **Adatum Web Server Certificate**
      - Request Handling: **Allow private key to be exported**
      - Authenticated Users permissions: under **Allow**, click **Enroll**

   b. Close the Certificate Templates console.

   c. In the Certification Authority console, choose to issue a New Certificate Template, and select the **Adatum Web Server Certificate** template.

   d. Restart the AdatumCA certification authority.

   e. Close the Certification Authority console.

3. Configure computer certificate autoenrollment by performing the following steps:

   a. On LON-DC1, open the Group Policy Management Console.

   b. In the Group Policy Management Console, navigate to **Forest: Adatum.com\Domains\Adatum.com**.

   c. Edit the **Default Domain Policy**.

   d. In the Group Policy Management Editor, navigate to **Computer Configuration\Policies\Windows Settings\Security Settings\Public Key Policies**.

   e. Under **Automatic Certificate Request Settings**, configure **Automatic Certificate Request** to issue the **Computer** certificate.

   f. Close both the Group Policy Management Editor and the Group Policy Management Console.

▶ Task 7: Configure internal resources

1. Start LON-SVR1 and then request a certificate for LON-SVR1 by performing the following steps:

   a. At the command prompt, type the following command, and then press Enter:

   ```
   Mmc
   ```

   b. In the MMC, add the **Certificates** snap-in for **Local computer**.

   c. In the Certificates snap-in console tree, navigate to **Certificates (Local Computer)\Personal\Certificates**, and request a new certificate.

   d. Under **Request Certificates**, click **Adatum Web Server Certificate** and specify the following setting:

   e. Subject name: Under **Common name**, type **nls.adatum.com**

   f. In the Certificates snap-in, in the details pane, verify that a new certificate with the name **nls.adatum.com** was enrolled with the **Intended Purposes** of **Server Authentication**.

   g. Close the console window. When you are prompted to save settings, click **No**.

2. To change the HTTPS bindings, perform the following steps:

   a. Open Internet Information Services (IIS) Manager.

   b. In the Internet Information Services (IIS) Manager console, navigate to and click **Default Web site**.

   c. Configure Site Bindings by selecting **nls.adatum.com** for **SSL Certificate**.

   d. Close the Internet Information Services (IIS) Manager console.

3. Create a shared folder on LON-SVR1 by performing the following steps:

   a. Create and share a folder named **C:\Files** with default values.

   b. Create a new text file in this folder named **DirectAccess.txt**.

   c. Open the text file in Notepad, and type **This is a corporate file**.

   d. Save the file.

▶ Task 8: Configure a DirectAccess server

1. Obtain required certificates for LON-RTR by performing the following steps:

   a. Switch to LON-RTR.

   b. Open a command prompt, and refresh group policy by typing the following command:

   ```
   gpupdate /force
   ```

   c. Open the MMC by typing **mmc** at a command prompt.

   d. Add the **Certificates** snap-in for **Local computer**.

e. In the Certificates snap-in, in the MMC, request a new certificate with the following settings:

   - Certificate template: **Adatum Web Server Certificate**
   - Common name: **131.107.0.2**
   - Friendly name: **IP-HTTPS Certificate**

f. Close the MMC.

2. Create a CRL distribution point on LON-RTR by performing the following steps:

   a. Switch to Server Manager.

   b. In Internet Information Services (IIS) Manager, create new virtual directory named **CRLD**, and assign **c:\crldist** as a home directory.

   c. Enable directory browsing, and then allow the double escaping feature.

   Question: Why did you make the CRL available on the edge server?

3. Share and secure permissions to the CRL distribution point by performing the following step:

   o In the File Explorer details pane, right-click the **CRLDist** folder, click **Properties**, and then grant **Full Control Share** and **NTFS** permissions.

4. Publish the CRL to LON-RTR by performing the following steps:

Note: This step makes the CRL available on the edge server for Internet-based DirectAccess clients.

   a. Switch to LON-DC1.

   b. Start the Certification Authority console.

   c. In the console tree, open **AdatumCA**, right-click **Revoked Certificates**, point to **All Tasks**, and then click **Publish**.

5. Complete the DirectAccess Getting Started Wizard on LON-RTR by performing the following steps:

   a. Restart LON-RTR and then sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.

   b. On LON-RTR, open Server Manager.

   c. In Server Manager, click **Tools**, and then click **Routing and Remote Access**.

   d. In Routing and Remote Access, disable the existing configuration and close the console.

   e. In Server Manager, open the Remote Management console, click **Configuration**, and then start the **Enable Direct Access Wizard**.

   f. Complete the wizard with following settings:

      - Network Topology: **Edge** is selected
      - **131.107.0.2** is used by clients to connect to the remote access server

   g. In the Remote Access Management console, under Step 1, click **Edit**.

   h. Add the **DA_Clients** group.

   i. In the Remote Access Management console details pane, under **Step 2**, click **Edit**.

   j. On the **Network Topology** page, verify that **Edge** is selected, and then type **131.107.0.2**.

k. On the **Network Adapters** page, verify that **CN=131.107.0.2** is used as a certificate to authenticate IP-HTTPS connection.

l. On the **Authentication** page, click **Use computer certificates**, click **Browse**, and then click **AdatumCA**.

m. On the **VPN Configuration** page, click **Finish**.

n. In details pane of the Remote Access Management console, under **Step 3**, click **Edit**.

o. On the **Network Location Server** page, click **The network location server is deployed on a remote web server (recommended)**, in the URL for the network location server, type **https://nls.adatum.com**, and then click **Validate**.

   Ensure that the URL validates.

p. On the **DNS** page, examine the values, and then click **Next**.

q. In the **DNS Suffix Search List**, click **Next**.

r. On the **Management** page, click **Finish**.

s. In the Remote Access Management console, in the details pane, review the setting for **Step 4**.

t. In the **Remote Access Review** dialog box, click **Apply**.

u. Under **Applying Remote Access Setup Wizard Settings**, click **Close**.

6. Update Group Policy settings on LON-RTR by performing the following step:

   o Open the command prompt and type the following commands, pressing Enter at the end of each line:

```
gpupdate /force
Ipconfig
```

7. Verify that LON-RTR has an IPv6 address for **Tunnel adapter IPHTTPSInterface** that begins with **2002**.

▶ Task 9: Configure DirectAccess Group Policy Object (GPO) settings

1. Switch to LON-DC1 and open the Group Policy Management Console.
2. Locate the **DirectAccess Client Settings** GPO.
3. In the results pane, remove the **DirectAccess – Laptop only WMI filter**.
4. Close the Group Policy Management Editor and Group Policy Management Console.
5. Start LON-CL1, and sign in as **Adatum\Administrator** with the password of **Pa$$w0rd**.
6. Open a command prompt window, and then type the following commands, pressing Enter at the end of each line:

```
gpupdate /force
gpresult /R
```

7. Verify that **DirectAccess Client Settings GPO** displays in the list of the Applied Policy objects for the Computer Settings.

▶ Task 10: Verify certificate distribution

1. On LON-CL1, open the Certificates snap-in console.
2. Verify that a certificate with the name **LON-CL1.adatum.com** displays with the **Intended Purposes** of **Client Authentication** and **Server Authentication**.
3. Close the console without saving.

▶ Task 11: Verify IP configuration

1. On LON-CL1, from the desktop, open an Internet Explorer window, and in the Address bar, type **http://lon-svr1.adatum.com/**.
2. Verify that the default IIS 8 webpage for LON-SVR1 displays.
3. In Internet Explorer, in the Address bar, type **https://nls.adatum.com/**.
4. Verify that the default IIS 8 web page for LON-SVR1 displays.
5. Open a File Explorer window. In the address bar, type **\\Lon-SVR1\Files**, and then press Enter.
6. Verify that the contents of the **Files** shared folder display.
7. Close all open windows except the command prompt.

▶ Task 12: Move LON-CL1, and verify connectivity to intranet resources

1. On LON-CL1, change the network adapter configuration to the following settings:
   - IP address: **131.107.0.10**
   - Subnet mask: **255.255.0.0**
   - Default gateway: **131.107.0.2**
2. Disable and then re-enable the **Ethernet** adapter.
3. Close the Network Connections window.
4. On your host, in Hyper-V Manager, right-click **20413C-LON-CL1**, and then click **Settings**. Change the Legacy Network Adapter to be on the **Private Network 2** network, and then click **OK**.
5. Open Internet Explorer, and in the Address bar, type **http://lon-svr1.adatum.com/**. You should see the default IIS 8 web page for LON-SVR1 display.
6. Open File Explorer, in the address bar, type **\\LON-SVR1\Files**, and then press Enter. A folder window with the contents of the Files shared folder should display.
7. At a command prompt, type the following command, and then press Enter:

   ```
   ping lon-dc1.adatum.com
   ```

   Verify that you are receiving replies from lon-dc1.adatum.com.
8. At the command prompt, type the following command, and then press Enter:

   ```
   gpupdate /force
   ```

9. Close all open windows.
10. Switch to LON-RTR.
11. Open the Remote Access Management console, and review the information on **Remote Client Status**.

Note: Notice that LON-CL1 is connected through IP-HTTPS. In the Connection Details pane, in the bottom-right of the screen, note the use of Kerberos authentication for the machine and the user.

12. Close all open windows.

▶ Task 13: Revert virtual machines

When you are finished with this portion of the lab, before you continue, you must revert all virtual machines to their initial state. To do this, perform the following steps:

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-CL1**, and then click **Revert**.
3. In the **Revert Virtual Machines** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20413C-LON-DC1**, **20413C-LON-SVR1**, and **20413C-LON-RTR**.
5. In the **Virtual Machines** list, right-click **20413C-LON-DC1**, and then click **Connect**.
6. Click **Start**.
7. Log on to LON-DC1 as **Adatum\Administrator** with the password of **Pa$$w0rd**.
8. Repeat steps 5 through 7 for **20413C-LON-SVR1**, **20413C-LON-RTR**, **20413C-LON-CL1**, and **20413C-LON-CL2**.

Results: After completing this exercise, you should have planned and implemented a DirectAccess solution.

Exercise 3: Planning and Implementing a VPN Solution

Scenario

For those computers that do not meet DirectAccess criteria, you must implement a VPN solution.

VPN Strategy	
Document Reference Number: BS0905/1	
Document Author Date	Brad Sutton 5th September
Requirements Overview Design a remote access strategy to support the following objectives: • Data security is very important. • The current service provider for Internet access provides no guarantees for availability. Availability guarantees are required for disaster recovery planning. • Where possible, computers must be brought into the management scope of the A. Datum or Trey Research networks. • Executives are allowed remote access to network resources, and are not restricted. • Branch management staff is allowed remote access only to resources in their main office site. For example, branch managers in Europe are allowed access only to Paris resources. • Customer Service staff are not allowed remote access.	

VPN Strategy

- Marketing staff are allowed remote access, but only for email.

Additional Information

- The current VPN deployment consists of a single VPN server.
- Clients use PPTP connections, and have connectivity to the entire network when connected.
- A.Datum has an infrastructure in place for deploying certificates and smart cards.
- Some executives have had problems with hotel firewalls blocking VPN connections.
- Users from non-European sites have complained about slow access to data over the VPN.
- There is only a single Internet connection for Contoso, which is currently located in Paris.

Proposals

1. What tunneling protocols will you use?
2. What authentication or encryption methods do you recommend?
3. How many network policies do you envision?
4. List the network policies and their characteristics.
5. In what order will these policies process?
6. What certificates are required?

The main tasks for this exercise are as follows:

1. Read the supporting documentation.
2. Update the proposal document with your planned course of action.
3. Examine the suggested proposals in the Lab Answer Key.
4. Discuss your proposed solution with the class, as guided by your instructor.
5. Install and configure the Remote Access role.
6. Create the required network policies.

7. Create a client VPN.
8. Test VPN access.

▶ Task 1: Read the supporting documentation

- Read the documentation provided.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the VPN Strategy document.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones in the Lab Answer Key.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

▶ Task 5: Install and configure the Remote Access role

1. Switch to LON-RTR.
2. Sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
3. Using Server Manager, install the **Network Policy and Access Services** role. using default values to complete the installation wizard.
4. Open the Network Policy Server console, and then register the server in AD DS.
5. Leave the Network Policy Server console open.
6. Open the Routing and Remote Access console, and then disable routing and remote access.
7. Reconfigure routing and remote access as a VPN server.

▶ Task 6: Create the required network policies

1. Switch to LON-RTR.
2. Switch to the Network Policy Server console.
3. Under Policies, locate **Network Policies**, and disable the two existing network policies.

Note: These two policies would interfere with the processing of the policy that you are about to create.

4. Create a new network policy by using the following properties:
   - Policy name: **Adatum VPN Policy**
   - Type of network access server: **Remote Access Server (VPN-Dial up)**
   - Condition: **NAS Port Type**
   - NAS Port Type: **Virtual (VPN)**
   - Permission: **Access granted**
   - Authentication methods: Default settings
   - Constraints: Default settings
   - Settings: Default settings

▶ Task 7: Create a client VPN

1. Switch to LON-CL2 and sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. Create a new VPN connection with the following properties:
    - Internet address to connect to: **10.10.0.1**
    - Destination name: **Adatum VPN**
    - Allow other people to use this connection: Enabled
3. After you have created the VPN, open the properties of the connection.
4. On the **Security** tab, use the following settings to reconfigure the VPN:
    - Authentication: **Allow these protocols =Microsoft CHAP Version 2 (MS-CHAP v2)**

▶ Task 8: Test VPN access

- Test the VPN connection by using the following credentials:
    - User name: **Adatum\Administrator**
    - Password: **Pa$$w0rd**

Results: After completing this exercise, you should have planned and implemented a VPN solution.

Exercise 4: Implementing Web Application Proxy

Scenario

You need to implement Web Application Proxy to enable external users to access applications at A. Datum. You will use the initial deployment as a proof of concept while the developers at A. Datum modify the internal applications to use claims-based authentication.

The main tasks for this exercise are as follows:

1. Install the Active Directory Federation Services (AD FS) role.
2. Install the Web Application Proxy role.
3. Configure access to an internal website.
4. Verify access to the internal website.

▶ Task 1: Install the Active Directory Federation Services (AD FS) role

1. On LON-DC1, open **Windows PowerShell**. At the command prompt, run the following command:

```
Add-KdsRootKey –EffectiveImmediately
```

2. In Server Manager, add the **Active Directory Federation Services** server role.
3. Run the AD FS Federation Server Configuration Wizard by using the following parameters:
    - Create a new Federation Service.
    - Create a stand-alone deployment.
    - Use the **LON-DC1.Adatum.com** certificate.
    - Choose a service name of **LON-DC1.Adatum.com**.

4. Open Windows PowerShell, type the following command to enable modification of the assigned certificates, and then press Enter:

   ```
   Set-ADFSProperties –AutoCertificateRollover $False command
   ```

5. In the AD FS Management console, add the **LON-DC1.Adatum.com** certificate as a new token-signing certificate. Verify that the certificate has a subject of **CN=LON-DC1.Adatum.com** and its purposes include **Proves your identity to a remote computer** and **Ensures the identity of a remote computer**.

6. Make the new certificate the primary certificate, and remove the old certificate.

▶ Task 2: Install the Web Application Proxy role

1. Switch to LON-RTR.

2. Open Server Manager, and then on the **Dashboard** page, click **Add roles and features**.

3. In the Add Roles and Features Wizard, on the **Select server roles** page, expand **Remote Access**, and then click **Web Application Proxy**.

4. Verify that the installation is successful.

▶ Task 3: Configure access to an internal website

1. On LON-RTR, open an MMC console.

2. Add the **Certificates - Computer account** snap-in, and then request a new certificate with the following settings:

   - Subject Name: **Common Name**: **lon-dc1.adatum.com**

   - Alternative name:
     - **DNS**: **lon-dc1.adatum.com**,
     - **enterpriseregistration.adatum.com**
     - **lon-svr1.adatum.com**

3. Switch to LON-SVR1.

4. Open an MMC console, and add the **Certificates - Computer account** snap-in.

5. Request a new certificate with the Subject Name **Common Name**: **lon-svr1.adatum.com**.

6. Open the Internet Information Services (IIS) Manager console, navigate to **LON-SVR1/Sites**, and then click **Default Web site**.

7. Configure site bindings by entering **lon-svr1.adatum.com** as a host name and selecting **lon-svr1.adatum.com** as the **SSL Certificate**.

8. Close the Internet Information Services (IIS) Manager console.

9. Switch to LON-RTR.

10. From Server Manager, open the Remote Access Management console. In the navigation pane, click **Web Application Proxy**, and run the Web Application Proxy Configuration Wizard.

11. In the Web Application Proxy Configuration Wizard, for **Federation service name**, type **lon-dc1.adatum.com** for the FQDN of the federation service.

12. In the **User name** and **Password** text boxes, type **Administrator** and **Pa$$w0rd**.

13. On the **AD FS Proxy Certificate** page, in the list of certificates currently installed on the Web Application Proxy server, select the **lon-dc1.adatum.com** certificate.

14. On the **Results** page, verify that the configuration was successful, and then close the wizard.

15. In the Remote Access Management console, in the navigation pane, start the Publish New Application Wizard.
16. On the **Preauthentication** page, click **Pass-through**.
17. In the **Name** text box, type **LON-SVR1 Web** as the friendly name for the application.
18. In the **External URL** box, type **https://lon-svr1.adatum.com** as the external URL for this application.
19. In the **External certificate** list, select the **lon-dc1.adatum.com** certificate.
20. In the **Backend server URL** box, ensure that **https://lon-svr1.adatum.com** is listed.

Note: Note that this value is entered automatically when you enter the external URL.

21. On the **Confirmation** page, review the settings, and then click **Publish**.
22. On the **Results** page, ensure that the application published successfully, and then click **Close**.
23. Switch to LON-SVR1.
24. From Server Manager, open the Internet Information Services (IIS) Manager console, and navigate to **Default Web Site**.
25. Configure **Authentication** for the **Default Web Site** with following settings:
    - Windows Authentication: Enabled
    - Anonymous Authentication: Disabled
26. Close the Internet Information Services (IIS) Manager console.

▶ Task 4: Verify access to the internal web site

1. Switch to LON-CL1.
2. Open Control Panel.
3. In Control Panel, remove LON-CL1 from the **Adatum.com** domain. Add LON-CL1 to a workgroup named **WORKGROUP**, and then restart LON-CL1.
4. Sign in with user name **Admin** and password **Pa$$w0rd**.
5. On LON-CL1, use Notepad to create a file named **Hosts**.
6. Add the following content to the hosts file: **131.107.0.2 lon-svr1.adatum.com**
7. Copy the hosts file to **C:\Windows\System32\drivers\etc**.
8. On LON-CL1, open Internet Explorer, and then type the following address **https://lon-svr1.adatum.com**.
9. When prompted, in the **Internet Explorer** dialog box, type **Adatum\Bill** for the user name and **Pa$$w0rd** for the password, and then click **OK**.
10. Verify that the default IIS 8.0 web page for LON-SVR1 opens.
11. If you are unable to connect to https://lon-svr1.adatum.com, perform the following steps:
    a. On LON-CL1, on the **Start** screen, type **cmd**, and then press Enter.
    b. In the command prompt window, type **regedit**, and then press Enter.
    c. In the **User Account Control** dialog box, click **Yes**.

d. In the Registry Editor window, in the navigation pane, navigate to **HKLM\Software\Policies\Microsoft\Windows NT\DNSClient\DNSPolicyConfig**.

e. Right-click each of the entries starting with **DA**, click **Delete**, and in the **Confirm Key Delete** dialog box, click **Yes**.

f. Close the Registry Editor window.

12. Restart LON-CL1 and perform steps 8 through 10 to verify connectivity to default IIS 8.0 web page on LON-SVR1.

Results: After completing this exercise, you should have implemented a Web Application Proxy solution.

▶ Task: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-CL1**, and then click **Revert**.
3. In the **Revert Virtual Machines** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for the following machines: 20413C-LON-DC1, 20413C-LON-CL2, 20413C-LON-RTR, and 20413C-LON-SVR1.

Module Review and Takeaways

Review Question(s)

Question: Which type of policy can you use to determine whether a network connection attempt will be successful?

Question: When configuring home computers to enable access to corporate email, which of the following is generally a better approach for enabling remote access: RPC over HTTPS, or a VPN?

Question: In a mixed client environment that requires strong levels of security, which of the following VPN tunnel types would you select: PPTP, L2TP/IPsec, SSTP, or IKEv2?

Question: True or False? The NPS server role can function as a RADIUS client.

Question: Which of the following is a more secure firewall solution: bastion host, multi-homed firewall, or back-to-back firewalls?

Question: What function does the network location server have in a DirectAccess solution?

Question: In what ways can DirectAccess clients connect to network resources?

Course Evaluation

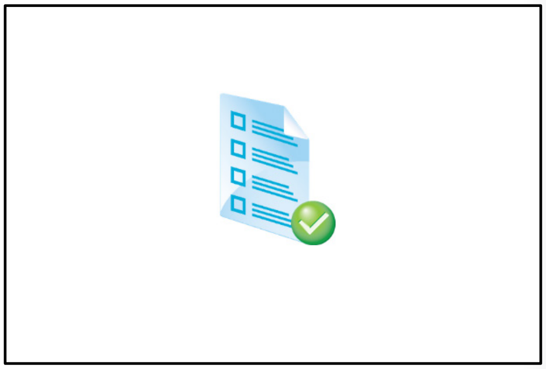

Keep this evaluation topic page if this is the final module in this course. Insert the Product_Evaluation.ppt on this page.

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: Planning Server Upgrade and Migration

Lab: Planning a Server Upgrade and Migration

Exercise 1: Planning a Strategy for Server Upgrade and Migration

▶ Task 1: Read the supporting documentation

- Read the supporting documentation in the lab exercise scenario.

▶ Task 2: Update the proposal document with your planned course of action

- Analyze the internal and perimeter networks separately. Do this because of the different security configurations and settings for each of the networks.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Examine the suggested proposals in the Lab Answer Key.

1. You plan to run the MAP to help you decide on a server-consolidation strategy. What result do you expect to get from this tool?

As a first step, you should run the MAP tool to analyze the inventory of the organizations' server infrastructure, perform an assessment, and create reports that you can use in your upgrade and migration plans. As an alternative to running MAP, you could analyze the average CPU and memory utilization of all server computers.

2. Besides using the MAP tool, what considerations would help you determine which machines you would move to the virtual environment?

Based on the MAP results, or manual analysis, you should plan which servers are candidates for virtualization.

Usually you would choose the physical machines with less than average processor utilization and disk storage.

3. What is your decision regarding virtualization of domain controllers?

The first candidates for virtualization would be the domain controllers, because of their small memory utilization. However, you must be careful not to store both of the domain controller virtual machines on the same host. This is because the domain controllers then would have a single point of failure: the physical host.

4. What is your decision regarding virtualization of the LON-IF1 and LON-IF2 infrastructure servers?

You should plan to virtualize those infrastructure servers that are running DNS and DHCP—LON-INF1 and LON-INF2—because both have an average utilization of 50 percent CPU. Be careful not to store both infrastructure server virtual machines on the same host. This is because the domain controllers would have single point of failure: the physical host.

5. Do the virtualized machines require high availability?

Yes. A. Datum has addressed high availability of domain controllers and infrastructure servers by deploying two domain controllers, LON-DC1 and LON-DC2, and two infrastructure servers, LON-INF1 and LON-INF2.

6. What system resources, such as processors, memory, or disk space should you allocate to server roles?

You should consider analyzing the amount of resources required for server roles that have been collocated. Furthermore, if you plan to virtualize some of the servers, you would allocate more physical resources to the host machine memory, as necessary.

7. What are the best virtualization candidates on the internal network, considering current physical server utilization and high availability requirements?

On the internal network, you might consider the following servers for virtualization, because of their low memory and the low CPU utilization on LON-CA: LON-DC1, LON-DC2, LON-DC3, LON-INF1, and LON-INF2. Because of high availability, you should place the following virtual machines on two or more physical hosts with Hyper-V failover clustering configuration: LON-CA: LON-DC1, LON-DC2, LON-DC3, LON-INF1, and LON-INF2.

8. What are your plan's licensing considerations for internal network servers? Do these considerations have an impact on the host operating system?

If the organization deploys those six virtual machines on two physical servers, you would require three instances of Windows Server 2012 Standard edition, where each physical server is licensed for up to two virtual machines. Consequently, using Windows Server 2012 Datacenter edition would be more suitable, because it supports additional virtual licenses.

9. What are the best virtualization candidates on the perimeter network, considering current physical server utilization and high availability requirements?

On the perimeter network, the candidates for virtualization are LON-PER-NS1, LON-PER-NS2, LON-RADIUS, LON-VPN1, and LON-VPN2, because of their low memory and CPU utilization. As with the internal network, because of high availability, you should not place virtual machines on the perimeter network onto a single host, but on two or more physical hosts.

10. What are your plan's licensing considerations for perimeter network servers? How does this impact the selection of the host operating system?

If the organization deploys those five virtual machines on two physical servers, they would require two instances of Windows Server 2012 Datacenter edition, where each physical server is licensed for unlimited virtual machines.

11. How would you manage licensing and activation?

To manage licensing and activation, you should implement volume licensing based on Active Directory activation. This is because the organization deploys Windows 8 and Windows Server 2012 operating systems.

12. Are there any servers which you should not cohost?

You should not cohost the enterprise root certification authority (CA) role because it would spend the majority of the time offline. However, any subordinate CA server role might be a good candidate for cohosting with another server role, such as the AD DS server role.

13. Sketch your plan's relevant network areas.

On the internal network, LON-HOST1 will host the following virtualized servers: LON-DC1, LON-INF1, and LON-CA. LON-HOST2 will host the following virtualized servers: LON-DC2 and LON-INF2. On the perimeter network, LON-HOST3 will host the following virtualized servers: LON-VPN1 and LON-PER-NS1. LON-HOST4 will host the following virtualized servers: LON-VPN2, LON-PER-NS2, and LON-RADIUS.

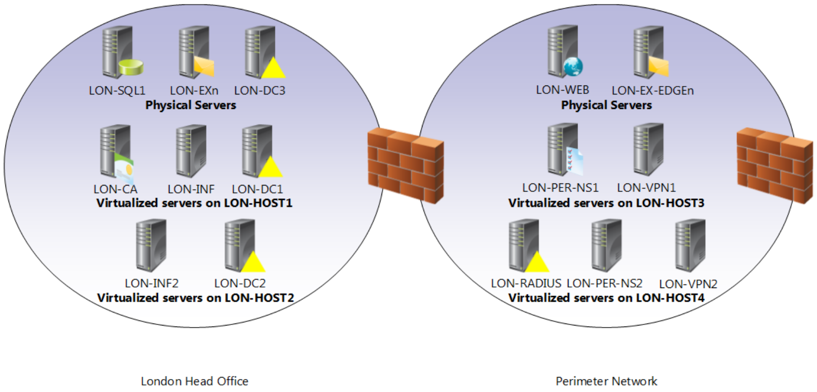

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor.

1. What was your approach to the design plan?

   Answers will vary.

2. Did your design plan differ from the suggested solution?

   Answers will vary.

Results: After completing this exercise, you will have planned a server upgrade and migration successfully.

Exercise 2: Evaluating Candidates for Server Virtualization

▶ Task 1: Evaluate server virtualization candidates

1. On LON-CL1, on the **Start** screen, click the **Microsoft Assessment and Planning Toolkit** tile.
2. In the **Microsoft Assessment and Planning Toolkit** console a dialog box will appear named **Microsoft Assessment and Planning Toolkit**.
3. In the **Microsoft Assessment and Planning Toolkit** dialog box, click **Manage**, click **Import**, and then click **Browse**.
4. In the **Microsoft Assessment and Planning Toolkit** dialog box, on the left pane expand **C:\Program Files\ Microsoft Assessment and Planning Toolkit\Sample**, and then on the right pane click on **MAP_SampleDB.bak** and then click **Open**.
5. In the **Microsoft Assessment and Planning Toolkit** dialog box, in the **Database Name** box, type **MAPDEMO**, and then click **OK**.
6. When the dialog box displays message that databases have been successfully imported, click **OK**, and then click **Close**.
7. In **Microsoft Assessment and Planning Toolkit** window, click **Use an existing database**, select **MAPDEMO**, click **OK**, and then click **Close**.
8. On LON-CL1, from the **Overview** page, on the navigation pane on the left, click **Server Virtualization**.
9. Under the **Steps to complete** section, click **Run the Server Consolidation Wizard**.
10. In the **Server Virtualization and Consolidation Wizard**, on the **Virtualization Technology** page, click **Windows Server 2012 Hyper-V**, and then click **Next**.
11. On the **Hardware Configuration** page, click **Sample host**, and then click **Next**.
12. On the **Utilization Settings** page, in each field, type **75**, and then click **Next**.
13. On the **Choose Computers** page, click **Choose the computers from a list on the next step of the wizard**, and then click **Next**.
14. On the **Computer List** page, select the **Computer Name** check box, and then click **Next**.
15. On the **Summary** page, review the settings, and then click **Finish**.
16. When the assessment process completes, click **Close**.
17. In the MAP console, on the **Server Virtualization** page, under Scenarios, click **Server Consolidation**, and review the **Details** section. Under the **Options** section, click the **Server Virtualization Report**, and when the proposal is generated, click **Close**.
18. On the **MAP console** menu, click **View**, and then click **Saved Reports and Proposals**.
19. In the File Explorer window, right-click the **ServerVirtRecommendation** workbook, and then click **Open**.
20. At the bottom of the Microsoft Excel® workbook, click each tab, and review the information in the report.
21. When finished close **Excel**, and then close **File Explorer**.
22. Spend time reviewing the MAP toolkit. Review at least three scenarios for which MAP provides information.

23. Be prepared to answer discussion questions based on your findings.

▶ Task 2: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for 20413C-LON-CL1.

Results: After completing this exercise, you will have installed MAP and used a sample database to evaluate which servers are virtualization candidates.

Module 2: Planning and Implementing a Server Deployment Strategy

Lab: Planning and Implementing a Server Deployment Infrastructure

Exercise 1: Planning an Automated Server Installation and Deployment Strategy

▶ Task 1: Read the supporting documentation

- Read the documentation in the lab Exercise Scenario.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the A. Datum Automated Server Installation and Deployment Strategy document.

1. What kind of image will you use: thin or thick?

   In this scenario, because there is no requirement for custom applications to be present in the Windows Server® 2012 R2 image, you can use thin images. For this deployment, you need the default boot image and the default install image from the Windows Server 2012 R2 media.

2. Would lite-touch or zero-touch deployment be applicable for this scenario?

   You can use both strategies to achieve the desired deployment scenario; however, the lite-touch scenario requires less infrastructure prerequisites.

3. Which deployment technologies would you consider to implement the server upgrade plan?

   In this scenario, because there are no requirements to upgrade settings from existing servers, you can consider implementing deployment by using Windows® Deployment Services (Windows DS). (Possible answers might include usage of Microsoft Deployment Toolkit (MDT) for customized installation, or enterprise companies will consider using deployment with Microsoft® System Center 2012 R2 Configuration Manager.)

4. What are the requirements for implementing this deployment technology?

   To implement the deployment with Windows DS, ensure the following requirements are met: Dynamic Host Configuration Protocol (DHCP), Domain Name System (DNS), and Windows DS. For Active Directory® integration, Active Directory Domain Services (AD DS) is required.

5. Produce a list of the deployment components necessary to support your server deployment plan.

   LON-DC1 is the server hosting the AD DS, DNS, and DHCP roles, which are required for deploying Windows Server 2012 R2. LON-SVR1 is the server configured with Windows DS to respond to all known and unknown computers. However, to increase security, you will require administrator approval for all unknown computers.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones shown previously.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you should have planned an automated server installation and deployment strategy for A. Datum Corporation.

Exercise 2: Preparing the Windows Server 2012 R2 Image

▶ Task 1: Create the image store, and map a network drive to the image store

1. Switch to LON-SVR1.
2. On the desktop, on the taskbar, click the **File Explorer** icon.
3. In File Explorer, in the navigation pane, expand **This PC**, click **Allfiles (E:)**, right-click the details pane, click **New**, click **Folder**, in the **New Folder** text box, type **Images**, and then press Enter.
4. In File Explorer, in the navigation pane, double-click **Images**, right-click the details pane, and then click **New**. Click **Folder**, in the **New Folder** text box, type **Custom Images**, and then press Enter.
5. On your host, in the 20413C-LON-SVR1 window, on the toolbar, click **Media**, point to **DVD Drive**, and then click **Insert Disk**.
6. In the **Open** dialog box, in the **File name** text box, type **D:\Program Files\Microsoft Learning\20413\Drives\Windows2012R2.iso**, and then click **Open**.
7. Copy **D:\sources\install.wim** into the **E:\Images\Custom Images** folder.
8. In File Explorer, right-click **E:\Images**, and then click **Properties**.
9. In the **Images Properties** dialog box, click the **Sharing** tab, and then click **Advanced Sharing**.
10. In the **Advanced Sharing** dialog box, select the **Share this folder** check box.
11. Click **Permissions**, and then click **Add**.
12. In the **Select Users, Computers, Service Accounts, or Groups** dialog box, in the **Enter the object names to select (examples)** text box, type **Administrator**, and then click **OK**.
13. In the **Permissions for Images** dialog box, click **Administrator (ADATUM\Administrator)**, in the **Allow** column, select the **Full Control** check box, and then click **OK**.
14. In the **Advanced Sharing** dialog box, click **OK**, and then click **Close**.
15. In File Explorer, right-click **This PC**, and then click **Map network drive**.
16. In the **Map network drive** dialog box, in the **Folder** text box, type **\\lon-svr1\Images**, and then click **Finish**.

▶ Task 2: Mount the relevant image

1. On LON-SVR1, on the taskbar, click the **Start** icon, and then type **cmd.exe**.
2. In the **Apps** list, right-click **cmd.exe**, and then click **Run as administrator**.
3. At the command prompt, type the following command, and then press Enter:

```
Mkdir c:\mounted
```

4. At the command prompt, type the following command, and then press Enter:

```
Dism /get-imageinfo /imagefile:”z:\Custom Images\install.wim”
```

Ensure that **The operation completed successfully** message displays.

Note: This command lists all images that are contained in install.wim. Notice the index for the Windows Server 2012 R2 Datacenter edition.

5. At the command prompt, type the following command, and then press Enter:

```
Dism /mount-wim /wimfile:”z:\Custom Images\install.wim” /index:4
/mountdir:c:\mounted
```

Ensure that **The operation completed successfully** message displays.

Note: This command will mount the install.wim image for offline servicing. After you mount the image, you can add drivers, add packages, or enable features. This step will take approximately five minutes for the mounting of the image finish.

▶ Task 3: Add the Web Server (IIS) role to the image

1. At the command prompt, type the following command, and then press Enter:

```
Dism /image:c:\mounted /get-features
```

Note: This command lists all available features and their state in the image file. If you want to see more detail, redirect the output in the text file by using the following command:
Dism /image:c:\mounted /get-features > c:\All Features.txt

2. At the command prompt, type the following command, and then press Enter:

```
Dism /image:c:\mounted /get-featureinfo /featurename:IIS-WebServerRole
```

Note: Notice that the state of Web Server (IIS) role is disabled.

3. At the command prompt, type the following command, and then press Enter:

```
Dism /image:c:\mounted /enable-feature /featurename:IIS-WebServerRole –all
```

Ensure that **The operation completed successfully** message displays.

Note: This command will install the Web Server (IIS) role with all the depended features. Note that the name of the role is case sensitive.

4. At the command prompt, type the following command, and then press Enter:

```
Dism /unmount-wim /mountdir:c:\mounted /commit
```

Ensure that **The operation completed successfully** message displays.

Note: This command commits the changes in the install.wim image, which you will use later for deploying to a machine that has no operating system installed. This step (to commit the changes in install.wim) can take approximately five minutes.

5. Do not close the Command Prompt window.

Results: After completing this exercise, you should have prepared the image, and added the Web Server (IIS) role to the image.

Exercise 3: Deploying Windows Server 2012 R2

▶ Task 1: Install the Windows DS role

1. On LON-SVR1, switch to Server Manager.
2. In the Server Manager console, click **Manage**, and then 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, in the **Roles** list, select the **Windows Deployment Services** check box.
7. In the **Add Roles and Features Wizard** dialog box, click **Add Features**, and then click **Next**.
8. On the **Select features** page, click **Next**.
9. On the **WDS** page, click **Next**.
10. On the **Select role services** page, verify that both the **Deployment Server** and **Transport Server** check boxes are selected, and then click **Next**.
11. Click **Install** to finish installation of Windows DS. When installation is complete, click **Close**.

▶ Task 2: Configure Windows DS

1. In Server Manager, click **Tools**, and then click **Windows Deployment Services**.
2. In the **Windows Deployment Services** console, in the navigation pane, expand **Servers**, right-click **LON-SVR1.Adatum.com**, and then click **Configure Server**.
3. In the Windows Deployment Services Configuration Wizard, on the **Before You Begin** page, click **Next**.
4. On the **Install Options** page, verify that **Integrated with Active Directory** is enabled, and then click **Next**.
5. On the **Remote Installation Folder Location** page, in the **Path** text box, type **E:\RemoteInstall**, and then click **Next**.
6. On the **PXE Server Initial Settings** page, verify that **Do not respond to any client computers** is enabled, and then click **Next**.
7. When the wizard finishes, clear the **Add images to the server now** check box, and then click **Finish**.

▶ Task 3: Use WDSUtil to add a boot image

1. Switch to the Command Prompt window.
2. At the command prompt, type the following command, and then press Enter:

   ```
   Wdsutil /add-image /ImageFile:”d:\sources\boot.wim” /ImageType:boot
   ```

3. Ensure that **The operation completed successfully** message displays.
4. Switch to the Windows Deployment Services console.
5. In the navigation pane, expand **LON-SVR1.Adatum.com**, and then click **Boot Images**.
6. Verify that one boot image for the 64-bit architecture is listed.

▶ Task 4: Add an install image

1. In the Windows Deployment Services console, right-click **Install Images**, and then click **Add Image Group**.
2. In the **Add Image Group** dialog box, in the **Name** text box, type **ImageGroup1**, and then click **OK**.
3. Right-click **ImageGroup1**, and then click **Add Install Image**.
4. In the Add Image Wizard, in the **File location** text box, type **Z:\custom images\install.wim**, and then click **Next**.
5. On the **Available Images** page, clear all check boxes, select the **Windows Server 2012 SERVERDATACENTER** check box, and then click **Next**.
6. On the **Summary** page, click **Next**, and then click **Finish** to import the boot image.

Note: This process of adding the install image will take 5–10 minutes.

▶ Task 5: Configure automatic naming

1. In Windows Deployment Services console, right-click **LON-SVR1.Adatum.com**, and then click **Properties**.
2. In the **Properties** dialog box, click the **PXE Response** tab, and then click **Respond to all client computers (known and unknown)**.
3. Click the **AD DS** tab, in the **Format** text box, type **LON-SVR%0#**, and then click **OK**.

▶ Task 6: Launch the deployment process

1. On the host computer, if necessary, switch to Hyper V Manager.
2. In Hyper V Manager, click **20413C-LON-SVR3**, and in the Actions pane, click **Connect**.
3. In the **20413C-LON-SVR3 on localhost – Virtual Machine Connection** dialog box, click **Start**.
4. When prompted, press the **F12** key for network service boot.
5. In the Windows Deployment Services Wizard, on the **Windows Deployment Services** page, click **Next**.
6. When the message box to connect to Windows DS server LON-SVR1.adatum.com displays, specify the following credentials, and then click **OK**:
   - Username: **Adatum\Administrator**
   - Password: **Pa$$w0rd**
7. On the **Select the operating system you want to install** page, click **Windows Server 2012 R2 SERVERDATACENTER**, and then click **Next**.
8. Leave the default drive selection, and then click **Next**.

Note: Steps 9 through 13 are optional. You can choose to not complete the deployment process, or you can follow these steps and complete the deployment.

9. When the computer is installed, on the **Settings** page, select the **I accept the license terms for using Windows** check box, and then click **Accept**.
10. On the **Region and Language** page, click **Next**.

11. For the built-in administrator account, in both the **Password** and **Reenter password** text boxes, type **Pa$$w0rd**, and then click **Finish**.
12. Sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
13. In Server Manager, in the navigation pane, verify the presence of the Web Server (IIS) role.

Results: After completing this exercise, you should have deployed Windows Server 2012 R2 by using Windows DS.

▶ Task: 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 Microsoft Hyper-V® Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20413C-LON-SVR1** and **20413C-LON-SVR3**.

Module 3: Planning and Deploying Servers Using Virtual Machine Manager

Lab: Planning and Deploying Virtual Machines by using Virtual Machine Manager

Exercise 1: Planning Microsoft® System Center 2012 R2 Virtual Machine Manager Components

▶ Task 1: Read the supporting documentation

Read the documentation provided in the Student Handbook.

▶ Task 2: Update the proposal document with your planned course of action

Answer the questions in the proposal section of the Adatum Virtualization Test Project:

Proposal Document for Adatum Virtualization Test Project

Domain controllers

Profile requirements: Eight domain controllers all running Windows Server® 2012 R2, four in London, and two each in Toronto and Sydney. What profiles do you recommend to meet these domain controller requirements?
Answer: You need two hardware profiles and one guest operating system profile: one profile for the requirements for the London domain controllers, and a hardware profile for the Toronto and Sydney domain controllers. You only need one guest operating system profile for all domain controllers because they all run Windows Server 2012 R2 Datacenter.

File servers

Profile requirements: Five file servers, three running Windows Server 2008 R2, and two running Windows Server 2012 R2; Hardware/performance requirements: two processor cores, 4 gigabytes (GB) of random access memory (RAM), two 2-terabyte (TB) drives for data. What profiles do you recommend to meet these file server requirements?
Answer: You need one hardware profile and two guest operating system profiles: one for Windows Server 2008 R2, and the other for Windows Server 2012 R2.

Web servers

Profile requirements: Three web servers, all running Windows Server 2012 R2. What profiles do you recommend to meet these web server requirements?
Answer: You need one hardware profile and one guest operating system profile.

Database servers

Profile requirements: Two Microsoft SQL Server® servers running Windows Server 2012 R2 and Microsoft SQL Server 2012. What profiles do you recommend to meet these requirements?
Answer: You need one hardware profile, one guest operating system profile, and one SQL Server profile.

Email servers

Profile requirements: Two Microsoft Exchange Server servers running Windows Server 2012 R2 and Exchange Server 2013. What profiles do you recommend to meet these email server requirements?
Answer: You need one hardware profile and one guest operating system profile.

Application servers

Profile requirements: Two servers running Windows Server 2008 R2 and Microsoft SharePoint® Server 2010. What profiles do you recommend to meet these application server requirements?
Answer: You need one hardware profile and one guest operating system profile.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

Compare your proposals with the ones shown in this document. Note that your answers may differ from the given answers, and are not necessarily incorrect, depending on the criteria you used.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

1. Go over your answers for the proposal with the class.
2. Explain any choices you made that differ from the answers given in this document.

Results: At the end of this exercise, you should have created hardware profiles, guest operating system profiles, and SQL Server profiles to meet requirements.

Exercise 2: Planning Virtual Machine and Service Templates

▶ Task 1: Read the supporting documentation

Read the documentation provided in the Student Handbook.

▶ Task 2: Update the proposal document with your planned course of action

Specify the types of profiles and templates needed and write them into the proposals section of the Adatum Server Service Template Requirement document.

Proposals

1. Create a hardware profile for the first virtual machine.
2. Create a hardware profile for the second virtual machine.
3. Create a guest operating system profile for the first virtual machine.
4. Create both guest operating system and SQL Server profiles for the second virtual machine.
5. Incorporate the above profiles into the service template.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

Compare your proposals with the ones in given previously.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

1. Go over your answers for the proposal with the class.
2. Explain any choices you made that differ from the lab answers.

Results: At the end of this exercise, you should have planned virtual machine and service templates.

Exercise 3: Implementing Virtual Machine Manager Components

▶ Task 1: Adding LON-HOST1 as a Host Server to VMM

Set the default domain Group Policy to allow domain members to become hosts

1. On LON-DC1, in Server Manager, click **Tools**, and then click **Group Policy Management**.
2. In the Group Policy Management Console, in the console tree, expand **Forest: Adatum.com**, expand **Domains**, and then expand **Adatum.com**. Under **Adatum.com**, right-click **Default Domain Policy**, and then click **Edit**.
3. In the Group Policy Management Editor, maximize the window. In the console tree, under Computer Configuration, expand **Policies**, and then navigate to the following location: **Administrative Templates\Network\Network Connections\Windows Firewall\Domain Profile**.
4. In the Domain Profile details pane, double-click **Windows Firewall: Allow inbound file and printer sharing exception**.
5. In the **Windows Firewall: Allow inbound file and printer sharing exception** pop-up window, click **Enabled**, in the **Options** text box, type an asterisk (*****), and then click **OK**.
6. In the Domain Profile details pane, double-click **Windows Firewall: Allow ICMP exceptions**.
7. In the **Windows Firewall: Allow ICMP exceptions** pop-up dialog box, select the **Enabled** radio button, in the Options area, select the **Allow inbound echo request** check box, and then click **OK**.
8. In the Domain Profile details pane, double-click **Windows Firewall**: **Define inbound port exceptions**.
9. In the **Windows Firewall: Define inbound port exceptions** pop-up dialog box, click **Enabled**. In the Options area, by Define port exceptions, click **Show**.
10. In the **Show Contents** pop-up dialog box, under **Value**, type **5985**, and then click **OK** twice.
11. In the Group Policy Management Editor, in the console tree, under Administrative Templates, expand **Windows Components**, expand **Windows Remote Management (WinRM)**, and then click **WinRM Service**.
12. In the WinRM Service details pane, double-click **Allow remote server management through WinRM**.
13. In the **Allow remote server management through WinRM** dialog box, click the **Enabled** radio button. In the Options area, in both the **IPv4** and **IPv6** text boxes, type an asterisk (*****), and then click **OK**.
14. Close the Group Policy Management Editor, and then close the Group Policy Management Console.
15. On LON-HOST1, on the taskbar, click the **Windows PowerShell** icon.
16. In the Windows PowerShell window, type the following cmdlet, and then press Enter:

    ```
    gpupdate /force
    ```

17. When both computer and user policies update successfully, close the Windows PowerShell window.

Add LON-HOST1 to VMM

1. On LON-VMM1, from the desktop, on the taskbar, click the **Virtual Machine Manager Console** icon.
2. On the **Connect to Server** page, click **Connect**.

3. In the Virtual Machine Manager console, click the **VMs and Services** workspace.
4. In the console tree, right-click **All Hosts**, and then click **Add Hyper-V Hosts and Clusters**.
5. In the Add Resource Wizard, on the **Resource Location** page, click the **Windows Server computers in a trusted Active Directory domain** option (it should be the default), and then click **Next**.

Note: On the **Credentials** page, note the two radio button options, **Use an existing Run As account** and **Manually enter the credentials**. The default setting is the **Use an existing Run As account** radio button, and then there is a field to type the **Run As account**, and a Browse button to browse to the account. Note that the **Run As account** must have local administrator permissions on the host machine being assigned. In this lab, you do not use a **Run As account**.

6. On the **Credentials** page, select the **Manually enter the credentials** radio button. In the **User name** text box, type **ADATUM\Administrator**, in the **Password** text box, type **Pa$$w0rd**, and then click **Next**.
7. On the **Discovery Scope** page, note that the **Specify Windows Server computers by names** radio button is already selected. In the **Computer names** text box, type **lon-host1.adatum.com**, and then click **Next**.
8. On the **Target resources** page, in the Discovered computers section, select the **lon-host1.adatum.com** check box, and then click **Next**.
9. When the **Virtual Machine Manager** pop-up window warns you that if the Microsoft Hyper-V server role is not enabled on the selected server, the VMM will enable Hyper-V, click **OK**.
10. On the **Host Settings** page, note that in the **Host group** drop-down list box, there is only one option, **All Hosts**. Note the **Reassociate this host with this VMM environment** check box. This setting takes hosts that have been assigned to a different VMM management server and assigns them to this one. Click **Next**.
11. On the **Summary** page, in the upper left, click the **View Script** button.
12. In Notepad, review the Windows PowerShell cmdlets that display.

Note: These are the cmdlets necessary to run the script in Windows PowerShell to add the LON-HOST1 host to this VMM management server. Saving these scripts can be very useful for documenting your work or for creating another host, perhaps at a later time.

13. Close **Notepad** without saving the script.
14. On the **Summary** page, click **Finish**.
15. A Jobs pop-up window appears, displaying all the individual steps being taken to add the host. The final step entitled **Add virtual machine host** takes the longest. It might take several minutes to complete the job.

Note: The Jobs pop-up window might display a yellow triangle with the text Add virtual machine host Completed w/ info. This occurs because Multipath I/O is not enabled for known storage arrays. This is expected.

16. When the job finishes, close the Jobs window.
17. In the VMs and Services console tree, under All Hosts, verify that LON-HOST1 now displays.

▶ Task 2: Create File Server guest operating system and hardware profiles

1. In the VMM console, click the Library workspace.
2. In the console tree, expand Profiles, and then click **Guest OS Profiles**. On the **Home** tab, click **Create**, and in the context menu, click **Guest OS Profile**.
3. In the New Guest OS Profile Wizard, on the **General** page, in the **Name** text box, type **FileServerGuestOS**, and then in the **Description** text box, type **Adatum File Server Guest OS profile**.
4. In the New Guest OS Profile Wizard console tree, click **Guest OS Profile**.
5. On the **General Settings** page, in the **Operating System** drop-down list box, click **Windows Server 2012 R2 Standard**.
6. Click the Identity Information section, and in the **Computer name** text box, type **FS1##**.
7. Click **Admin Password**, and in the details pane, select the **Specify the password of the local administrator account** check box.
8. In the text boxes for **Password** and **Confirm**, type **Pa$$w0rd**.
9. Click the **Roles** section, select the **File and Storage Services** check box, and then click **OK**.
10. In the console tree, click **Hardware Profiles**.
11. On the **Home** tab, click **Create**, and in the context menu, click **Hardware Profile**.
12. In the New Hardware Profile Wizard, on the **General** page, in the **Name** text box, type **FileServerHWProfile**, and then in the **Description** text box, type **Create File Server hardware profile**.
13. In the New Hardware Profile Wizard console tree, click **Hardware Profile**.
14. In the Compatibility section, select the **Hyper-V** check box.
15. In the central console tree, click **Memory**.
16. In the Memory details pane, click **Dynamic**, and in the Maximum memory area, replace the value that displays by typing **2048**.
17. In the center console tree, scroll down and click **Network Adapter 1**.
18. In the Network Adapter 1 details pane, click **Connected to a VM network**.
19. In the VM network area, click **Browse**, in the pop-up window, click **External Network**, and then click **OK**.
20. Click **OK** again.

▶ Task 3: Create developers' group web server guest operating system profile and hardware profile

1. In the VMM console tree, click **Guest OS Profiles**, on the **Home** tab, click **Create**, and in the context menu, click **Guest OS Profile**.
2. In the New Guest OS Profile Wizard, on the **General** page, in the **Name** text box, type **DevGroupWebGuestOS**, and then in the **Description** text box, type **Developers' Group Web Server Guest OS profile**.
3. In the New Guest OS Profile Wizard console tree, click **Guest OS Profile**.
4. On the **General** page, in the **Operating System** drop-down list box, click **Windows Server 2012 R2 Standard**.

5. Click the Identity Information section, and in the **Computer name** text box, type **DevWeb##**.
6. Click **Admin Password**, and in the details pane, select the **Specify the password of the local administrator account** check box.
7. In the text boxes for **Password** and **Confirm**, type **Pa$$w0rd**.
8. Click the **Roles** section, select the **Web Server (IIS)** check box, and then click **OK**.
9. In the console tree, click **Hardware Profiles**.
10. On the **Home** tab, click **Create**, and in the context menu, click **Hardware Profile**.
11. In the New Hardware Profile Wizard, on the **General** page, in the **Name** text box, type **DevGroupWebHWProfile**, and then in the **Description** text box, type **Create Developers' Group Web Server hardware profile**.
12. In the New Hardware Profile Wizard console tree, click **Hardware Profile**.
13. In the Compatibility section, select the **Hyper-V** check box.
14. In the central console tree, click **Memory**.
15. In the Memory details pane, click **Dynamic**, and in the **Maximum memory** area, replace the value that displays by typing **2048**.
16. In the center console tree, scroll down and click **Network Adapter 1**.
17. In the Network Adapter 1 details pane, click **Connected to a VM network**.
18. In the VM network area, click **Browse**, in the pop-up window, click **External Network**, and then click **OK**.
19. Click **OK** again.

▶ Task 4: Create developer's group database guest operating system profile and hardware profile

1. In the console tree, click **Guest OS Profiles**.
2. On the **Home** tab, click **Create**, and in the context menu, click **Guest OS Profile**.
3. In the New Guest OS Profile Wizard, on the **General** page, in the **Name** text box, type **DevGroupDBGuestOS**, and in the **Description** text box, type **Developers' Group Database Server Guest OS profile**.
4. In the New Guest OS Profile Wizard console tree, click **Guest OS Profile**.
5. On the **General** page, in the **Operating System** drop-down list box, click **Windows Server 2012 R2 Standard**.
6. Click the Identity Information section, and in the **Computer name** text box, type **DevDB##**.
7. Click **Admin Password**, and in the details pane, select the **Specify the password of the local administrator account** check box.
8. In the text boxes for **Password** and **Confirm**, type **Pa$$w0rd**, and then click **OK**.
9. In the console tree, click **Hardware Profiles**.
10. On the **Home** tab, click **Create**, and in the context menu, click **Hardware Profile**.
11. In the New Hardware Profile Wizard, on the **General** page, in the **Name** text box, type **DevGroupDBHWProfile**, and then in the **Description** text box, type **Create Developers' Group Database Server hardware profile**.

12. In the New Hardware Profile Wizard console tree, click **Hardware Profile**, and in the Compatibility section, select the **Hyper-V** check box.
13. In the central console tree, click **Memory**.
14. In the Memory details pane, click **Dynamic**, and in the Maximum memory area, replace the value that displays by typing **2048**.
15. Scroll down in the center console tree, and then click **Network Adapter 1**.
16. In the Network Adapter 1 details pane, click **Connected to a VM network**.
17. In the VM network area, click **Browse**.
18. In the pop-up window, click **External Network**, and then click **OK**.
19. Click **OK** again.

▶ Task 5: Create a developers' group database SQL Server profile

1. In the console tree, click **SQL Server Profiles**.
2. On the **Home** tab, click **Create**, and on the context menu, click **SQL Server Profile**.
3. In the New SQL Server Profile Wizard, on the **General** page, in the **Name** text box, type **DevGroupDBSQL**, and then in the **Description** text box, type **Developers' Group Database Server SQL Server profile**.
4. In the wizard's console tree, click **SQL Server Configuration**.
5. Next to Add, click **SQL Server Deployment**, and complete the form as follows and then click OK:

- Name: **DevDeploy**
- Instance Name: **MSSQLSERVER**
- Instance ID: **1**

6. In the console tree, click **Configuration**, in the **SQL Server administrators** text box, type **adatum\DevDBO**, and then click **Add**.
7. In the **Media Source** text box, type **D:\**.
8. In the System Administrator (SA) Password Run As Account section, click **Browse**.
9. In the Browse Run As Accounts pop-up window, click **Create Run As Account**.
10. In the **Create Run As Account** pop-up dialog box, complete the fields with the following settings:

- Name: **DevDBO**
- Description: **Developer's Group Database Owner**
- Username
- **Adatum\DevDBO**
- Password and Confirm Password: **Pa$$w0rd**

11. Clear the **Validate domain credentials** check box, and then click **OK** twice.
12. Select the **Use TCP/IP for remote connections** check box.
13. In the console tree, click **Service Accounts**. For each account, click **Browse**, click the **NT AUTHORITY\System** hyperlink, and then click **OK**.
14. On the **New SQL Server Profile** page, click **OK**.

▶ Task 6: Configure a virtual machine template

1. In the Virtual Machine Manager console, click the Library workspace.
2. In the console tree, expand **Templates**, and then click **VM Templates**.
3. On the ribbon, on the **Home** tab, click **Create VM Template**.
4. In the Create VM Template Wizard, on the **Select Source** page, to the right of the **Use an existing VM template or a virtual hard disk stored in the library** option, click **Browse**.
5. In the Select VM Template Source window, click **SmallCore.vhd**, and then click **OK**.
6. On the **Select Source** page, click **Next**.
7. On the **Identity** page, in the **VM Template name** text box, type **FSVMTemplate**, in the **Description** text box, type **Create the File Server VM template**, and then click **Next**.
8. On the **Configure Hardware** page, in the **Hardware profile** drop-down area, in the drop-down list box, click **FileServerHWProfile**, and then click **Next**.
9. On the **Configure Operating System** page, in the **Guest OS profile:** drop-down list box, click **FileServerGuestOS**, and then click **Next**.
10. On the **Application Configuration** page, in the **Application profile** drop-down list box, click **None – do not install any applications**, and then click **Next**.
11. On the **SQL Server Configuration** page, in the **SQL Server profile** drop-down list box, click **None – no SQL Server configuration settings**, and then click **Next**.
12. On the **Summary** page, click **Create**.
13. When the jobs finish, close the Jobs window.
14. Examine the **FSVMTemplate** in the Templates details pane. Note the items in the **Template** tab of the ribbon.

Note: From the Templates details pane, you can enable and disable the template, export its settings, and even delete it.

15. On the **Template** tab, click **Properties**.
16. In the **Properties** dialog box, observe that the **Hardware** and **OS Configuration** pages no longer point to the profiles created earlier, but now contain all the settings that you configured in the profiles.
17. In the **FSVMTemplate Properties** dialog box, click **Cancel.**

▶ Task 7: Create a virtual machine template from the developers' group web server profiles

1. In the Virtual Machine Manager console, click the **Library** workspace.
2. In the console tree, expand **Templates**, and then click **VM Templates**.
3. On the ribbon, on the **Home** tab, click **Create VM Template**.
4. In the Create VM Template Wizard, on the **Select Source** page, to the right of the **Use an existing VM template or a virtual hard disk stored in the library** option, click **Browse**.
5. In the Select VM Template Source window, click **SmallCore.vhd**, and then click **OK**.
6. On the **Select Source** page, click **Next**.

7. On the **Identity** page, in the **VM Template name** text box, type **DevGrpWebVMTemplate**, in the **Description** text box, type **Create the Developers' Group Web Server VM template**, and then click **Next**.

8. On the **Configure Hardware** page, in the **Hardware profile** drop-down area, in the drop-down list box, click **DevGroupWebHWProfile**, and then click **Next**.

9. On the **Configure Operating System** page, in the **Guest OS profile** drop-down list box, click **DevGroupWebGuestOS**, and then click **Next**.

10. On the **Application Configuration** page, in the **Application profile** drop-down list box, click **None – do not install any applications**, and then click **Next**.

11. On the **SQL Server Configuration** page, in the **SQL Server profile** drop-down list box, click **None – no SQL Server configuration settings**, and then click **Next**.

12. On the **Summary** page, click **Create**.

13. When the jobs finish, close the Jobs window.

14. In the Templates details pane, examine the template **DevGrpWebVMTemplate**. Note the items in the **Template** tab of the ribbon.

▶ Task 8: Create the developers' group web service template

1. In the Virtual Machine Manager console, click the Library workspace.

2. On the ribbon, on the **Home** tab, click **Create Service Template**.

3. In the **New Service Template** dialog box, in the **Name** text box, type **Dev Group Web Service Template**. In the **Release** text box, type **1**. In the Patterns section, click the **Blank** icon, and then click **OK**.

4. In the Virtual Machine Manager Service Template Designer console, note that the name you selected, **Dev Group Web Service Template**, is part of the overall name.

Note: This is the template that you are currently designing. The number 1 next to the name is the release version.

5. Click and drag the **DevGrpWebVMTemplate** name onto the Designer canvas area that displays the following message: **Drag VM Templates onto the canvas to create a new Tier and copy the VM Template settings into that tier**.

6. Note the box labeled **DevGrpWebVMTemplate - Machine Tier 1**.

7. The Designer canvas might display the **External Network** box with a connector spread out across the canvas. Drag the **External Network** box to beside the **NIC 1** box. It will shorten the connector.

▶ Task 9: Deploy virtual machines using the service template

1. On the **Home** tab, click **Save and Validate**, and then click **Configure Deployment**.

2. In the **Select name and destination** pop-up dialog box, in the **Name** text box, type **DevGroup Web Service**, and then click **OK**.

3. When the **Deploy Service – DevGroup Web Service** console displays, if you see a message in a pink shaded area in the middle of the screen saying it could not find a host, on the ribbon, click **Refresh Preview**.

Note: The host that VMM deploys to is based on placement ratings. However, since we only have one host, VMM deploys to lon-host1.adatum.com.

4. On the ribbon, click **Deploy Service**, and in the Deploy service pop-up window, click **Deploy**.
5. When the Jobs window displays, notice that you can see the Create Service Instance job running.

Note: This process can take between 15 to 30 minutes to complete.

6. After about 10 minutes, on LON-HOST1, open the Hyper-V console and verify that you see the virtual machine with a name of a **DEVWeb01**.
7. On **LON-VMM1**, close the Jobs window.
8. Close all open windows.

▶ Task 10: 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 LON-HOST1, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-DC1-B**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for 20413C-LON-VMM1.
5. Restart the host computer, and at the boot menu, select the Windows Server 2012 installation.

Results: At the end of this exercise, you should have implemented Virtual Machine Manager components.

Module 4: Designing and Maintaining an IP Configuration and Address Management Solution

Lab: Designing and Maintaining an IP Configuration and IP Address Management Solution

Exercise 1: Planning Dynamic Host Configuration Protocol (DHCP) to Support Your Proposal

▶ Task 1: Read the supporting documentation

- Read the documentation provided in the Student Handbook.

▶ Task 2: Update the proposals document with your planned course of action

- Answer the questions in the proposals section of the Contoso DHCP Deployment Strategy document.

Proposals

1. How should clients and servers in the head office in Paris obtain an IP configuration?

   You should configure all computers to obtain an IP address automatically from a DHCP server. The exception will be routers that you must configure manually. You should configure servers that need a specific IP address with a reservation. This would enable you to configure the server centrally, while still fixing the IP address.

2. How should clients in regional hub offices obtain an IP configuration?

   Client computers should obtain an IP configuration from a DHCP server.

3. How will you provide high availability for DHCP in the Paris office?

   You can configure high availability by configuring DHCP failover between two DHCP servers.

4. How will you provide high availability for DHCP in the regional hub sites?

   Potentially, the client and server computers in the regional hub sites can obtain their IP configurations from the DHCP servers in Paris. However, this does not provide protection against line failure. If the link between Paris and Athens fails, then DHCP would be unavailable.
   A possible solution would be to provide a DHCP server in each regional office to provide for address allocation in the regions.

5. How many scopes do you need to configure on the DHCP servers in the regional hub sites?

   You must configure a scope for each subnet. For example, Athens has five branches offices attached. This means you would need a minimum of six scopes.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones shown previously.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you will have planned DHCP to support the Contoso IP addressing scheme.

Exercise 2: Planning an IP Address Management (IPAM) Deployment

▶ Task 1: Read the supporting documentation

- Read the documentation provided in the Student Handbook.

▶ Task 2: Update the proposals document with your planned course of action

- Answer the questions in the proposals section of the A. Datum IPAM Deployment Plan document.

Proposals

1. Is a centralized or distributed topology better?

   Currently, you can implement only the A. Datum portion of the IPAM strategy, because Windows Server® 2012 is not yet deployed at Contoso and Trey Research. However, IPAM servers can monitor and manage DHCP, Domain Name System (DNS), domain controllers, and Network Policy Server (NPS) servers running Windows Server 2008.
   Because there are no Windows Server 2012 servers at either Contoso or Trey Research, initially at least, a centralized model would be better. In this way, the IPAM server at A. Datum (based in London) can manage the required roles elsewhere.
   As the Windows Server 2012 deployment plan proceeds and extends to Contoso and Trey Research, you could deploy additional IPAM servers at these locations.

2. Which server roles will you manage?

   You would manage all server roles: DNS, DHCP, and NPS servers.

3. What Active Directory® Domain Services (AD DS) considerations should you take into account for the inclusion of the Contoso and Trey Research organizations?

   IPAM can only manage roles within a single Active Directory forest. If Contoso is integrated as part of the same Active Directory forest, then there are no additional IPAM considerations. When you deploy AD DS, if you deploy Contoso as a separate Active Directory forest, then you will require additional IPAM servers.
   Trey Research currently has a separate Active Directory forest and requires its own IPAM deployment. However, Trey Research currently has no Windows Server 2012 servers.

4. What are the organizational and server-level prerequisites for IPAM?

   To ensure a successful IPAM implementation, your infrastructure must meet several prerequisites:

   - 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 enable IPv6 on the IPAM server.
   - You must sign in to 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.
   - You must ensure that the server meets the following requirements:
     - Dual core processor of 2.0 gigahertz (GHz) or higher
     - Windows Server 2012 operating system or newer

        - 4 or more gigabytes (GB) of random access memory (RAM)
        - 80 GB of free hard disk space
    - You must ensure servers running Windows Server 2008 R2 and Windows Server 2008 meet the following requirements:
        - 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 Beta must be installed (KB2506146)
        - For Windows Server 2008 SP2, Windows Management Framework Core (KB968930) is also required.
        - Windows Remote Management must be enabled.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones shown previously.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you will have planned an IPAM deployment strategy for A. Datum.

Exercise 3: Implementing DHCP and IPAM

▶ Task 1: Install the DHCP Server role

1. If necessary, sign in to LON-SVR1 as **Adatum\Administrator** with a password of **Pa$$w0rd**.
2. In Server Manager, in the results pane, click **Add roles and features**.
3. Click **Next** three times.
4. On the **Select server roles** page, click the **DHCP Server** role, and then click **Add Features**.
5. Click **Next** three times, and then click **Install**.
6. Once the role installs, click **Complete DHCP configuration**.
7. In the **DHCP Post-Install Configuration Wizard** dialog box, click **Next**.
8. On the **Authorization** page, click **Commit**.
9. On the **Summary** page, click **Close**.
10. To close the Add Roles and Features Wizard, click **Close**.

▶ Task 2: Configure a DHCP failover relationship

1. Switch to LON-DC1.
2. If necessary, sign in to LON-DC1 as **Adatum\Administrator** with a password of **Pa$$w0rd**.
3. In Server Manager, click **Tools**, and then in the drop-down list box, click **DHCP**.
4. In the DHCP console, expand **lon-dc1.adatum.com**, select and right-click **IPv4**, and then click **Configure Failover**.
5. In the Configuration Failover Wizard, click **Next**.
6. On the **Specify the partner server to use for failover** page, in the **Partner Server** text box, type **172.16.0.11**, and then click **Next**.
7. On the **Create a new failover relationship** page, in the **Relationship Name** text box, type **Adatum DHCP Failover**.
8. In the **Maximum Client Lead Time** field, set the hours to **0**, and then set the minutes to **15**.
9. Ensure that the **Mode** field is set to **Load balance**.
10. Ensure that the **Load Balance Percentage** is set to **50%**.
11. Select the **State Switchover Interval** check box, and then in the text box next to the check box, type **45**.
12. Select the **Enable Message Authentication** check box, in the **Shared Secret** text box, type **Pa$$w0rd**, and then click **Next**.
13. Click **Finish**, and then click **Close**.
14. Switch to LON-SVR1, open the DHCP console, and note that the IPv4 node is active.
15. Expand the **IPv4** node, and then expand **Scope**.
16. Click **Address Pool**, and note that the address pool is configured.
17. Click **Scope Options**, and note that the scope options are configured.
18. Close the DHCP console on both LON-SVR1 and LON-DC1.

▶ Task 3: Install IPAM

1. If necessary, sign in to LON-SVR2 as **Adatum\Administrator** with a password of **Pa$$w0rd**.
2. In Server Manager, in the results pane, click **Add roles and features**.
3. In the Add Roles and Features Wizard, click **Next** four times.
4. On the **Select features** page, select the **IP Address Management (IPAM) Server** check box.
5. In the **Add features that are required for IP Address Management (IPAM) Server** pop-up dialog box, click **Add Features**, and then click **Next**.
6. On the **Confirm installation selections** page, click **Install**.
7. When the Add Roles and Features Wizard completes, close the wizard.

▶ Task 4: Configure Group Policy Object (GPO) Settings

1. In the Server Manager navigation pane, click **IPAM**.
2. Click **Provision the IPAM server**.
3. In the Provision IPAM Wizard, click **Next**.
4. On the **Configure database** page, ensure that **Windows Internal Database (WID)** is selected, and then click **Next**.
5. On the **Select provisioning method** page, ensure that **Group Policy Based** is selected, in the **GPO name prefix** text box, type **IPAM**, and then click **Next**.
6. On the **Summary** page, click **Apply**.

Note: Provisioning may take five or more minutes to complete.

7. When provisioning completes, click **Close**.

▶ Task 5: Configure IP management server discovery

1. In the IPAM Overview pane, click **Configure server discovery**.
2. In the **Configure Server Discovery** dialog box, click **Add** to add the Adatum.com domain, and then click **OK**.
3. In the IPAM Overview pane, click **Start server discovery**.

Note: Discovery may take 5 to 10 minutes to run. The yellow bar indicates when discovery is complete.

▶ Task 6: 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 for both servers.
2. Scroll down to the Details view, and note the status report.

Note: The IPAM server has not yet been granted permission to manage LON-DC1 and LON-SVR1 through Group Policy.

3. On the taskbar, right-click the **Windows PowerShell** icon, and then click **Run as Administrator**.
4. At the Windows PowerShell® prompt, type the following command:

   ```
   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.

Note: The command may take five minutes or more to complete.

6. Close Windows PowerShell.
7. Switch back to Server Manager.
8. In the IPv4 details pane, right-click **lon-dc1**, and then click **Edit Server**.
9. In the **Add or Edit Server** dialog box, in the **Manageability status** drop-down list box, click **Managed**, and then click **OK**.
10. In the IPv4 details pane, right-click **lon-svr1**, and then click **Edit Server**.
11. In the **Add or Edit Server** dialog box, in the **Manageability status** drop-down list box, click **Managed**, and then click **OK**.
12. Switch to LON-DC1.
13. On the taskbar, click the **Windows PowerShell** icon.
14. At the Windows PowerShell command prompt, type the following command, and then press Enter:

    ```
    Gpupdate /force
    ```

15. After this command is successful, at the Windows PowerShell prompt, type the following command, and then press Enter:

    ```
    Gpresult /R
    ```

16. Verify that the IPAM_DNS, IPAM_DC_NPS, and IPAM_DHCP GPOs have been applied to the computer.
17. Close Windows PowerShell.
18. Switch to LON-SVR1.
19. On the taskbar, click the **Windows PowerShell** icon.
20. At the Windows PowerShell prompt, type the following command, and then press Enter:

    ```
    Gpupdate /force
    ```

21. After this command is successful, at the Windows PowerShell prompt, type the following command, and then press Enter:

    ```
    Gpresult /R
    ```

22. Verify that the IPAM_DHCP GPO has been applied to the computer.
23. Close the Windows PowerShell window.

24. Switch back to LON-SVR2.
25. In Server Manager, right-click **LON-DC1**, and then click **Refresh Server Access Status**.
26. In Server Manager, right-click **LON-SVR1**, and then click **Refresh Server Access Status**.
27. In Server Manager, refresh IPv4 by clicking the **Refresh** icon.

Note: This may take up to five minutes for the status to change.

28. In the IPAM Overview pane, click **Retrieve data from managed servers**.

Note: This action may take five or more minutes to complete.

▶ Task 7: Configure and verify a DHCP scope with IPAM

1. Switch to LON-SVR2.
2. In Server Manager, on the IPAM navigation pane, click **DNS and DHCP Servers**.
3. Right-click the record for **lon-dc1.adatum.com** for the **DHCP** server role, and then click **Create DHCP scope**.
4. In the **Create DHCP Scope** dialog box, in the **Scope name** text box, type **Paris Office**.
5. In the **Start IP address** text box, type **172.32.32.2**.
6. In the **End IP address** text box, type **172.32.32.200**.
7. In the **Subnet mask** text box, type **255.255.224.0**.
8. On the navigation pane, click **Options**.
9. In the **DHCP Scope Options** section, under **Configure options**, click **New**.
10. In the New Configuration pane, in the **Option** list, click **003 Router**.
11. In the **Values** list, under **IP address**, type **172.32.32.1**, and then click **Add Configuration**.
12. Under **Configure options**, click **New**.
13. In the New Configuration pane, in the **Option** list, click **006 DNS Servers**.
14. In the **Values** list, under **IP address**, type **172.32.32.2**, click **Add Configuration**, and then click **OK**.
15. In Server Manager, in the IPAM navigation pane, click **DHCP Scopes**.

Note: Notice the new DHCP scope that displays.

16. Switch to LON-DC1.
17. In Server Manager, click **Tools**, and then click **DHCP**.
18. In the DHCP console, expand **lon-dc1.adatum.com**, and then expand **IPv4**.

Note: Notice that the Paris Office scope is listed.

19. Switch to LON-SVR2.

20. In Server Manager, under IPAM, click **DHCP Scopes**, right-click the **Paris Office** scope, and then click **Configure DHCP Failover**.
21. In the **Configure DHCP Failover Relationship** dialog box, in the **Configuration option** list, click **Use an existing relationship**.
22. In the **Relationship name** list, click **Adatum DHCP Failover**, and then click **OK**.

Note: Notice that the added Paris Office scope displays in the list of scopes.

▶ Task 8: Configure IP address blocks, record IP addresses, and create DHCP reservations

1. On LON-SVR2, in Server Manager, in the IPAM pane, click **IP address Blocks**.
2. In the **Current view** list, click **IP Address Ranges**. Note that due to DHCP failover, 172.32.32.0 is listed twice.
3. Right-click the **172.32.32.0/19** range for **lon-dc1.adatum.com**, and then click **Edit IP Address Range**.
4. In the **Edit IP Address Range** dialog box, click **Reservations**.
5. In the **Reservation** text box, type **172.32.32.2**, and then click **Add**.
6. In the **Edit IP Address Range** dialog box, click **OK**.

Results: After completing this exercise, you will have deployed DHCP and IPAM to support your proposals.

▶ Task: 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 Microsoft Hyper-V® Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 through 3 for 20413C-LON-SVR1 and 20413C-LON-SVR2.

Module 5: Designing and Implementing Name Resolution

Lab: Designing and Implementing Name Resolution

Exercise 1: Designing a Strategy for DNS Name Resolution

▶ Task 1: Read the supporting documentation

- Read the documentation that the student handbook provides.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the Contoso DNS Name Resolution Strategy document.

Proposals

1. If you create a new design, what is your preferred namespace for Active Directory® Domain Services (AD DS)?

   The most simple approach is to use split DNS, so that both the internal Active Directory domain and the DNS domain have the same name, Contoso.com. An external DNS infrastructure exists already, as does a public web site, therefore you can house these resources in the perimeter network between the firewalls. You also can house the he virtual private network (VPN) and secure customer website in this location. However, the VPN would allow connections to pass through to the internal Active Directory–integrated DNS servers.

2. What additional factors should you consider when you are modifying an existing design?

   The most important consideration would be security for the internal Active Directory domain and DNS infrastructure. It is vitally important that unencrypted packets outside of the firewall do not resolve DNS names for the internal Active Directory resources.

3. What DNS namespace do you recommend that A. Datum Corporation use for AD DS?

   There are a number of factors to consider, foremost of which is, what the plan is for integrating Contoso into the A. Datum Active Directory forest. Currently that information does not exist, so you can only make suggestions. However, it would be logical to make the Active Directory domain names match the DNS domain names that you choose. In the Adatum.com DNS infrastructure, you can make Contoso.com and TreyResarch.Net DNS servers conditional forwarders for their respective domain names.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones shown previously.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you will have created a DNS name resolution design.

Exercise 2: Designing a Strategy for DNS Server Placement

▶ Task 1: Read the supporting documentation

Read the documentation provided in the student handbook.

▶ Task 2: Update the proposal document with your planned course of action

Answer the questions in the proposals section of the Contoso DNS Server Placement Strategy document.

Proposals

1. How many DNS servers do you require at the head office in Paris?

   Answers will vary. Currently, there are only two DNS servers for internal resolution. However, there could be nearly three times as many client computers at the Paris offices in the near future, once the A. Datum merger occurs. Therefore, it would be logical to measure the workload on the existing DNS servers, and then make a determination about additional servers. For now, two servers seem insufficient. When the Active Directory infrastructure is in place, an Active Directory–integrated zone should be part of the solution. Therefore, each domain controller in Paris also will be a DNS Server.

2. Do the branch locations require DNS servers?

   Currently, name resolution is possible only when the wide area network (WAN) link is operational to the regional hub. You could argue that with the link down, access to services could be affected to the extent that it would be largely irrelevant whether name resolution was functional in those circumstances.
   However, because each branch uses a local server, adding a DNS role to that server seems sensible. To support AD DS, deploying a read-only domain controller (RODC) to these branches seems highly probable. Therefore, there is no reason why you could not combine this role with that of the DNS Server role.
   Adding a name server at the branches reduces the need for query traffic over the WAN links, but does generate DNS zone-replication traffic. You can mitigate this largely by using Active Directory–integrated zones that rely on Active Directory replication for updates.

3. Are additional DNS servers required at each regional hub site?

   Currently, there is a single regional hub DNS server for each regional hub. These hubs each support only a few hundred users. A single DNS server is capable of supporting that workload. However, from a high availability perspective, this may not be sufficient. Therefore, you should deploy at least one more DNS server to the regional hub locations.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

Compare your proposals with the ones shown previously.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you will have determined where to place Contoso DNS servers to support your initial DNS design.

Exercise 3: Designing DNS Zones and DNS Zone Replication

▶ Task 1: Read the supporting documentation

Read the documentation that is provided in the student handbook.

▶ Task 2: Update the proposal document with your planned course of action

Answer the questions in the proposals section of the Contoso DNS Zones and Zones Replication Strategy document.

Proposals

1. Which zones do you need to create on internal DNS servers?

   Only one zone is required to support the single domain name for internal use. It does not matter whether this is a different domain name from the external name of Contoso.com or a subdomain of it. The internal DNS servers will host only a single zone.

2. Which zones do you need to create on external DNS servers?

   The Contoso.com zone exists already and is the only one that contains external records.

3. In which regional hub sites will you place each DNS zone?

   All of the regional hub sites will be part of the Contoso.com domain. By using Active Directory–integrated zones, the regional hub locations would get DNS updates through the Active Directory replication process.

4. How will you configure replication or zone transfers for each zone?

   Once you implement Active Directory, you can configure the zones as Active Directory–integrated zones. This means that Active Directory replication manages all zone replication automatically and securely.
   Until AD DS is deployed, you should configure zone transfers from the primary server to each secondary server at the head office and at the regional hubs. You should configure the regional hubs as master servers for the branch office DNS servers. Additionally, you should secure all DNS traffic with Internet Protocol security (IPsec).

5. How would implementing AD DS affect your design?

   As previously discussed, this will simplify the overall implementation of DNS, and of zones and replication.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

Compare your proposals with the ones shown previously.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you will have a DNS zone design that you can use to implement DNS.

Exercise 4: Implementing DNS

▶ Task 1: Install the DNS Server role

1. Switch to LON-SVR1.
2. If necessary, sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
3. In Server Manager, in the results pane, click **Add roles and features**.
4. In the Add Roles and Features Wizard, on the **Before you begin** page, click **Next**.
5. On the **Select installation type** page, click **Role-based or feature-based installation**, and then click **Next**.
6. On the **Select destination server** page, click **Select a server from the server pool**, and then click **Next**.
7. On the **Select server roles** page, in the **Roles** list, select the **DNS Server** check box, click **Add Features**, and then click **Next**.
8. On the **Select features** page, click **Next**.
9. On the **DNS Server** page, click **Next**.
10. On the **Confirmation** page, click **Install**.
11. When the role has been added successfully, click **Close**.

▶ Task 2: Create and configure secondary zones

1. On LON-SVR1, on the taskbar, right-click the **Windows PowerShell®** icon, and then click **Run as administrator**.
2. At the command prompt, type the following cmdlets, pressing Enter at the end of each row:

```
add-dnsserversecondaryzone –masterservers 172.16.0.10 –Name Adatum.com –Zonefile
”Adatum.com.dns”
Register-DnsClient
```

3. Switch to LON-DC1. If necessary, sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
4. If necessary, switch to Server Manager.
5. Click **Tools**, and then click **DNS**.
6. In the DNS console, expand **Forward Lookup Zones**, and then click **Adatum.com**.
7. Right-click **Adatum.com**, and then click **Properties**.
8. In the **Adatum.com Properties** dialog box, click the **Zone Transfers** tab.
9. Select the **Allow zone transfers** check box, click **Only to servers listed on the Name Servers tab**, and then click **Notify**.
10. In the **Notify** dialog box, in the **The following servers** list, type **172.16.0.11**, and then press Enter.
11. Click **OK**, and then click the **Name Servers** tab.
12. On the **Name Servers** tab, click **Add**.
13. In the **New Name Server Record** dialog box, in the **Server fully qualified domain name (FQDN)** box, type **LON-SVR1.Adatum.com**, click **Resolve**, and then click **OK** twice.

14. Switch to LON-SVR1.
15. In Server Manager, click **Tools**, and then click **DNS**.
16. In DNS Manager, expand **LON-SVR1**, expand **Forward Lookup Zones**, and then click **Adatum.com**.
17. In the DNS console, right-click **Adatum.com**, and then click **Transfer from Master**.
18. In DNS Manager, press the **F5** key. The zone data should display. If it does not, then right-click **Adatum.com**, and then click **Transfer from Master**.

▶ Task 3: Enable and configure zone transfers

1. On LON-SVR1, in the DNS Manager console tree, right-click **LON-SVR1**, and then click **Properties**.
2. In the **LON-SVR1 Properties** dialog box, click the **Forwarders** tab.
3. On the **Forwarders** tab, click **Edit**.
4. In the **Edit Forwarders** dialog box, in the **IP Address** list, type **172.16.0.10**, and then press Enter.
5. When validation completes, click **OK**.
6. Click the **Advanced** tab. Clear the **Enable round robin** check box, and then click **Apply**.
7. Click the **Root Hints** tab. Click **Remove** repeatedly to empty the **Name servers** list, and then click **OK**.
8. Switch to LON-DC1.
9. On the taskbar, right-click the **Windows PowerShell** icon, and then click **Run as administrator**.
10. At the Windows PowerShell command prompt, type the following command, and then press Enter:

```
set-dnsserverglobalnamezone
-enable $true
```

11. In the DNS console, expand **LON-DC1**, expand **Forward Lookup Zones**, right-click **Forward Lookup Zones**, and then click **New Zone**.
12. Click **Next**.
13. On the **Zone Type** page, click **Primary zone**, and then click **Next**.
14. On the **Active Directory Zone Replication Scope** page, click **To all DNS server running on domain controllers in this forest: Adatum.com**, and then click **Next**.
15. On the **Zone Name** page, in the **Zone name** box, type **GlobalNames**, and then click **Next**.
16. On the **Dynamic Update** page, click **Next**, and then click **Finish**.
17. In the DNS console, click **GlobalNames**, right-click the **GlobalNames zone**, and then click **New Alias**.
18. In the **New Resource Records in the Alias name** text box, type **Server1**, in the **Fully qualified domain name (FQDN) for the target host** (bottom) text box, type **lon-svr1.adatum.com**, and then click **OK**.

▶ Task 4: Test DNS resolution from a client

1. Sign in on LON-CL1 as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. On the Start screen, click any empty area, and then type **PowerShell**.
3. In the **Search** panel, right-click **Windows PowerShell**, and then click **Run as administrato**r.

4. In the Windows PowerShell window, at the command prompt, type the following cmdlet, and then press Enter:

```
Get-DnsClientServerAddress
```

5. Observe the results.

6. In the lower left of the taskbar, right-click the **Windows** icon, and then on the context menu, click **Network Connections**.

7. Right-click the **Ethernet** (Adatum.com) network connection object, and then click **Properties**.

8. In the Ethernet Properties pop-up, in the **This connection uses the following items** box, select the **Internet Protocol Version 4 (TCPIPv4)** checkbox, and then click **Properties**.

9. In the **Internet Protocol Version 4 (TCPIPv4) Properties** dialog box, in the **Alternate DNS server address**, type **172.16.0.11**. Click **OK**, click **Close**, and then close the Network Connections window.

10. Switch back to the Windows PowerShell window.

11. In the Windows PowerShell window, type the following cmdlet, and then press Enter:

```
Get-DnsClientServerAddress
```

Note: You should see the IP address added for LON-SVR1.

12. In the Windows PowerShell window, type the following cmdlet, and then press Enter:

```
Get-DnsClientCache
```

Note: You should see the Lon-dc1 DNS server name and addresses, but not lon-svr1.

13. In the Windows PowerShell window, type the following cmdlet, and then press Enter:

```
Resolve-DnsName LON-SVR1
```

Note: This should return the **A** record entry for LON-SVR1.

14. In the Windows PowerShell window, type the following cmdlet, and then press Enter:

```
Get-DnsClientCache
```

Note: You now should see both the DNS server names and addresses.

▶ Task 5: 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 Microsoft® Hyper-V® Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for 20413C-LON-SVR1 and 20413C-LON-CL1.

Results: After completing this exercise, you will have implemented DNS successfully.

Module 6: Designing and Implementing an Active Directory Domain Services Forest and Domain Infrastructure

Lab A: Designing and Implementing an Active Directory Domain Services Forest Infrastructure

Exercise 1: Designing an Active Directory® Forest Infrastructure

▶ Task 1: Read the supporting documentation

- Read the documentation provided in the Student Handbook.

▶ Task 2: Update the Proposals section with your planned course of action

- Answer the questions in the Proposals section of the A. Datum, Trey Research, and Contoso Forest Integration Strategy document.

Proposals

1. How many forests does the current deployment have?

   The current deployment has two forests: adatum.com, and treyresearch.net. These two forests are not linked.

2. Is that number of forests sufficient?

   It depends on whether you propose to implement an additional forest for Contoso, Ltd, or implement Active Directory Domain Services (AD DS) at Contoso by configuring Contoso as part of the A. Datum forest. It should be clear that the Trey Research forest must remain as a separate forest to provide the degree of administrative and resource separation required by the design objectives. In addition, the schema changes made to the forest may not be suitable in the A. Datum environment, and could be costly to test. For this reason, the number of forests should remain at two.

3. What forest design and forest trust design will enable collaboration between A. Datum, Contoso, and Trey Research? Are there any special requirements for this scenario?

   By implementing Contoso into the existing A. Datum forest, collaboration will be straightforward. Users in different parts of the same forest can share resources easily.

   Regarding integrating Trey Research, an easier option to enable collaboration would be to create a forest trust between A. Datum and Trey Research. However, to separate administration of the two organizations, it would be better to implement nontransitive trusts.

   Alternatives to the trusts between A. Datum and Trey Research also exist, such as implementing Active Directory Federation Services (AD FS).

4. How can you address the requirement to protect confidential data in Trey Research from unauthorized access?

   If you implement Trey Research as a separate forest, they can continue to control access to their sensitive research data in the same way they have always done. The design does not compromise this significant design goal.

5. How should you plan for A. Datum perimeter server requirements?

   You can use one of two methods: you should deploy the perimeter services as part of a stand-alone Active Directory forest on the perimeter network, or you should deploy Active Directory Lightweight Directory Services (AD LDS) to the perimeter, and then configure it to interact with the internal A. Datum forest.

6. Are there any alternatives to the forest design that you would consider?

   Answers will vary, but you might have considered implementing three forests—one for each organization. Alternatively, you may have selected a single forest for the entire merged organization.

7. What are the benefits and drawbacks of the alternative design, if any?

   A three-forest design would be more complicated to configure and manage than is required by the objectives. The best design would be to merge Contoso into the existing forest. Because this design is easier to configure and administer, it will likely be less expensive, and it fulfills the other design objectives of collaboration and simplicity.

   Merging Trey Research would be more complicated, and therefore more expensive than the other methods. It also compromises the critical security design consideration. Therefore, Trey Research requires administrative and resource separation.

8. Draw a simple diagram of your proposed Active Directory forest design in the space provided.

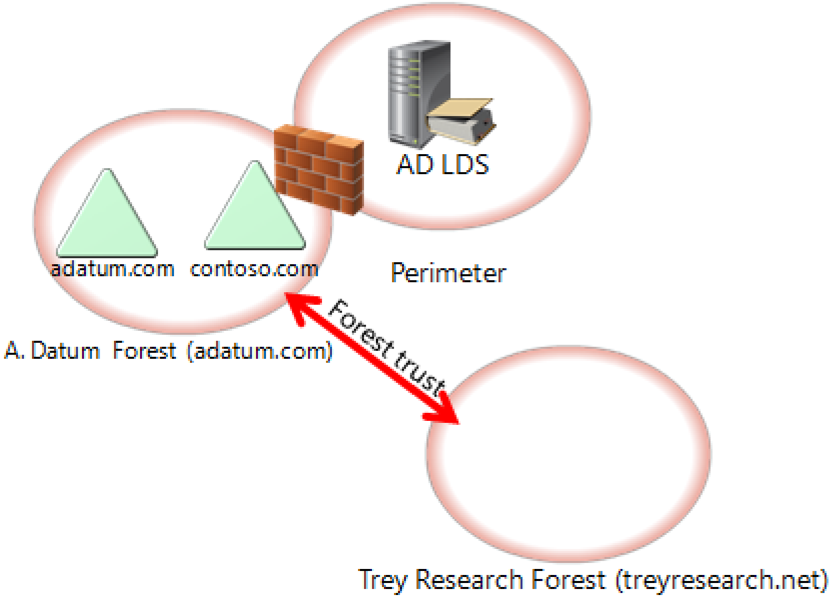
AD LDS
adatum.com
contoso.com
Perimeter
Forest trust
A. Datum Forest (adatum.com)
Trey Research Forest (treyresearch.net)

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones shown previously.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you should have created a forest design that incorporates A. Datum Corporation, Trey Research, and Contoso, Ltd.

Exercise 2: Implementing Active Directory Forest Trusts

▶ Task 1: Configure DNS to support the forest trusts

1. Switch to LON-DC1 and, if necessary, sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. In Server Manager, click **Tools**, and then click **DNS**.
3. In Domain Name System (DNS), in the navigation pane, expand **LON-DC1**, expand **Conditional Forwarders**, right-click **Conditional Forwarders**, and then click **New Conditional Forwarder**.
4. In the **New Conditional Forwarder** dialog box, in the **DNS Domain** text box, type **treyresearch.net**. In the IP addresses of the master servers section, under **IP Address**, type **172.16.10.10**, press Enter, and then click **OK**.
5. Click the **Start** button.
6. Type **cmd.exe**, and then press Enter.
7. At the command prompt, type the following command, and then press Enter:

```
nslookup trey-dc1.treyresearch.net
```

8. Verify that the query is successful, returning the IP address of 172.16.10.10.
9. Switch to TREY-DC1.
10. If necessary, sign in as **Treyresearch\Administrator** with the password **Pa$$w0rd**.
11. Click the **Start** button, point to **Administrative Tools**, and then click **DNS**.
12. In the DNS management console, in the navigation pane, expand **TREY-DC1**, expand **Conditional Forwarders**, right-click **Conditional Forwarders**, and then click **New Conditional Forwarder**.
13. In the **New Conditional Forwarder** dialog box, in the **DNS Domain** text box, type **Adatum.com**, in the IP addresses of the master servers section, under **IP Address**, type **172.16.0.10**, press Enter, and then click **OK**.
14. Click the **Start** button, type **cmd.exe**, and then press Enter.
15. At the command prompt, type the following command, and then press Enter:

```
nslookup lon-svr1.adatum.com
```

16. Verify that the query is successful, returning the IP address of 172.16.0.11.

▶ Task 2: Create the required forest trusts between A. Datum and Trey Research

1. Switch to LON-DC1.
2. In Server Manager, click **Tools**, and then click **Active Directory Domains and Trusts**.
3. In Active Directory Domains and Trusts, click **Adatum.com**, right-click **Adatum.com**, and then click **Properties**.
4. In the **Adatum.com Properties** dialog box, click the **Trusts** tab.
5. On the **Trusts** tab, click **New Trust**.
6. In the **New Trust Wizard** dialog box, click **Next**.
7. On the **Trust Name** page, in the **Name** text box, type **treyresearch.net**, and then click **Next**.
8. On the **Trust Type** page, click **Forest trust**, and then click **Next**.

9. On the **Direction of Trust** page, ensure that the **Two-way** option is selected, and then click **Next**.
10. On the **Sides of Trust** page, click **Both this domain and the specified domain**, and then click **Next**.
11. On the **User Name and Password** page, in the **User name** text box, type **Treyresearch\administrator**.
12. In the **Password** text box, type **Pa$$w0rd**, and then click **Next**.
13. On the **Outgoing Trust Authentication Level--Local Forest** page, click **Next**.
14. On the **Outgoing Trust Authentication Level--Specified Forest** page, click **Next**.
15. On the **Trust Selections Complete** page, click **Next**.
16. On the **Trust Creation Complete** page, click **Next**.
17. On the **Confirm Outgoing Trust** page, click **Yes, confirm the outgoing trust**, and then click **Next**.
18. On the **Confirm Incoming Trust** page, click **Yes, confirm the incoming trust**, and then click **Next**.
19. On the **Completing the New Trust Wizard** page, click **Finish**.
20. In the **Adatum.com Properties** dialog box, click **OK**.

Results: After completing this exercise, you should have successfully implemented part of the forest infrastructure strategy that you designed.

▶ Task: To prepare for the next lab

- Leave all the virtual machines running.

Lab B: Designing and Implementing an Active Directory Domain Infrastructure

Exercise 1: Designing an Active Directory Domain Infrastructure

▶ Task 1: Read the supporting documentation

- Read the documentation provided in the Student Handbook.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the Proposals section of the Contoso AD DS Integration Strategy document.

Proposals

1. Should you create a separate forest to accommodate the Contoso organization?

   No. Where possible, a single forest is always preferable because it is simpler to manage, and provides the best integration. Circumstances do exist where multiple forests are useful, but these circumstances do not apply to this scenario. Therefore, a design that consolidates Contoso into the adatum.com forest is the best approach.

2. If you decide to implement Contoso as part of the existing A. Datum forest, is it better to implement Contoso as a separate domain or as an organizational unit (OU) in the existing A. Datum domain?

   It is better to implement Contoso as a separate domain. Using an OU limits the options available for administrative separation and delegation.

3. Assuming you choose a separate domain for Contoso, what Active Directory domain name will you use for Contoso? (Contoso already uses contoso.com for DNS purposes.)

   Contoso maintains an external domain name of contoso.com. (This is not an Active Directory domain name, but a DNS domain name.)
   Currently, the internal domain name matches the external namespace, and split-DNS has been configured to accommodate this configuration. However, as a best practice, you should use an internal namespace that does not match the external namespace. Therefore, two strategies exist: implement a different domain name (such as contoso.priv), or implement AD DS as a subdomain of the external namespace (for example, ad.contoso.com).
   However, changing the domain name from contoso.com to something else for all internal resources may be a time-consuming process, and hence, costly. Given that adatum.com and treyresearch.net use the same domain name internally and externally, the name contoso.com is adequate.
   There is no requirement for the Contoso namespace to be subordinate to that of the adatum.com namespace. They can each be configured as separate Active Directory trees.

4. Which is the forest root domain?

   The forest root domain is always the first domain created in a forest. In this instance, the first domain was adatum.com, and therefore, this is the forest root. As the forest root, adatum.com cannot be changed.

5. Is it a good idea to deploy additional domain controllers from the adatum.com domain in Paris? Why or why not?

   Yes, it is a good idea. It is useful to ensure that domain controllers for the forest root are accessible from all domains that have the greatest number of users. For example, when processing Group Policy Objects (GPOs) based on site containers, the policies process from domain controllers in the forest root domain.

6. How do you plan to address the requirement that users in Contoso need to access resources in the Trey Research organization?

   A forest trust is already established between treyresearch.net and adatum.com. This may suffice. However, if performance is a concern, you can improve access times by creating shortcut trusts between the treyresearch.net and contoso.com domains (depending upon what domain name was chosen for Contoso).

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones shown previously.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you should have successfully designed a domain infrastructure strategy for the integration of Contoso into the A. Datum organization.

Exercise 2: Implementing an Active Directory Domain Infrastructure

▶ Task 1: Verify that the prerequisites for adding a new domain are satisfied

1. Switch to CON-SVR.
2. Sign in as **Administrator** with a password of **Pa$$w0rd**.
3. In Server Manager, in the details pane, click **Add roles and features**.
4. In the Add Roles and Features Wizard, on the **Before you begin** page, click **Next**.
5. On the **Select installation type** page, click **Next**.
6. On the **Select destination server** page, click **Next**.
7. On the **Select server roles** page, in the **Roles** list, select the **Active Directory Domain Services** check box.
8. Click **Add Features**, and then click **Next**.
9. On the **Select features** page, click **Next**.
10. On the **Active Directory Domain Services** page, click **Next**.
11. On the **Confirm installation selections** page, click **Install**.
12. When the role installation completes, click **Close**.

▶ Task 2: Add CON-SVR as a domain controller in a new domain in an existing forest

1. In Server Manager, in the navigation pane, click **AD DS**.
2. In the details pane, click **More**.
3. In the **All Servers Task Details and Notifications** dialog box, click **Promote this server to a domain controller**.
4. In the Active Directory Domain Services Configuration Wizard, on the **Deployment Configuration** page, click **Add a new domain to an existing forest**.
5. In the **Select domain type** list, click **Tree Domain**.
6. In the **Forest name** text box, type **adatum.com**.
7. In the **New domain name** text box, type **contoso.com**, and then click **Change**.
8. In the **Windows Security** dialog box, in the **User name** text box, type **Adatum\Administrator**. In the **Password** text box, type **Pa$$w0rd**.
9. Click **OK**, and then click **Next**.
10. On the **Domain Controller Options** page, in the **Password** and **Confirm password** text boxes, type **Pa$$w0rd**, and then click **Next**.
11. On the **DNS Options** page, click **Next**.
12. On the **Additional Options** page, click **Next**.
13. On the **Paths** page, click **Next**.

14. On the **Review Options** page, click **Next**.
15. When the prerequisites check completes, click **Install**.
16. After CON-SVR restarts, when prompted, sign in as **Contoso\Administrator** with the password **Pa$$w0rd**.

Results: After completing this exercise, you should have successfully implemented a part of the domain infrastructure strategy that you devised.

▶ Task: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 through 4 for **20413C-TREY-DC1** and **20413C-CON-SVR**.

Module 7: Designing and Implementing an AD DS Organizational Unit Infrastructure

Lab: Designing and Implementing an Active Directory OU Infrastructure and Delegation Model

Exercise 1: Designing an Organizational Unit Infrastructure

▶ Task 1: Read the supporting documentation

- Read the documentation provided in the Student Handbook.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposal section of the A. Datum OU Redesign Proposal document.

Proposals

1. Which organizational unit (OU) design model do you think would work in this situation?

Multiple possible models could work. However, a hybrid model with a multiple tenancy–based approach with additional OUs at the top level is the best solution. A. Datum wants to provide central services in the future for Contoso, Ltd and Trey Research. In addition, A. Datum wants to separate their infrastructure administration from the data (User/Group) administration.

2. How does centralized administration in London affect the overall OU design?

In the A. Datum part of the OU model, an OU for Users and Groups should be at the next lower level. There is also a requirement to separate departments. However, because a single team should have rights to create and delete all Users and Groups, the preferred solution is to separate these object types at a top level prior to splitting into departments (or group types, which is likely but not requested in the current proposal).

3. How can you place Active Directory® objects within the OU structure?

The easiest way to place Active Directory objects within the OU structure is by using scripting. From a command prompt, use **dsquery** piped to **dsmove**. In a Windows PowerShell® command-line interface, use the cmdlets **get-ADObject** and **move-ADObject**. You can query for the **objecttype**, and, for example, for the department, and then move those objects to their new respective OUs.

You can also use the Active Directory Administrative Center to place objects. When evaluating the tool to use for an administrative task, you should consider the number of objects that you are working with, whether you want to reuse the solution, and whether you are familiar with the command line and scripting methods.

4. Which OUs do you suggest for the top- level OUs?

There are multiple valid options. The following is one approach:

- Central-IT (hosts central admin accounts and delegation groups)
- Enterprise (enterprise-wide services, in this case, only Servers)
- Adatum (all regular users, groups, clients, and servers of A. Datum)
- TreyResearch (all regular users, groups, clients, and servers of Trey Research)
- Contoso (all regular users, groups, clients, and servers of Contoso)

5. Which OUs should you create to delegate administration?

- Central-IT: no delegation at all, changes are only to be made by domain admins.
- Enterprise\Servers\SQL, WEB, APP: the appropriate server or application admin team obtains delegated rights or local rights (Server Operators) via Group Policy Object (GPO).
- Adatum\Users and Adatum\Groups: for the Adatum Team in London that is creating and deleting users or groups.
- Adatum\Users\Marketing (and other departments): for the departmental administration of the users' passwords resets.
- Adatum\Clients\Sydney (and London, Toronto, and later possibly some of the smaller branches): for the local admin team that administers clients, they will be granted local admins' rights via GPO to support user issues on the client computers.
- Adatum\Servers\Sydney (London, Toronto): for the local admin team that is administering servers, they will be granted server operator rights via GPO on the servers.
- Trey Research and Contoso can copy the Adatum model, or incorporate their own specific requirements.

6. Which OUs should you create so that you can manage the local client computers and servers?

- Adatum\Clients\Sydney
- Adatum\Servers\Sydney

7. Which OUs should you create for managing the enterprise application servers?

- Enterprise\Servers\SQL (WEB, APP)

8. Create a drawing of your draft design.

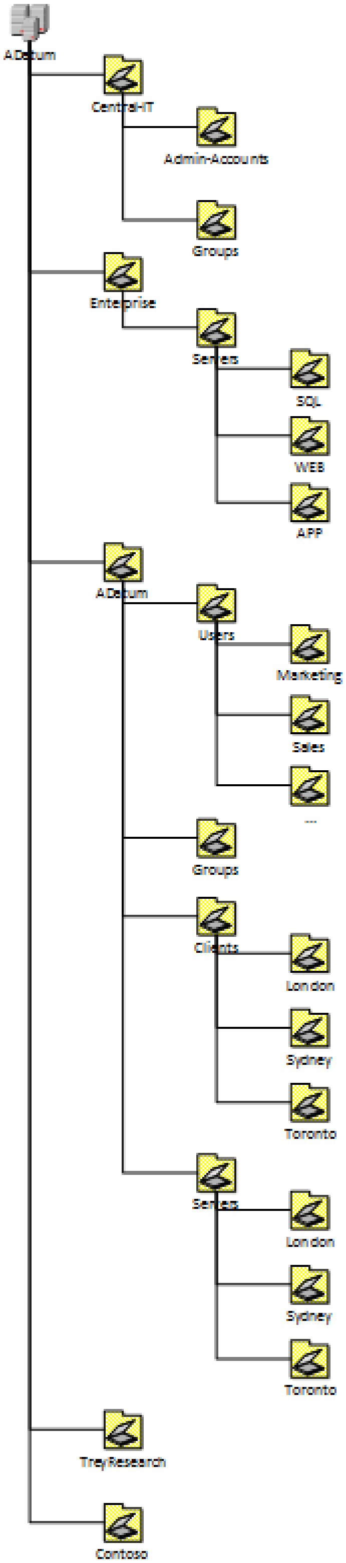
ADatum
Central-IT
Admin-Accounts
Groups
Enterprise
Servers
SQL
WEB
APP
ADatum
Users
Marketing
Sales
...
Groups
Clients
London
Sydney
Toronto
Servers
London
Sydney
Toronto
TreyResearch
Contoso

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones shown previously.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you should have created an OU design that reflects the A. Datum Corporation's administrative task model.

Exercise 2: Implementing the OU Design

▶ Task 1: Create the new OU structure

1. Switch to LON-DC1 and, if necessary, sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.

Note: For the following steps, you can also use the Windows PowerShell cmdlet **New-ADOrganizationalUnit**. The syntax is:

```
New-ADOrganizationalUnit –name OU-Name –path “parentDN”
```

2. In Server Manager, click **Tools**, and then click **Active Directory Administrative Center**.
3. In Active Directory Administrative Center, in the navigation pane, click **Adatum (local)**.
4. In the tasks pane, in the Adatum (local) section, click **New**, and then click **Organizational Unit**.
5. In the **Create Organizational Unit** dialog box, in the **Name** text box, type **Central-IT**.
6. In the **Description** text box, type **Admin accounts and Groups for Delegation. Administered from DOMAIN ADMINS ONLY**, and then click **OK**.
7. Repeat Steps 4 through 6 to create the following top-level OUs:

- Name: **Enterprise**
  - Description: **Enterprise-wide managed objects**
- Name: **Adatum**
  - Description: **Regular objects for A. Datum**
- Name: **Contoso**
  - Description: **Regular objects for Contoso**
- Name: **TreyResearch**
  - Description: **Regular objects for Trey Research**

8. Use the provided Windows PowerShell Script to create the sub-OUs:
   a. On the taskbar, right-click the **Windows PowerShell** icon, and then click **Run ISE as Administrator**.
   b. In the **Administrator: Windows PowerShell ISE** dialog box, click **File**, and then click **Open**.
   c. In the **Open** dialog box, in the **File name** text box, type **E:\Labfiles\Create-SubOUs.ps1**, and then click **Open**.
   d. Click **File**, and then click **Run**.
   e. Leave Windows PowerShell open for the next part of the lab.

Use the provided Windows PowerShell script to create the sub-OUs.

Create-SubOUs.ps1

```
$subous = @(`
("Admin-Accounts","ou=Central-IT,dc=adatum,dc=com","Admin-accounts only"),`
("Groups","ou=Central-IT,dc=adatum,dc=com","Groups for administrative tasks
delegation"),`
("Servers","ou=Enterprise,dc=adatum,dc=com","Enterprise-wide managed Servers"),`
("SQL","ou=Servers,ou=Enterprise,dc=adatum,dc=com",""),`
("WEB","ou=Servers,ou=Enterprise,dc=adatum,dc=com",""),`
("APP","ou=Servers,ou=Enterprise,dc=adatum,dc=com",""),`
("Users","ou=Adatum,dc=adatum,dc=com","Adatum regular user accounts"),`
("Groups","ou=Adatum,dc=adatum,dc=com","Adatum regular group accounts"),`
("Clients","ou=Adatum,dc=adatum,dc=com","Adatum clients"),`
("Servers","ou=Adatum,dc=adatum,dc=com","Adatum Servers"),`
("Marketing","ou=Users,ou=Adatum,dc=adatum,dc=com",""),`
("Sales","ou=Users,ou=Adatum,dc=adatum,dc=com",""),`
("Development","ou=Users,ou=Adatum,dc=adatum,dc=com",""),`
("IT","ou=Users,ou=Adatum,dc=adatum,dc=com",""),`
("Research","ou=Users,ou=Adatum,dc=adatum,dc=com",""),`
("London","ou=Clients,ou=Adatum,dc=adatum,dc=com",""),`
("Sydney","ou=Clients,ou=Adatum,dc=adatum,dc=com",""),`
("Toronto","ou=Clients,ou=Adatum,dc=adatum,dc=com",""),`
("London","ou=Servers,ou=Adatum,dc=adatum,dc=com",""),`
("Sydney","ou=Servers,ou=Adatum,dc=adatum,dc=com",""),`
("Toronto","ou=Servers,ou=Adatum,dc=adatum,dc=com",""))

$subous | %{New-ADOrganizationalUnit -Name $_[0] -Path $_[1] -Description $_[2]}
```

9. Switch back to Active Directory Administrative Center, and verify the presence of the OUs. The OU structure should match the design that you created in the last exercise.

▶ Task 2: Migrate the user and group accounts to the new OUs

1. On LON-DC1, switch to the Windows PowerShell Integrated Scripting Environment (ISE).

2. Click **File**, and then click **Close**.

3. Click **File**, and then click **Open**.

4. In the **Open** dialog box, in the **File name** text box, type **E:\Labfiles\Move-AdatumUserGroups.ps1**, and then click **Open**.

Use the provided Windows PowerShell script to move users and groups into their new OUs.

Move-AdatumUserGroups.ps1

```
$depts = @("Marketing","Sales","IT","Development","Research")
ForEach ($dept in $depts) {
$userDN = "OU=" + $dept + ",ou=users,ou=adatum,dc=adatum,dc=com"
$users = Get-ADUser -Properties Department -Filter {department -eq $dept}
    ForEach ($user in $users) {
        Move-ADObject $user -TargetPath $userDN
    }
$group = Get-ADGroup -Filter {Name -eq $dept}
Move-ADObject $group -TargetPath "ou=groups,ou=adatum,dc=adatum,dc=com"
}
```

5. Click **File**, and then click **Run**.

6. Click **File**, and then click **Close**.

7. Switch back to Active Directory Administrative Center, and verify that the user and group objects have been moved:

    a. In the navigation pane, click **Adatum (local)**.

    b. In the details pane, double-click **Adatum**, double-click **Users**, and then double-click **Sales**.

Note: The Sales OU should be populated with user objects. If it is not, ask your instructor for guidance with the script. If you want, you can open each OU and check for users.

c. In the navigation pane, click **Adatum (local)**.

d. In the details pane, double-click **Adatum**, and then double-click **Groups**.

Note: The Groups OU should be populated with group objects. If it is not, ask your instructor for guidance with the script.

Results: After completing this exercise, you should have successfully implemented a part of the OU design.

Exercise 3: Designing and Implementing an Active Directory Permissions Model

▶ Task 1: Read the supporting documentation

- Read the documentation provided in the Student Handbook.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the A. Datum OU Administrative Tasks Delegation Model document.

Proposals

1. Which groups must you create to facilitate the high-level requirements?

   You must configure the following groups:

Name	Type	Members	Description
London-UserGroupProvisioning	Global	To be determined	Create and delete User and Group accounts in Adatum
acl_adatum-users_ccdc	Domain Local	London-UserGroupProvisioning	
acl_adatum-groups_ccdc	Domain Local	London-UserGroupProvisioning	
Enterprise-ServerOps	Global	To be determined	Manage enterprise application servers
acl_enterprise_serveroperator	Domain Local	Enterprise-ServerOps	
Marketing-Admins	Global	To be determined	Manage the Marketing users
acl_Users-Marketing_resetpwd	Domain Local	Marketing-Admins	
Xyz-Department-Admins	Global	To be determined	Manage the xyz-Department users
acl_xyzDepartment_resetpwd	Domain Local	Xyz-Department-Admins	
London-Admins	Global	To be determined	Manage London clients and servers

Name	Type	Members	Description
acl_clients-London_computer_ccdc	Domain Local	London-Admins	
acl_servers-London_computer_ccdc	Domain Local	London-Admins	

2. What delegations must you configure?

   You must configure the following delegations:

Group	Where	Permissions	Applies to
acl_adatum-users_ccdc	Adatum\Users	**Create User objects** and **Delete User objects**	All descendant objects
acl_adatum-groups_ccdc	Adatum\Group	**Create Group objects** and **Delete Group objects**	All descendant objects
acl_enterprise-serveroperator		GPO for granting server operators rights only	
acl_Users-Marketing_resetpwd	Adatum\Users\Marketing	Reset passwords	Descendant User objects
acl_clients-London_computer_ccdc	Adatum\Clients\London	**Create Computer objects** and **Delete Computer objects**, additional GPO to grant local admin rights	
acl_servers-London_computer_ccdc	Adatum\Servers\London	**Create Computer objects** and **Delete Computer objects**, additional GPO to grant server	

Group	Where	Permissions	Applies to
		operators rights	

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones shown previously.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

▶ Task 5: Create the required groups

1. On LON-DC1, switch to Active Directory Administrative Center.
2. In Active Directory Administrative Center, in the navigation pane, click **Adatum (local)**.
3. In the details pane, double-click the **Central-IT** OU.
4. In the details pane, click the **Groups** OU.
5. In the tasks pane, in the Groups section, click **New**, and then click **Group**.
6. In the **Create Group** dialog box, in the **Group name** text box, type **London-UserGroupprovisioning**, and then click **OK**.
7. Repeat steps 5 and 6 to create these additional groups:

- Name: **Enterprise-ServerOps**
  - Type: **Global/Security**
- Name: **Marketing-Admins**
  - Type: **Global/Security**
- Name: **London-Admins**
  - Type: **Global/Security**

8. In the tasks pane, in the Groups section, click **New**, and then click **Group**.
9. In the **Create Group** dialog box, in the **Group name** text box, type **acl_adatum-users_ccdc**.
10. Under Group scope, click **Domain Local**.
11. Scroll down to the Members section, and then click **Add**.
12. In the **Select Users, Contacts, Computers, Service accounts, or Groups** dialog box, in the **Enter the object name to select (examples)** text box, type **London-UserGroupProvisioning**, click **Check Names**, and then click **OK**.
13. In the **Create Group: acl_adatum-users_ccdc** dialog box, click **OK**.
14. Repeat steps 8 through 13 to create the following groups:

- Name: **acl_adatum-groups_ccdc**
  - Type: **Domain Local/Security**
  - Member: **London-UserGroupProvisioning**
- Name: **acl_enterprise-serveroperator**
  - Type: **Domain Local/Security**

  - Member: **Enterprise-ServerOps**
- Name: **acl_Users-Marketing_resetpwd**
  - Type: **Domain Local/Security**
  - Member: **Marketing-Admins**
- Name: **acl_clients-London_computer_ccdc**
  - Type: **Domain Local/Security**
  - Member: **London-Admins**
- Name: **acl_servers-London_computer_ccdc**
  - Type: **Domain Local/Security**
  - Member: **London-Admins**

▶ Task 6: Delegate permissions to management groups

Note: You can optionally complete this task by using the command-line tool Dsacls.exe. For example, to grant rights to create or delete group objects, the syntax is as follows:

```
dsacls ou=… /G adatum\acl_...:CCDC;group
```

For help with the tool Dsacls.exe, at a command prompt type **dsacls /?**.

1. Switch to Active Directory Administrative Center.
2. In Active Directory Administrative Center, in the navigation pane, switch to **Tree View**.
3. Expand **Adatum (local)**, and then click **Users**.
4. In the task pane, in the Users section, click **Properties**.
5. In the **Users** properties dialog box, scroll down to the Extensions section.
6. On the **Security** tab, click **Advanced**.
7. In the **Advanced Security Settings for Users** dialog box, on the **Permissions** tab, click **Add**.
8. In the **Permission Entry for Users** dialog box, click **Select a principal**.
9. In the **Select User, Computer, Service Account or Group** dialog box, in the **Enter the object name to select (examples)** text box, type **acl_adatum-users_ccdc**.
10. Click **Check Names**, and then click **OK**.
11. In the **Permission Entry for Users** dialog box, scroll down, and then click **Clear All**.
12. In the Permissions section, select both the **Create User objects** and **Delete User objects** check boxes, and then click **OK**.
13. In the **Advanced Security for Users** dialog box, click **OK**.
14. In the **Users** dialog box, click **Cancel**.

Note: By using the **dsacls** command, you can perform the same task of granting permissions with the following command:

```
dsacls ou=users,ou=adatum,dc=adatum,dc=com /G adatum\acl_adatum-users_ccdc:CCDC;group
```

15. Repeat steps 3 through 13 to configure the following settings:

- OU: **Adatum\Groups**
  Group: **acl_adatum-groups_ccdc**
  Permissions: **Create Group objects, Delete Group objects**
- OU: **Adatum\Users\Marketing**
  Group: **acl_users-Marketing_resetPwd**
  Applies to: **Descendant User objects**
  Permissions: **Reset password**
- OU: **Adatum\Clients\London**
  Group: **acl_clients-London_computer_ccdc**
  Permissions: **Create Computer objects, Delete Computer objects**
- OU: Adatum\Servers\London
  Group: acl_servers-London_computer_ccdc
  Permissions: Create Computer objects, Delete Computer Objects

▶ 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 Microsoft® Hyper-V® Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.

Results: After completing this exercise, you should have successfully designed and implemented an administrative permissions model.

Module 8: Designing and Implementing a Group Policy Object Strategy

Lab: Designing and Implementing a Group Policy Object Strategy

Exercise 1: Designing a Group Policy Object (GPO) Strategy

Task 1: Read the supporting documentation

- Read the documentation provided.

Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the A. Datum GPO Strategy Proposal document.

Proposals

1. Which of the requirements will necessitate creating one or more GPOs?

   The central IT administrators in London must be able to manage all GPOs and settings in the organization. Administrators in each office should be able to manage only GPOs that apply to that office. Although you can complete any of the remaining tasks manually on each computer, using GPOs requires the least effort. You could implement some of the other requirements, such as the security warning or preventing access to registry editing tools, by using local policies only. However, because local policies are hard to manage, GPOs are also beneficial for these settings.

2. Can you fulfill any of the requirements without creating GPOs?

   You can fulfill all of the requirements without creating GPOs.

3. Are there any exceptions to the default GPO application that you must consider?

   Yes, there is one exception: security filtering of administrator desktops so that they will not be prevented from accessing registry editing tools.

4. List the GPOs that you must create to fulfill the lab scenario's requirements. Provide the following information in the table provided:

   - Name of the GPO
   - The requirements that the GPO fulfills
   - The configuration settings (user policies, computer policies, user preferences, or computer preferences) the GPO will contain
   - The container (domain, organizational unit (OU), site) to which the GPO will be linked

Name	Requirements fulfilled	Configuration settings	Applies to
All_Clients	Configures the local admin accounts	Computer Configuration \ Policies \ Windows Settings \ Security Settings \ Restricted Groups	OU=Clients

Name	Requirements fulfilled	Configuration settings	Applies to
All_Clients	Configures general Windows® Update settings	Computer Configuration \ Policies \ Administrative Templates \ Windows Components \ Windows Update \ Configure Automatic Updates	OU=Clients
All_Users_but_Admins	Prevents editing of the registry	User Configuration \ Policies \ Administrative Templates \ System \ Prevent access to registry editing tools	DC=adatum
London_ Clients	Displays a compliance message	Computer Configuration \ Windows Settings \ Security Settings \ Local Policies \ Security Options \ Interactive Logon: Message text for users attempting to log on Interactive Logon: Message title for users attempting to log on	OU=London, OU=Clients
Marketing_Share	Users must have a default set of mapped drives	User Configuration \ Preferences \ Windows Settings \ Drive Maps	OU=Marketing

5. List other configuration tasks that you must perform within the Group Policy Management Console to fulfill the scenario requirements.

   Other configuration tasks include:

   - The All_Users_but_Admins policy needs security filtering to deny access. This will apply the policy to the users but not to administrators (Group IT).
   - You must configure the administration of GPOs as desired.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones shown previously.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you should have created a GPO design that meets the GPO requirements of A. Datum Corporation.

Exercise 2: Implementing the GPO Design

▶ Task 1: Prepare the environment

1. Switch to LON-DC1, and if necessary, sign is as **Adatum\Administrator** with the password **Pa$$w0rd**.

2. On the taskbar, right-click the **Windows PowerShell®** icon, and then click **Run ISE as Administrator**.

3. At the Administrator: Windows PowerShell ISE prompt, type the following Windows PowerShell script, pressing Enter at the end of each line:

```
New-ADOrganizationalUnit –name Clients –path “dc=adatum,dc=com”
New-ADOrganizationalUnit –name London
–path “ou=clients,dc=adatum,dc=com”
Get-ADObject -Filter {name –eq ‘LON-CL1’} | Move-ADObject -TargetPath
“ou=London,ou=Clients,dc=adatum,dc=com”
```

4. Click **File**, and then click **Run**.

5. Click **File**, and then click **Close**. When prompted, click **No**.

6. Click **File**, and then click **New**.

7. In Untitled1.ps1, type the following Windows PowerShell script, pressing Enter at the end of each line:

```
New-Item c:\shares –ItemType Directory
New-Item c:\shares\Marketing –ItemType Directory
New-SmbShare –Name Marketing –Path c:\shares\Marketing –FullAccess ADatum\Marketing
```

8. Click **File**, and then click **Run**.

▶ Task 2: Create the required GPOs and link them to the required domain containers

1. On LON-DC1, switch to Server Manager.

2. In Server Manager, click **Tools**, and then click **Group Policy Management**.

3. In the Group Policy Management Console, in the navigation pane, expand **Forest Adatum.com**, expand **Domains**, expand **Adatum.com**, and then click **Group Policy Objects**.

4. On the **Action** menu, click **New**.

5. In the **New GPO** dialog box, in the **Name** text box, type **All_Clients**, and then click **OK**.

6. Right-click the **All_Clients** policy, and then click **Edit**.

7. In the Group Policy Management Editor, in the navigation pane, expand **Computer Configuration**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, and then click **Restricted Groups**.

8. In the details pane, right-click in the empty area, and then click **Add Group**.

9. In the **Add Group** dialog box, in the **Group** text box, type **Administrators**, and then click **OK**.

10. In the **Administrators Properties** dialog box, next to Members of this group, click **Add**.

11. In the **Add Member** dialog box, click the **Browse** button.

12. In the **Select Users, Service Accounts, or Groups** dialog box, in the **Enter the object names to select** text box, type **IT**, and then click **Check Names**.

13. Click **OK** three times to close the dialog boxes.
14. In the Group Policy Management Editor, under **Computer Configuration**, expand **Policies**, expand **Administrative Templates**, expand **Windows Components**, and then click **Windows Update**.
15. Double-click the **Configure Automatic Updates** Group Policy setting.
16. Click **Enable**, in the **Configure automatic updating** drop-down list box, click **4 – Auto download and schedule the install**, and then click **OK** to close the Group Policy setting.
17. Close the Group Policy Management Editor.
18. In the Group Policy Management Console, in the navigation pane, click the **Clients** OU.
19. Right-click the **Clients** OU, and then click **Link an Existing GPO**.
20. In the **Select GPO** dialog box, click the **All_Clients** GPO, and then click **OK**.
21. Right-click **Group Policy Objects**, and then click **New**.
22. In the **New GPO** dialog box, in the **Name** text box, type **All_Users_but_Admins**, and then click **OK**.
23. In the navigation pane, right-click the **All_Users_but_Admins** policy, and then click **Edit**.
24. In the Group Policy Management Editor, in the navigation pane, expand **User Configuration**, expand **Policies**, expand **Administrative Templates**, click **System** node, and then double-click the **Prevent access to registry editing tools** policy.
25. Click **Enable**, and then click **OK** to close the dialog box.
26. Close the Group Policy Management Editor.
27. In the Group Policy Management Console, right-click the **Adatum** domain, and then click **Link an Existing GPO**.
28. In the **Select GPO** dialog box, click the **All_Users_but_Admins** GPO, and then click **OK**.
29. Right-click **Group Policy Objects**, and then click **New**.
30. In the **New GPO** dialog box, in the **Name** text box, type **London_Clients**, and then click **OK**.
31. Right-click the **London_Clients** policy, and then click **Edit**.
32. In the Group Policy Management Editor, expand **Computer Configuration**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, expand **Local Policies**, and then click **Security Options**.
33. In the details pane, double-click the **Interactive Logon: Message text for users attempting to log on** policy setting.
34. Select the **Define this policy setting in the template** check box, and then type the message **Only A. Datum Employees are allowed to log on to this computer**.
35. Click **OK** to close the policy setting.
36. In the Group Policy Management Editor, in the details pane, double-click the **Interactive Logon: Message title for users attempting to log on** policy setting.
37. Select the **Define this policy setting** check box, and then type the message title as **Property of A. Datum**.
38. Click **OK** to close the policy setting.
39. Close the Group Policy Management Editor.
40. In the Group Policy Management Console, under the Clients OU, right-click the **London** OU, and then click **Link an Existing GPO**.

41. In the **Select GPO** dialog box, click the **London_Clients** GPO, and then click **OK**.
42. Right-click **Group Policy Objects**, and then click **New**.
43. In the **New GPO** dialog box, in the **Name** text box, type **Marketing_Share**, and then click **OK**.
44. Right-click the **Marketing_Share** policy, and then click **Edit**.
45. In the Group Policy Management Editor, expand **User Configuration**, expand **Preferences**, expand **Windows Settings**, and then click **Drive Maps**.
46. In the details pane, right-click in the empty area, click **New**, and then click **Mapped Drive**.
47. Under New Drive Properties, in the **Location** text box, type **\\LON-DC1\Marketing**.
48. In the **Label as** text box, type **Marketing-Materials**.
49. In the **Drive Letter** drop-down list box, click drive **M**, and then click **OK**.
50. Close the Group Policy Management Editor.
51. In the Group Policy Management Console, right-click the **Marketing** OU, and then click **Link an Existing GPO** on the context menu.
52. In the **Select GPO** dialog box, click the **Marketing_Share** GPO, and then click **OK**.

▶ Task 3: Configure filtering

1. In the Group Policy Management Console, in the navigation pane, expand **Group Policy Objects**, and then click the **All_Users_but_Admins** Group Policy.
2. In the details pane, click the **Delegation** tab, and then click the **Advanced** button.
3. In the **All_Users_but_Admins Security Settings** dialog box, click the **Add** button.
4. In the **Select Users, Computers, Service Accounts, or Groups** dialog box, in the **Enter the object names to select** text box, type **IT**, and then click **Check Names**.
5. Click **OK** to close the dialog box.
6. In the **All_Users_but_Admins** Security Settings, in the **Group or user names** box, ensure that **IT** is selected.
7. In the **Permissions for IT** box, for the **Apply group policy** setting, select the **Deny** check box, and then click **OK**.
8. When the Windows Security window displays, asking if you want to continue, click **Yes**.

▶ Task 4: Test the design

1. Switch to LON-CL1, and sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. On the **Start** screen, type **cmd**, right-click **Command Prompt**, and then click **Run as Administrator**.
3. In the **User Account Control** dialog box, click **Yes**.
4. In the Administrator: Command Prompt window, type the following command, and then press Enter:

```
gpupdate /force.
```

5. Close the Command Prompt window.
6. Right-click the **Start** button, expand the **Shut down or sign out** submenu, and then click **Restart**.
7. After LON-CL1 restarts, click **OK**, and then sign in as **Adatum\Adam** with the password **Pa$$w0rd**.

Note: Adam Barr is a member of the Marketing Group. Also, note that prior to signing in, Adam is receiving a compliance message.

8. After signing in, on the desktop, on the charms menu, click **Settings**, and then click **Control Panel**.
9. In Control Panel, click **System and Security**.
10. In System and Security, click **Windows Update**.
11. On the left side, click **Change settings**.
12. In the **Change Settings Control Panel** dialog box, note the following message that displays: **Some settings are managed by your system administrator**. Also note that in the **Important Updates** section, the **Install updates automatically (recommended)** drop-down list box is grayed out, indicating that access is denied.
13. Click **Cancel**, and then close Control Panel.
14. On the **Start** screen, type **Regedit**.
15. Click the **Regedit** tile.
16. Note that you receive the following message: **Registry editing has been disabled by your administrator**.
17. On the taskbar, open File Explorer.
18. In File Explorer, click the **This PC** node.
19. In the **Network Locations** section, note that Marketing-Share is connected to drive M.
20. Right-click the **Start** button, expand the **Shut down or sign out** submenu, and then click **Sign out**.
21. Switch to LON-CL1, click **OK**, and then sign in as **Adatum\Brad** with the password **Pa$$w0rd**.

Note: Notice that prior to signing in, Brad is receiving a compliance message.

22. After signing in, on the desktop, on the charms menu, click **Settings**, and then click **Control Panel**.
23. In Control Panel, click **System and Security**.
24. In System and Security, click **Windows Update**.
25. On the left side, click **Change settings**.
26. In the **Change Settings Control Panel** dialog box, note the following message that displays: **Some settings are managed by your system administrator**. In addition, notice that in the **Important Updates** section, the **Install updates automatically (recommended)** drop-down list box is grayed out, indicating that access is denied.
27. Click **Cancel**, and close Control Panel.
28. On the **Start** screen, type **Regedit**.
29. Click the **Regedit** tile.

    If prompted to confirm **User Account Control**, click **Yes**.

Note: The tool for editing the Registry opens.

30. On the taskbar, open File Explorer.
31. In File Explorer, click the **This PC** node.

Note: The Marketing Share does not display as connected.

32. Close File Explorer.
33. On the **Start** screen, type **cmd**, and then click the **Command Prompt** tile.
34. At the command prompt, type the following command, and then press Enter:

```
whoami /all
```

35. Verify that Brad is a member of the local Administrator group.
36. Right-click the **Start** button, expand the **Shut down or sign out** submenu, and then click **Sign out**.

Results: After completing this exercise, you should have successfully implemented the GPO design that you created.

▶ Task: To prepare for the next module

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

1. On the host computer, start Microsoft® Hyper-V® Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20413C-LON-CL1**.

Module 9: Designing and Implementing an AD DS Physical Topology

Lab: Designing and Implementing an Active Directory Domain Services Physical Topology

Exercise 1: Designing Active Directory Sites and Replication

▶ Task 1: Read the supporting documentation

- Read the supporting documentation provided in the Student Handbook.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the Initial A. Datum Site Design document.

Proposals

1. How can you address problems that users currently have with logon times?

   Because the various network locations currently do not have associated Active Directory sites, it is likely that geographically remote domain controllers are servicing authentication requests from users at those locations, even though a local domain controller exists at some locations. The wide area network (WAN) links between the branches, regional hubs, and Paris have limited capacity, so that is a possible reason for the slow logon time. You need to create Active Directory sites that match the physical locations on the network diagram, and where required, to deploy additional domain controllers.

2. How can you address problems that users currently have with Microsoft® Exchange Server?

   Slow message delivery can indicate an issue with Exchange Server or with name resolution, but it also can indicate that Exchange Transport servers cannot reach appropriate domain controllers and global catalog servers. Because A. Datum does not have a defined Active Directory site topology and it has deployed Exchange Server, it is likely that the email delays occur because the Exchange servers in Paris are contacting other domain controllers over remote links.
   If you establish an Active Directory site topology that is representative of the physical locations in the network, and then deploy additional domain controllers where required, email message delivery problems will reduce.
   You should deploy a domain controller in Paris so that the Exchange server can access a domain controller over a local link.
   Additionally you should plan to deploy domain controllers to the regional hubs as their user counts grow.
   You also should ensure that the local domain controller in the regional hub is a global catalog server.

3. Do you need to create new Active Directory sites? If yes, for which offices?

   Yes, you must create additional sites. You require a site for Paris, and then you need to move the Paris domain controller to the site. This prevents all the European authentication traffic from routing to London. Additionally, you should consider additional sites for each regional hub: Rome, Barcelona, Munich, Athens, Toronto, and Sydney. If you intend to deploy domain controllers to the branches, then you also must configure site objects for each branch.

4. Do you need to restructure any existing sites?

   Yes, you must either rename and reconfigure the default site, or discontinue its use.

5. What else do you need to plan for the new sites?

   After establishing the new sites, you should plan the placement of domain controllers for these sites. Each new site should have at least one domain controller. Additionally, you should verify that the new sites are well connected with other sites with site links, and that the new sites are associated with proper IP subnets. If you do not configure global catalog servers in the new sites, you should consider the option to cache universal group membership.

6. Are there any alternative solutions?

   Yes, there are several alternative solutions. To address the issue with logon times, you can try to increase the bandwidth between Paris and the locations that have the logon problem and that do not have a defined Active Directory site, such as London. For locations that do have an Active Directory site, you should configure at least one domain controller as a global catalog server. For Exchange Server, you can deploy a dedicated Active Directory site for Exchange servers, and then move the Exchange servers to that site. You must also deploy at least one domain controller for that site. By using this approach, you effectively force Exchange Server to communicate only to locally deployed domain controllers, which speeds up directory search and email delivery.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones shown previously.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you should have created a suitable Active Directory site design for A. Datum Corporation.

Exercise 2: Planning the Placement of Domain Controllers and Active Directory Replication

▶ Task 1: Read the supporting documentation

- Read the supporting documentation provided in the Student Handbook.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the Domain Controller Placement Strategy document.

Proposals

1. How will you address problems that users currently have with logon times?

   Users are most likely experiencing slow logon times because of a lack of domain controllers at their locations. Authentication requests travel over WAN links, and when these links are used for other purposes, users experience slow logon times. The solution to this problem is to deploy at least one local domain controller at each location where users are experiencing these problems.

2. How will you avoid having a single point of failure in Active Directory Domain Services (AD DS)?

   Currently, a single point of failure has been detected on operations master (otherwise known as *flexible single master operations*) server roles. All of the operations master roles are located on only one server in the main office. To mitigate this, you should distribute the operations master roles over multiple domain controllers. You can collocate the schema master and the domain-naming master on one domain controller, and the primary domain controller (PDC) emulator and relative ID (RID) master on another domain controller. Additionally, you can locate the infrastructure master on a third domain controller, providing you did not configure it as a global catalog server.

3. How will you provide permanent authentication and Active Directory search services that do not depend on links between locations?

   You can provide authentication that is independent of WAN links by deploying domain controllers on each site. By designating at least one domain controller per site as a global catalog server, you will enable users to search AD DS.

4. How will you provide management and maintenance of IT services in new sales locations?

   These locations will not have dedicated IT staff. Therefore, you should consider splitting the administrator role. Some local employees might perform basic maintenance and administration of local IT services, while a remote administrator performs the remaining administration.

5. How will you address security considerations in the new sales offices?

   Because branch office locations should have a local authentication service provider, this implies installation of a domain controller. The most appropriate solution with consideration to security issues is to deploy a read-only domain controller (RODC). You also can consider increasing security even more by deploying an RODC on a Server Core installation of Windows Server 2012 and newer.

6. How will you achieve rapid deployment of IT services in new sales locations?

   By deploying a domain controller in each location, Active Directory–based services such as Domain Name System (DNS) will provide a faster response time for local users.

7. Are there any alternative solutions?

   You could decide not to install additional domain controllers in each location. To achieve the desired response times, you could increase the speed of the links between the sales office and the site to which they are authenticating. Additionally, to meet the redundancy requirements, you could install backup links between locations.

- Answer the questions in the Active Directory replication Considerations section of the Domain Controller Placement Strategy document.

Active Directory replication Considerations

1. Which sites should you link with Active Directory site links?

   You must create and configure site links between all the newly created sites.

2. Do you need to configure any additional site link attributes? If yes, which attributes should you configure, and how?

   You may want to configure a replication schedule for site links that are established over slower links (for example, from the branch offices to the regional hubs).

3. Do you need to configure bridgehead servers?

   You may want to configure a specific server in London and Paris as a preferred bridgehead server. Both London and Paris effectively are hub sites, with replication traffic passing through them to other sites. Selecting a bridgehead server that has a higher specification may improve Active Directory replication.

4. Do you need to configure site link bridging?

   No. By default, all site links are transitive, and the network is routed fully. Therefore, there is no need for you to configure site link bridging.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones shown previously.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you should have created a suitable domain controller placement strategy.

Exercise 3: Implementing Active Directory Sites and Domain Controllers

▶ Task 1: Install the Paris domain controller

1. On LON-SVR4, on the taskbar, right-click **Start**, and then click **Control Panel**.
2. In Control Panel, click **System and Security**, and then click **System**.
3. In System, click **Change settings**.
4. In the **System Properties** dialog box, click **Change**.
5. In the **Computer Name/Domain Changes** dialog box, in the **Computer name** text box, type **PAR-DC1**, and then click **OK**.
6. In the **Computer Name/Domain Changes** dialog box, click **OK**.
7. In System Properties, click **Close**.
8. Click **Restart Now**.

Note: The lab steps will now refer to 20413C-LON-SVR4 as PAR-DC1.

9. On PAR-DC1 (20413C-LON-SVR4), sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
10. On PAR-DC1, in Server Manager, click **Manage**, and in the drop-down list box, click **Add Roles and Features**.
11. On the **Before you begin** page, click **Next**.
12. On the **Select installation type** page, confirm that **Role-based or feature-based installation** is selected, and then click **Next**.
13. On the **Select destination server** page, ensure that **Select a server from the server pool** is selected, and that **PAR-DC1.adatum.com** is highlighted, and then click **Next**.
14. On the **Select server roles** page, click **Active Directory Domain Services**.
15. On the **Add features that are required for Active Directory Domain Services** page, click **Add Features**, and then click **Next**.
16. On the **Select features** page, click **Next**.
17. On the **Active Directory Domain Services** page, click **Next**.
18. On the **Confirm installation selections** page, click **Install**. This may take a few minutes to complete.
19. When the Active Directory binaries have installed, click **Close**.
20. At the top of the results pane, click the yellow triangle.
21. In the **Post-deployment Configuration** dialog box, click **Promote this server to a domain controller**.
22. In the Active Directory Domain Services Configuration Wizard, on the **Deployment Configuration** page, ensure that the **Add a domain controller to an existing domain** option is selected, and then click **Next**.
23. On the **Domain Controller Options** page, ensure that both the **Domain Name System (DNS) server** and **Global Catalog (GC)** check boxes are selected, and confirm that **Site name** is set to **AdatumHQ**.

24. Under **Type the Directory Services Restore Mode (DSRM) password**, in both the **Password** and **Confirm password** text boxes, type **Pa$$w0rd**, and then click **Next**.
25. On the **DNS Options** page, click **Next**.
26. On the **Additional Options** page, click **Next**.
27. On the **Paths** page, click **Next**.
28. On the **Review Options** page, click **Next**.
29. On the **Prerequisites Check** page, confirm that there are no errors, and then click **Install**. The server will restart automatically.
30. After PAR-DC1 restarts, sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.

▶ Task 2: Create the Active Directory sites

1. On PAR-DC1, in Server Manager, 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, in the **Name** text box, type **Paris**. Under **Select a site link object for this site**, click **DEFAULTIPSITELINK**, and then click **OK**.
4. In the **Active Directory Domain Services** dialog box, click **OK**. The Paris site displays in the navigation pane.
5. In the Active Directory Sites and Services console, in the navigation pane, expand **Sites**, and then click the **Subnets** folder.
6. Right-click **Subnets**, and then click **New Subnet**.
7. In the **New Object – Subnet** dialog box, under **Prefix**, type **10.10.0.0/16**.
8. Under **Select a site object for this prefix**, click **Paris**, and then click **OK**.
9. In the Active Directory Sites and Services console, in the navigation pane, under **Sites**, click the **Subnets** folder.
10. In the Subnets folder, in the details pane, verify that the subnets display and are associated with their correct site.

▶ Task 3: Configure site links

1. On PAR-DC1, in the Active Directory Sites and Services console, in the navigation pane, expand **Sites**, expand **Inter-Site Transports**, and then click the **IP** folder.
2. Right-click the **IP** folder, and then click **New Site Link**.
3. In the **Name** text box, type **LONDON-PARIS**.
4. In the **Sites not in this site link** box, click **AdatumHQ**, press and hold the **Ctrl** key, and then click **PARIS**. Click **Add**, and then click **OK**.
5. Right-click **LONDON-PARIS**, and then click **Properties**.
6. In the **LONDON-PARIS Properties** dialog box, next to **Cost**, change the value to **80**. Next to **Replicate Every**, change the value to **60** minutes, and then click **OK**.

▶ Task 4: Move the new domain controller to the appropriate site

1. On PAR-DC1, in the Active Directory Sites and Services console, in the navigation pane, expand **Sites**, expand **AdatumHQ**, and then expand the **Servers** folder.
2. If the **PAR-DC1** server does not display, right-click the **Servers** folder, and then click **Refresh**.
3. Right-click **PAR-DC1**, and then click **Move**.
4. In the **Move Server** dialog box, click **Paris**, and then click **OK**.
5. In the navigation pane, expand **Paris**, expand **Servers**, right-click **PAR-DC1**, and then click **Properties**.
6. In the **PAR-DC1 Properties** dialog box, in the **Transports available for inter-site data transfer** box, click **IP**, click **Add**, and then click **OK**.

▶ Task 5: To prepare for the next module

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

1. On the host computer, start Microsoft Hyper-V® Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machines** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for 20413C-LON-RTR and 20413C-LON-SVR4.

Results: After completing this exercise, you should have successfully configured Active Directory sites, domain controllers, and replication.

Module 10: Planning and Implementing Storage and File Services

Lab: Planning and Implementing Storage

Exercise 1: Planning a Storage Solution

▶ Task 1: Read the supporting documentation

- Read the documentation provided in the Student Handbook.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the Sales Application Storage Strategy document.

1. How will you configure storage?

   You should configure storage by implementing an Internet small computer system interface (iSCSI) storage area network (SAN).

2. What type of storage is indicated?

   Network storage is indicated.

3. How will you try to ensure that the storage is made highly available?

   Implementing Multipath I/O (MPIO) helps to ensure high availability. Alternatively, you could implement two-way or three-way mirroring or parity mode with Storage Spaces.

4. How could Storage Spaces help address the requirements?

   Storage Spaces can provide an inexpensive storage solution that addresses the need to manage costs.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones shown previously.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you will have planned storage for the sales application.

Exercise 2: Implementing iSCSI Storage

▶ Task 1: Install the iSCSI target

1. If necessary, sign in to LON-DC1 with the user name **Adatum\Administrator** and the password **Pa$$w0rd**.
2. In Server Manager, click **Add roles and features**.
3. In the Add Roles and Features Wizard, on the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, click **Next**.
5. On the **Select destination server** page, ensure that **Select server from the server pool** is selected, and then click **Next**.
6. On the **Select server roles** page, expand **File And Storage Services (2 of 12 Installed)**, expand **File and iSCSI Services (1 of 11 installed)**, select the **iSCSI Target Server** check box, and then click **Next**.
7. On the **Select features** page, click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. When the installation completes, click **Close**.
10. If prompted to restart, click **Restart Now**.
11. Sign in to LON-DC1 with the username **Adatum\Administrator** and the password **Pa$$w0rd**.

▶ Task 2: Configure iSCSI targets

1. On LON-DC1, in Server Manager, in the navigation pane, click **File and Storage Services**.
2. In the File and Storage Services pane, click **iSCSI**.
3. In the iSCSI VIRTUAL DISKS pane, click **TASKS**, and then in the **TASKS** drop-down list 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 **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**, ensure **GB** is selected in the drop-down list box, 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** text 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 **172.16.0.11**, and then click **OK**.
11. On the **Specify access servers** page, click **Next**.
12. On the **Enable Authentication** page, click **Next**.
13. On the **Confirm selections** page, click **Create**.

14. On the **View results** page, wait until creation completes, and then click **Close**.
15. In the iSCSI VIRTUAL DISKS pane, click **TASKS**, and then in the **TASKS** drop-down list box, click **New iSCSI Virtual Disk**.
16. In the New iSCSI Virtual Disk Wizard, on the **Select iSCSI virtual disk location** page, under **Storage location**, click **C:**, and then click **Next**.
17. On the **Specify iSCSI virtual disk name** page, in the **Name** text box, type **iSCSIDisk2**, and then click **Next**.
18. On the **Specify iSCSI virtual disk size** page, in the **Size** text box, type **5**, in the drop-down list box, ensure **GB** is selected, and then click **Next**.
19. On the **Assign iSCSI target** page, click **lon-dc1**, and then click **Next**.
20. On the **Confirm selection** page, click **Create**.
21. On the **View results** page, wait until creation completes, and then click **Close**.

▶ Task 3: Connect to and configure iSCSI targets

1. Switch to LON-SVR1.
2. In Server Manager, click the **Tools** menu, and then click **iSCSI Initiator**.
3. In the **Microsoft iSCSI** message box, click **Yes**.
4. In the **iSCSI Initiator Properties** dialog box, in the **Targets** text box, type **LON-DC1**, and then click **Quick Connect**.
5. In the Quick Connect window, in the **Discovered targets** section, click **iqn.1991-05.com.microsoft:lon-dc1-lon-dc1-target**, and then click **Done**.
6. In the **iSCSI Initiator Properties** dialog box, click **OK**.
7. On LON-SVR1, in Server Manager, click the **Tools** menu, and then click **Computer Management**.
8. In the Computer Management console, under Storage node, click **Disk Management**. Notice that the new disks are added. However, they all are currently offline and not formatted.
9. Close the Computer Management console.

Results: After completing this exercise, you will have successfully implemented an iSCSI SAN.

Exercise 3: Configuring a Redundant Storage Space

▶ Task 1: Create a storage pool by using the iSCSI disks that are attached to the server

1. On LON-SVR1, switch to Server Manager.
2. In Server Manager, in the navigation pane, click **File and Storage Services**, and then in the Servers pane, click **Storage Pools**.
3. In the STORAGE POOLS pane, click **TASKS**, and then in the **TASKS** drop-down list box, click **New Storage Pool**.
4. In the New Storage Pool Wizard, on the **Before you begin** page, click **Next**.
5. On the **Specify a storage pool name and subsystem** page, in the **Name** text box, type **StoragePool1**, and then click **Next**.
6. On the **Select physical disks for the storage pool** page, select both physical disks, and then click **Next**.
7. On the **Confirm selections** page, click **Create**.
8. On the **View results** page, wait until the creation completes, and then click **Close**.

▶ Task 2: Create a mirrored disk

1. On LON-SVR1, in Server Manager, in the STORAGE POOLS pane, click **StoragePool1**.
2. In the VIRTUAL DISKS pane, click **TASKS**, and then in the **TASKS** drop-down list box, click **New Virtual Disk**.
3. In the New Virtual Disk Wizard, on the **Before you begin** page, click **Next**.
4. On the **Select the storage pool** page, click **StoragePool1**, and then click **Next**.
5. On the **Specify the virtual disk name** page, in the **Name** text box, type **Mirrored vDisk**, and then click **Next**.
6. On the **Select the storage layout** page, in the **Layout** drop-down list, click **Mirror**, and then click **Next**.
7. On the **Specify the provisioning type** page, click **Thin**, and then click **Next**.
8. On the **Specify the size of the virtual disk** page, in the **Virtual disk size** text box, type **8**, and then click **Next**.
9. On the **Confirm selections** page, click **Create**.
10. On the **View results** page, wait until the creation completes, ensure that **Create a volume when this wizard closes** is selected, and then click **Close**.
11. In the New Volume Wizard, on the **Before you begin** page, click **Next**.
12. On the **Select the server and disk** page, in the Disk pane, click the **Mirrored vDisk** virtual disk, and then click **Next**.
13. On the **Specify the size of the volume** page, click **Next** to confirm the default selection.
14. On the **Assign to a drive letter or folder** page, in the **Drive letter** drop-down list box, ensure that drive **F** is selected, and then click **Next**.
15. On the **Select file system settings** page, in the **File system** drop-down list box, click **ReFS**. In the **Volume label** text box, type **Mirrored Volume**, and then click **Next**.

16. On the **Confirm selections** page, click **Create**.

17. On the **Completion** page, wait until the creation completes, and then click **Close**.

▶ Task 3: Copy a file and verify that it displays

1. Pause your mouse pointer in the lower left of the taskbar, and then click **Start**.

2. On the Start screen, type **command prompt**, and then press Enter.

3. At the command prompt, type the following command, and then press Enter:

```
Copy C:\windows\system32\write.exe F:\
```

4. Close the command prompt.

5. On the taskbar, click the **File Explorer** icon.

6. In File Explorer, expand **Computer**, and then click **Mirrored Volume (F:)**. You should now see write.exe in the file list.

7. Close File Explorer.

▶ Task 4: Disconnect an iSCSI disk and verify that the file is still accessible

1. Switch to LON-DC1.

2. In Server Manager, in the iSCSI VIRTUAL DISKS pane, in the LON-DC1 list, right-click **iSCSIDisk1.vhd**, and then click **Disable iSCSI Virtual Disk**.

3. In the **Disable iSCSI Virtual Disk** warning message box, click **Yes**.

4. Switch to LON-SVR1.

5. On the taskbar, click the **File Explorer** icon.

6. In File Explorer, click **Mirrored Volume (F:)**.

7. In the file list pane, double-click **write.exe** to verify that access to the file is still available.

8. Close the Document - WordPad window.

9. Close File Explorer.

10. In Server Manager, in the STORAGE POOLS pane, on the menu bar, click **Refresh "Storage Pools"**. Wait until all panes refresh. Notice the warning that displays next to Mirrored vDisk.

11. In the VIRTUAL DISK pane, right-click **Mirrored vDisk**, and then in the drop-down list box, click **Properties**.

12. In the **Mirrored vDisk Properties** dialog box, in the navigation pane, click **Health**.

Note: Notice that the Health Status indicates a Warning. The Operational Status should indicate Degraded.

13. Click **OK** to close the **Mirrored vDisk Properties** dialog box.

Results: After completing this exercise, you will have configured a redundant storage space.

▶ Task: To prepare for the next module

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

1. On the host computer, start Microsoft® Hyper-V® Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-SVR1**, and then click **Revert**.
3. In the **Revert Virtual Machines** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20413C-LON-DC1**.

Module 11: Designing and Implementing Network Protection

Lab: Designing and Implementing Network Protection

Exercise 1: Designing a Windows Firewall Solution

▶ Task 1: Read the supporting documentation

- Read the documentation provided in the Student Handbook.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the Windows Firewall Deployment Strategy document.

Proposals

1. What inbound rules should you implement on servers?

   On each Windows Server® 2012 server, you must implement rules to allow access to all services that run on a particular server. This includes file and printer sharing, Domain Name System (DNS) lookups, domain logons, and any other specific applications. For Windows® client operating system services, the rules are configured already, and you do not need to create any. This applies to file and printer sharing, DNS, and domain logons.
   Rules that you should create to support specific applications are as follows:

- For the Research custom application, create an incoming program rule that allows the Researchsrv.exe executable to receive connections. This is more secure than using the port, because malicious software (otherwise known as *malware*) could also attempt to use the port.

- For the email server, create an incoming program rule that allows the Store.exe executable to receive connections. This is necessary, because the port number is not predictable for each connection.

- For the Customer Service web application, create an incoming port rule that allows access to port 8080. The default rules for Internet Information Services (IIS) allow access to ports 80 and port 443, but not port 8080. This is because you cannot create program rules for IIS.

2. What outbound rules should you implement on servers?

   There are no specific requirements for outbound rules listed on the servers. Windows Firewall is a stateful firewall, and does not require that you create corresponding outbound rules for communication already established by inbound rules. You must configure outbound rules for basic network services, such as DNS lookups, and domain authentication. These are in place by default.

3. What inbound rules should you implement on Windows 8 workstations?

   Client computers have no applications listed that require inbound communication. However, client computers require inbound communication for basic network communication. These rules are in place by default.

4. What outbound rules should you implement on Windows 8 workstations?

   The outbound rules necessary for basic network communication are in place by default. However, you must create outbound rules for other applications:

- For the Research custom application, you should create a program rule to allow Research.exe to communicate on the network. This is more secure than creating a port rule that allows communication to port 10101.
- For Windows Internet Explorer®, you should create a program rule that allows the Iexplore.exe executable access to the network. This prevents use of unsupported web browsers. After you create the program rule, you then edit it to restrict communication to ports 80, 443, and 8080.

5. How will you deploy Windows Firewall on servers that are running Windows Server?

   For the highest level of security, you should implement only the rules necessary on each server. Therefore, you should configure each server individually. In an organization with many servers running the same applications, you could apply rules by using Group Policy.

6. How will you deploy Windows Firewall on workstations?

   You should configure workstations with the necessary Windows Firewall rules by using Group Policy. If desired, you can create customized group policies for various workgroups that include only the necessary applications for each workgroup. To support this, each workgroup should have a separate organizational unit (OU) in Active Directory® Domain Services (AD DS).

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones shown previously.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you should have completed the Windows Firewall design for A. Datum.

Exercise 2: Implementing a Windows Firewall solution

▶ Task 1: Move the computers to the Research OU

1. Switch to LON-DC1.
2. In Server Manager, click **Tools**, and then click **Active Directory Users and Computers**.
3. In Active Directory Users and Computers, expand **Adatum.com**, and then click **Computers**.
4. Right-click **LON-CL1**, and then click **Move**.
5. In the **Move** dialog box, click **Research**, and then click **OK**.
6. Right-click **LON-SVR1**, and then click **Move**.
7. In the **Move** dialog box, click **Research**, and then click **OK**.

▶ Task 2: Create a Group Policy Object (GPO) and link it to the Research OU

1. Switch to Server Manager.
2. In Server Manager, click **Tools**, and then click **Group Policy Management**.
3. In Group Policy Management, expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**, and then click **Research**.
4. Right-click **Research**, and then click **Create a GPO in this domain, and link it here**.
5. In the **New GPO** dialog box, in the **Name** text box, type **Research Department Application Security Policy**, and then click **OK**.

▶ Task 3: Create the required connection security rules

1. In Group Policy Management, expand **Research**.
2. Right-click **Research Department Application Security Policy**, and then click **Edit**.
3. In the Group Policy Management Editor, under **Computer Configuration**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, expand **Windows Firewall with Advanced Security**, expand **Windows Firewall with Advanced Security – LDAP://CN={GUID}**, and then click **Connection Security Rules**.
4. Right-click **Connection Security Rules**, and then click **New Rule**.
5. In the New Connection Security Rule Wizard, on the **Rule Type** page, click **Custom**, and then click **Next**.
6. On the **Endpoints** page, click **Next**.
7. On the **Requirements** page, click **Require authentication for inbound connections and request authentication for outbound connections**, and then click **Next**.
8. On the **Authentication Method** page, click **Computer and user (Kerberos V5)**, and then click **Next**.
9. On the **Protocol and Ports** page, in the **Protocol type** list, click **TCP**.
10. In the **Endpoint 1 port** list, click **Specific Ports**, in the text box, type **80**, and then click **Next**.
11. On the **Profile** page, clear the **Private** and **Public** check boxes, and then click **Next**.
12. On the **Name** page, in the **Name** text box, type **Research Department Application Security rule**, and then click **Finish**.

▶ Task 4: Create the firewall rules

1. In the Group Policy Management Editor, right-click **Inbound Rules**, and then click **New Rule**.
2. In the New Inbound Rule Wizard, on the **Rule Type** page, click **Custom**, and then click **Next**.
3. On the **Program** page, click **Next**.
4. On the **Protocol and Ports** page, in the **Protocol type** list, click **TCP**.
5. In the **Local port** list, click **Specific Ports**, in the text box, type **80**, and then click **Next**.
6. On the **Scope** page, click **Next**.
7. On the **Action** page, click **Allow the connection if it is secure**, and then click **Customize**. Ensure that **Allow the connection if it is authenticated and integrity-protected** is selected, click **OK**, and then click **Next**.
8. On the **Users** page, click **Next**.
9. On the **Computers** page, click **Only allow connections from these computers**, and then click **Add**.
10. In the **Select Computers, or Groups** dialog box, in the **Enter the object names to select (examples)** text box, type **LON-CL1; LON-SVR1**, click **Check Names**, click **OK**, and then click **Next**.
11. On the **Profile** page, clear the **Private** and **Public** check boxes, and then click **Next**.
12. On the **Name** page, in the **Name** text box, type **Research Department Application Firewall rule**, and then click **Finish**.

▶ Task 5: Refresh the Group Policy settings

1. Switch to LON-CL1.
2. On the Start screen, type **cmd.exe**, and then press Enter.
3. At the command prompt, type the following command, and then press Enter:

```
Gpupdate /force
```

4. At the command prompt, type the following command, and then press Enter:

```
Shutdown /r /t 0
```

5. Switch to LON-SVR1.
6. Pause your mouse pointer over the lower left of the taskbar, and then click **Start**.
7. On the Start screen, type **cmd.exe**, and then press Enter.
8. At the command prompt, type the following command, and then press Enter:

```
Gpupdate /force
```

9. At the command prompt, type the following command, and then press Enter:

```
Shutdown /r /t 0
```

▶ Task 6: Attempt to connect to the web server

1. Switch to LON-CL1. You must wait until LON-SVR1 has restarted before proceeding.
2. Sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
3. On the **Start** screen, click **Desktop**.

4. On the taskbar, click the **Internet Explorer** icon.
5. In Internet Explorer, in the Address bar, type **http://LON-svr1**, and then press Enter.
6. Verify that the default IIS 8 webpage displays.

▶ Task 7: Verify settings

1. Click **Start**, on the Start screen, type **Windows Firewall**, and click **Windows Firewall**.
2. In Windows Firewall, click **Advanced settings**.
3. In Windows Firewall with Advanced Security, in the navigation pane, expand **Monitoring**, expand **Security Associations**, and then click **Main Mode**.
4. In the right pane, double-click the item listed.

   What is the First authentication method?
   Answer: Computer (Kerberos V5)
5. Click **OK**, and then click **Quick Mode**.
6. In the right pane, double-click the item listed.

   What is the Remote port?
   Answer: TCP port 80
7. Click **OK**.

▶ Task 8: To prepare for the next exercise

When you are finished with this exercise, revert two of the virtual machines to their initial state, and start two additional virtual machines. To do this perform the following steps:

1. On the host computer, start Microsoft® Hyper-V® Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-CL1**, and then click **Revert**.
3. In the **Revert Virtual Machines** dialog box, click **Revert**.
4. Repeat step 2 and 3 for 20413C-LON-SVR1 and 20413C-LON-DC1.
5. Click **20413C-LON-DC1**, and in the Actions pane, click **Start**.
6. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
7. Sign in using the following credentials:

- User name: **Administrator**
- Password: **Pa$$w0rd**
- Domain: **Adatum**

8. Repeat steps 5 through 7 for 20413C-LON-SVR2, and 20413C-LON-CL1.

Results: After completing this exercise, you should have configured the required firewall rules.

Exercise 3: Designing a Network Access Protection (NAP) Solution

▶ Task 1: Read the supporting documentation

- Read the documentation provided in the Student Handbook.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the NAP Deployment Strategy document.

Proposals

1. Which NAP enforcement method is appropriate for the given scenario?

   The Internet Protocol security (IPsec) enforcement method is appropriate for this scenario. You can configure the servers in the Research department so that they will not communicate with workstations that have not passed the health check.

2. What is the simplest way to apply the necessary client configuration to several computers simultaneously?

   You can use Group Policy to enable the enforcement client for each type of enforcement. You can do this with any OUs that contain client computers.

3. How will you ensure that the configuration applies to only client computers, and not to servers?

   If client computers are in separate OUs, then you can link the GPO only to those OUs with client computers. Alternatively, if client computers and servers exist in the same OU, you can use security filtering to ensure that only client computers can apply the policy. You should create a group for the client computers, and then ensure that only that group has the necessary permissions to apply the GPO.

4. What determines the options available for verifying client computer status? How can you expand these options?

   The system health agents (SHA) and system health validators (SHV) that are included with Windows 8, Windows 8.1, Windows Server 2012, and Windows Server 2012 R2 determine the options available for verifying client computer status. You use additional SHAs and SHVs to expand the NAP monitoring capabilities. You add an SHA and SHV as a pair, with the SHA on the client computer side, and the SHV on the server side.

5. How do noncompliant computers access remediation servers?

   You should configure remediation servers to accept traffic that has not been authenticated. This enables the noncompliant computer to make contact with remediation servers even though its overall network access is restricted.

6. Which servers should you configure as remediation servers?

   You should configure all servers that are necessary to bring a computer into compliance as remediation servers. This can include domain controllers, DNS servers, and Windows Server Update Services (WSUS) servers.

7. In addition to the existing servers hosting the Domain Controller, DNS, DHCP, and WSUS server roles, what additional roles should you plan to deploy to implement this solution?

   You should plan to deploy an enterprise certification authority (CA) and a Network Policy Server (NPS). The CA will work in conjunction with the Health Registration Authority (HRA) role service on the server with the NPS role installed, to issue health certificates to NAP-compliant client computers.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones shown previously.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you should have completed the NAP design.

Exercise 4: Implementing NAP with IPsec Enforcement

▶ Task 1: Configure NAP health compliance certificates

1. Switch to LON-DC1.
2. On the Start screen, type **certsrv.msc**, and then press Enter.
3. In the Certificate Services window, in the left pane, expand **AdatumCA**.
4. Right-click **Certificate Templates**, and then click **Manage**.
5. In the Certificate Templates console, right-click **Workstation Authentication**, and then click **Duplicate Template**.
6. On the **General** tab, in the **Template display name** text box, type **NAP Health Certificate**.
7. On the **Subject Name** tab, click **Supply in the request**.
8. When the Certificate Templates warning displays, click **OK**.
9. On the **Extensions** tab, click **Application Policies**, and then click **Edit**.
10. In the **Edit Application Policies Extension** dialog box, click **Add**, click **System Health Authentication**, and then to return to the **Properties of New Template** dialog box click **OK** twice.
11. In the **Properties of New Template** dialog box, click **OK**.
12. Close the Certificate Templates console.
13. In the Certification Authority console, right-click **Certificate Templates**, click **New**, and then click **Certificate Template to Issue**.
14. Click **NAP Health Certificate**, and then click **OK**.
15. In the Certification Authority console, in the left pane, right-click **AdatumCA**, and then click **Properties**.
16. In the **AdatumCA Properties** dialog box, on the **Security** tab, click **Add**.
17. Click **Object Types**, select the **Computers** check box, and then click **OK**.
18. In the **Select Users, Computers, Service Accounts, or Groups** dialog box, in the **Enter the object names to select** text box, type **LON-SVR2**, and then click **OK**.
19. On the **Security** tab, click **LON-SVR2**, for the **Read**, **Issue and Manage Certificates**, **Manage CA**, and **Request Certificates** check boxes, select **Allow**, and then click **OK**.
20. Close the Certification Authority console.
21. On the taskbar, right-click **Windows PowerShell**, and then click **Run as Administrator**.
22. In Windows PowerShell, at the command prompt, type the following command, and then press Enter:

    ```
    Certutil -setreg policy\EditFlags +EDITF_ATTRIBUTEENDDATE
    ```

23. At the command prompt, type the following command, and then press Enter:

    ```
    net stop certsvc
    ```

24. At the command prompt, type the following command, and then press Enter:

    ```
    net start certsvc
    ```

▶ Task 2: Install the Network Policy and Access Services role on LON-SVR2

1. Switch to LON-SVR2.
2. Click **Start**, and on the Start screen, type **mmc.exe**, and then press Enter.
3. On the **File** menu, click **Add/Remove Snap-in**.
4. In the **Add or Remove Snap-ins** dialog box, click **Certificates**, click **Add**, click **Computer account**, click **Next**, and then click **Finish**.
5. In the **Add or Remove Snap-ins** dialog box, click **OK**.
6. In the console, expand **Certificates**, right-click **Personal**, point to **All Tasks**, and then click **Request New Certificate**.
7. In the **Certificate Enrollment** dialog box, click **Next**.
8. On the **Select Certificate Enrollment Policy** page, click **Active Directory Enrollment Policy**, and then click **Next**.
9. Select the **Computer** check box, and then click **Enroll**.
10. Verify that the status of certificate installation is **Succeeded**, and then click **Finish**.
11. Close the Console1 window.
12. When prompted to save console settings, click **No**.
13. In the Server Manager console, on the Dashboard, click **Add roles and features**.
14. To get to the server role selection screen, click **Next** three times.
15. In the **Select server roles** dialog box, click **Network Policy and Access Services**, click **Add Features**, and then click **Next**.
16. Click **Next** two more times.
17. In the **Select role services** dialog box, select the **Network Policy Server** and the **Health Registration Authority** check boxes, click **Add Features**, and then click **Next**.
18. On the **Certification Authority** page, click **Use an existing remote CA**, and then click **Select**.
19. On the **Select Certification Authority** page, click **AdatumCA**, click **OK**, and then click **Next**.
20. On the **Authentication requirements** page, click **Yes, require requestors to be authenticated as members of a domain**, and then click **Next**.
21. On the **Server Authentication Certificate** page, select the **LON-SVR2.Adatum.com** certificate that displays in the list, and then click **Next**.
22. Click **Next** two more times.
23. On the **Confirm installation selections** page, click **Install**.
24. On the **Installation progress** page, verify that the installation is successful, and then click **Close**.

▶ Task 3: Configure the HRA

1. On LON-SVR2, in Server Manager, click **Tools**, and then click **Health Registration Authority**.
2. In the left pane, expand **Health Registration Authority (Local Computer)**, and then click **Certification Authority**.
3. In the details pane, confirm that **LON-DC1.Adatum.com\AdatumCA** is listed.
4. In the left pane, right-click **Certification Authority**, and then click **Properties**.

5. In the **Certificate Authorities Properties** dialog box, click **Use enterprise certification authority**.
6. In both the **Authenticated compliant certificate template** and **Anonymous compliant certificate template**, select the **NAP Health Certificate** template.
7. Verify that the validity period for certificates approved by this Health Registration Authority is set to **4** hours, and then click **OK**.
8. Close the Health Registration Authority window.

▶ Task 4: Configure health policies

1. On LON-SVR2, in Server Manager, click **Tools**, and then click **Network Policy Server**.
2. In the details pane, click **Configure NAP**.
3. In the Configure NAP Wizard, on the **Select Network Connection Method For Use with NAP** page, in **Network connection method**, click **IPsec with Health Registration Authority (HRA)**, and then click **Next**.
4. On the **Specify Enforcement Servers Running HRA** page, click **Next**.
5. On the **Configure Machine Groups** page, click **Next**.
6. On the **Define NAP Health Policy** page, clear the **Enable auto-remediation of client computers** check box, and then click **Next**.
7. Click **Finish**.

▶ Task 5: Configure the health validator

1. On LON-SVR2, in the Network Policy Server console, in the left pane, expand **Network Access Protection**.
2. Expand **System Health Validators**, expand **Windows Security Health Validators**, and then click **Settings**.
3. In the details pane, double-click **Default configuration**.
4. On the **Windows 8/Windows 7/Windows Vista** tab, clear all check boxes, and then select the **A firewall is enabled for all network connections** check box.
5. On the **Windows Security Validator** page, click **OK**.

▶ Task 6: Create the NAP client configuration GPO

1. Switch to LON-DC1.
2. In Server Manager, click **Tools**, and then click **Active Directory Users and Computers**.
3. In the Active Directory Users and Computers console, click **Adatum.com**. On the Action menu, click **New**, and then click **Organizational Unit**.
4. In the **New Object – Organizational Unit** dialog box, type **NAP_Clients**, and then click **OK**.
5. Click the **Computers** node. Right-click **LON-CL1**, and then click **Move**.
6. In the **Move** dialog box, click **NAP_Clients**, and then click **OK**.
7. In Server Manager, click **Tools**, and then click **Group Policy Management**.
8. In the Group Policy Management Console, in the left panel, expand **Forest: Adatum.com**, expand **Domains**, and then expand **Adatum.com**.
9. Right-click **NAP_Clients**, and then click **Create a GPO in this domain, and Link it here**.

10. In the **New GPO** dialog box, in the **Name** text box, type **NAP Client Configuration**, and then click **OK**.
11. Click **NAP Client Configuration**, and then click **OK**.
12. Expand **NAP_Clients**, in the console tree, right-click **NAP Client Configuration**, and then click **Edit**.
13. In the Group Policy Management Editor, in the console tree, expand **Computer Configuration**.
14. Expand **Policies**, expand **Windows Settings**, expand **Security Settings**, expand **Network Access Protection**, expand **NAP Client Configuration**, and then click **Enforcement Clients**.
15. In the details pane, right-click **IPsec Relying Party**, and then click **Enable**.
16. In the console tree, expand **Health Registration Settings**, right-click **Trusted Server Groups**, and then click **New**.
17. On the **Group Name** page, type **Internal**, and then click **Next**.
18. On the **Add Servers** page, type **https://lon-svr2.adatum.com/domainhra/hcsrvext.dll**, click **Add**, and then click **Finish**.
19. In the console tree, under **Security Settings**, click **System Services**.
20. In the details pane, double-click **Network Access Protection Agent**.
21. In the **Network Access Protection Agent Properties** dialog box, select the **Define this policy setting** check box.
22. Under **Select services startup mode**, click **Automatic**, and then click **OK**.
23. In the console tree, expand **Administrative Templates: Policy definitions (ADMX files) retrieved from the local computer**, expand **Windows Components**, and then click **Security Center**.
24. In the details pane, double-click **Turn on Security Center**.
25. In the **Turn on Security Center** dialog box, click **Enabled**, and then click **OK**.
26. Close the editor.

▶ Task 7: Test a client on the internal network

1. Restart LON-CL1, and sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
2. On LON-CL1, on the **Start** screen, type **Windows PowerShell**, and then press Enter.
3. At the Windows PowerShell command prompt, type the following command, and then press Enter:

```
gpupdate /force
```

4. At the command prompt, type the following command, and then press Enter:

```
get-service napagent
```

Note: This cmdlet verifies that the NAP agent is running.

5. At the command prompt, type the following command, and then press Enter:

```
netsh nap client show grouppolicy
```

Note: This command verifies that the correct GPO is being applied to the client, and that the IPsec Relying Party agent is running.

▶ Task 8: Test compliance

1. Switch to Windows PowerShell.
2. At the command prompt, type the following command, and then press Enter:

   ```
   napstat
   ```

3. Notice that the client is fully compliant.
4. Click **Start**, type **Services.msc**, and then press Enter.
5. Double-click **Windows Firewall**.
6. In the **Startup type** list, click **Disabled**.
7. Click **Stop**, and then click **OK**.
8. Switch back to the Windows PowerShell window.
9. At the Windows PowerShell command prompt, type the following command, and then press Enter:

   ```
   Napstat
   ```

10. Notice that the client is no longer compliant.

▶ Task 9: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. In the Virtual Machines list, right-click **20413C-LON-CL1**, and then click **Revert**.
3. In the **Revert Virtual Machines** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for 20413C-LON-SVR2 and 20413C-LON-DC1.

Results: After completing this exercise, you should have implemented NAP with IPsec enforcement.

Module 12: Designing and Implementing Remote Access Services

Lab: Designing and Implementing Network Access Services

Exercise 1: Designing a Remote Access Strategy

▶ Task 1: Read the supporting documentation

- Read the documentation provided in the Student Handbook.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the Remote Access Services Strategy document.

Proposals

1. What technology should you use to provide automatic client connectivity?

   You can use DirectAccess to provide automatic connectivity for clients running the Windows® 8.1, Windows 8, and Windows 7 operating systems.

2. How can you provide secure access for computers that are not joined to any domain?

   You can integrate a virtual private network (VPN) and Network Access Protection (NAP) to provide secure access to internal resources from computers that are not joined to any domain.

3. How can you provide Contoso users with access to the secure website on the A. Datum internal network?

   You can use Web Application Proxy to publish the secure website on the A. Datum internal network to Contoso users that require remote access to the site.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones shown previously.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

Results: After completing this exercise, you should have successfully designed a remote access strategy.

Exercise 2: Planning and Implementing a DirectAccess Solution

▶ Task 1: Read the supporting documentation

- Read the documentation provided in the Student Handbook.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the DirectAccess Strategy document.

Proposals

1. What components must be in place to support DirectAccess?

   The following components must be in place to support DirectAccess:

   - Active Directory® Domain Services (AD DS). You must deploy at least one Active Directory domain. Workgroups are not supported.
   - Group Policy. You need Group Policy for centralized administration and deployment of DirectAccess client settings.
   - Domain Name System (DNS) and a domain controller. You must have at least one domain controller and at least one DNS server running Windows Server® 2012 R2, Windows Server 2012, Windows Server 2008 Service Pack 2 (SP2), or Windows Server 2008 R2.
   - Optionally, a PKI. If you deploy NAP, you will need to use a public key infrastructure (PKI) to issue computer certificates for authentication and health certificates.
   - IPsec policies. DirectAccess utilizes Internet Protocol security (IPsec) policies that you configure and administer as part of Windows Firewall with Advanced Security.
   - ICMPv6 Echo Request traffic. You must create separate inbound and outbound rules that allow Internet Control Message Protocol version 6 (ICMPv6) Echo Request messages.
   - IPv6 and transition technologies. IPv6 and the transition technologies such as ISATAP, Teredo, and 6to4 must be available for use on the DirectAccess server.

2. Will you implement a PKI?

   A PKI deployment is optional for simplified configuration and management. DirectAccess in Windows Server 2012 R2 enables client authentication requests to be sent over a HTTPS–based Kerberos authentication proxy service that is running on the DirectAccess server for clients running Windows 8.1 or Windows 8. This eliminates the need for establishing a second IPsec tunnel between clients and domain controllers. The Kerberos authentication proxy will send Kerberos requests to domain controllers on behalf of the client.

   However, for a full DirectAccess configuration that allows NAP integration, two-factor authentication, and force tunneling, you still need to implement certificates for authentication for every client that will participate in DirectAccess communication.

3. What must you configure on the DNS servers to support your planned deployment?

   You must configure a DNS host (A or AAAA) resource record for the network location service, and remove ISATAP from the network.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones shown previously.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

▶ Task 5: Configure AD DS and DNS

1. Create a security group for DirectAccess client computers by performing the following steps:
    a. Switch to LON-DC1.
    b. In Server Manager, click **Tools**, and then click **Active Directory Users and Computers**.
    c. In Active Directory Users and Computers, right-click **Adatum.com**, click **New**, and then click **Organizational Unit**
    d. In the **New Object – Organizational Unit** dialog box, in the **Name** text box, type **DA_Clients OU**, and then click **OK**.
    e. In Active Directory Users and Computers, expand **Adatum.com**, right-click **DA_Clients OU**, click **New**, and then click **Group**.
    f. In the **New Object - Group** dialog box, in the **Group name** text box, type **DA_Clients**.
    g. Under **Group scope**, click **Global**, under **Group type**, click **Security**, and then click **OK**.
    h. In the details pane, double-click **DA_Clients**.
    i. In the **DA_Clients Properties** dialog box, click the **Members** tab, and then click **Add**.
    j. In the **Select Users, Contacts, Computers, Service Accounts, or Groups** dialog box, click **Object Types**, select the **Computers** check box, and then click **OK**.
    k. In the **DA_Clients Properties** dialog box, under **Enter the object names to select (examples)**, type **LON-CL1**, and then click **OK**.
    l. Verify that LON-CL1 displays below **Members**, and then click **OK**.
    m. Close Active Directory Users and Computers.
2. Configure firewall rules for ICMPv6 traffic by performing the following steps:

Note: You must configure firewall rules for ICMPv6 traffic to enable subsequent testing of DirectAccess in the lab environment.

    a. In Server Manager, click **Tools**, and then click **Group Policy Management**.
    b. In the Group Policy Management Console, expand **Forest: Adatum.com**, expand **Domains**, and then expand **Adatum.com**.
    c. Under **Adatum.com**, right-click **Default Domain Policy**, and then click **Edit**.
    d. In the Group Policy Management Editor, expand **Computer Configuration**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, expand **Windows Firewall with Advanced Security**, and then click **Windows Firewall with Advanced Security**.
    e. In Windows Firewall with Advanced Security, click **Inbound Rules**, right-click **Inbound Rules**, and then click **New Rule**.
    f. On the **Rule Type** page, click **Custom**, and then click **Next**.
    g. On the **Program** page, click **Next**.
    h. On the **Protocols and Ports** page, under **Protocol type**, click **ICMPv6**, and then click **Customize**.

i. In the **Customize ICMP Settings** dialog box, click **Specific ICMP types**, click **Echo Request**, click **OK**, and then click **Next**.

j. On the **Scope** page, click **Next**.

k. On the **Action** page, click **Next**.

l. On the **Profile** page, click **Next**.

m. On the **Name** page, in the **Name** text box, type **Inbound ICMPv6 Echo Requests**, and then click **Finish**.

n. Switch to the Group Policy Management Editor. In the console tree, click **Outbound Rules**, right-click **Outbound Rules**, and then click **New Rule**.

o. On the **Rule Type** page, click **Custom**, and then click **Next**.

p. On the **Program** page, click **Next**.

q. On the **Protocols and Ports** page, under **Protocol type**, click **ICMPv6**, and then click **Customize**.

r. In the **Customize ICMP Settings** dialog box, click **Specific ICMP types**, click **Echo Request**, click **OK**, and then click **Next**.

s. On the **Scope** page, click **Next**.

t. On the **Action** page, click **Allow the connection**, and then click **Next**.

u. On the **Profile** page, click **Next**.

v. On the **Name** page, in the **Name** text box, type **Outbound ICMPv6 Echo Requests**, and then click **Finish**.

w. Close the Group Policy Management Editor and the Group Policy Management Console.

3. Create required DNS records by performing the following steps:

a. In Server Manager, click **Tools**, and then click **DNS**.

b. In the DNS Manager console, expand **LON-DC1**, expand **Forward Lookup Zones**, and then click **Adatum.com**.

c. Right-click **Adatum.com**, and then click **New Host (A or AAAA)**.

d. In the **Name** text box, type **nls**. In the **IP address** text box, type **172.16.0.11**, click **Add Host**, and then click **OK**.

e. In the **DNS** dialog box informing you that the record was created, click **OK**.

f. In the **New Host** dialog box, click **Done**.

g. Close the DNS Manager console.

4. Remove ISATAP from the DNS global query block list by performing the following steps:

a. Move the mouse pointer to the lower-right corner, and then click **search**.

b. In the **Search** field, type **cmd.exe**, and then press Enter.

c. At the command prompt, type the following command, and then press Enter:

```
dnscmd /config /globalqueryblocklist wpad
```

Note: Ensure that the **Command completed successfully** message displays.

d. Close the Command Prompt window.

5. Configure the DNS suffix on LON-RTR by performing the following steps:

   a. Switch to LON-RTR.

   b. Move the mouse to the lower right corner of the screen, click **Settings**, and then click **Control Panel**.

   c. In Control Panel, click **View network status and tasks**.

   d. In the Network and Sharing Center window, click **Change adapter settings**.

   e. In the Network Connection window, right-click **Ethernet**, and then click **Properties**.

   f. In the **Ethernet Properties** dialog box, double-click **Internet Protocol Version 4 (TCP/IPv4)**.

   g. In the **Internet Protocol Version 4 (TCP/IPv4)** dialog box, click **Advanced**.

   h. On the **DNS** tab, in the **DNS suffix for this connection** text box, type **Adatum.com**, and then click **OK**.

   i. In the **Internet Protocol Version 4 (TCP/IPv4)** dialog box, click **OK**.

   j. In the **Ethernet Properties** dialog box, click **OK**.

6. Configure the Ethernet 2 properties on LON-RTR by performing the following steps:

   a. In the Network Connection window, right-click **Ethernet 2**, and then click **Properties**.

   b. In the **Ethernet 2 Properties** dialog box, double-click **Internet Protocol Version 4 (TCP/IPv4)**.

   c. In the **Internet Protocol Version 4 (TCP/IPv4)** dialog box, in the **IP address** text box, type **131.107.0.2**, and in the **Subnet mask** text box, type **255.255.0.0**.

   d. Click **OK**, and then click **OK** again.

   e. Close the Network Connections window.

▶ Task 6: Configure required certificates

1. Configure the CRL distribution settings by performing the following steps:

   a. On LON-DC1, in Server Manager, on the **Tools** menu, click **Certification Authority**.

   b. In the details pane, right-click **AdatumCA**, and then click **Properties**.

   c. In the **AdatumCA Properties** dialog box, click the **Extensions** tab.

   d. On the **Extensions** tab, click **Add**. In the **Location** text box, type **http://crl.adatum.com/crld/**.

   e. Under **Variable**, click <**CaName**>, and then click **Insert**.

   f. Under **Variable**, click <**CRLNameSuffix**>, and then click **Insert**.

   g. Under **Variable**, click <**DeltaCRLAllowed**>, and then click **Insert**.

   h. In the **Location** text box, at the end of the **Location** string, type **.crl**, and then click **OK**.

   i. Select both the **Include in CRLs. Clients use this to find Delta CRL locations** and **Include in the CDP extension of issued certificates** check boxes, and then click **Apply**.

   j. In the dialog box asking you to restart Active Directory Certificate Services, click **No**.

   k. Click **Add**.

   l. In the **Location** text box, type **\\LON-RTR\crldist$\**.

   m. Under **Variable**, click <**CaName**>, and then click **Insert**.

   n. Under **Variable**, click <**CRLNameSuffix**>, and then click **Insert**.

o. Under **Variable**, click <**DeltaCRLAllowed**>, and then click **Insert**.

p. In the **Location** text box, at the end of the string, type **.crl**, and then click **OK**.

q. Select both the **Publish CRLs to this location** and **Publish Delta CRLs to this location** check boxes, and then click **OK**.

r. Click **Yes** to restart Active Directory Certificate Services (AD CS).

2. Duplicate the web certificate template and configure appropriate permission by performing the following steps:

a. In the Certification Authority console, expand **AdatumCA**, right-click **Certificate Templates**, and then click **Manage**.

Note: Users require the Enroll permission on the certificate.

b. In the Certificate Templates console, in the content pane, right-click the **Web Server** template, and then click **Duplicate Template**.

c. Click the **General** tab, and in the **Template display name** text box, type **Adatum Web Server Certificate**.

d. Click the **Request Handling** tab, and then click **Allow private key to be exported**.

e. Click the **Security** tab, and then click **Authenticated Users**.

f. In the Permissions for Authenticated Users window, under **Allow**, click **Enroll**, and then click **OK**.

g. Close the Certificate Templates console.

h. In the Certification Authority console, right-click **Certificate Templates**, point to **New**, and then click **Certificate Template to Issue**.

i. Click **Adatum Web Server Certificate**, and then click **OK**.

j. Right-click **AdatumCA**, point to **All Tasks**, and then click **Stop Service**.

k. Right-click **AdatumCA**, point to **All Tasks**, and then click **Start Service**.

l. Close the Certification Authority console.

3. Configure computer certificate autoenrollment by performing the following steps:

a. On LON-DC1, switch to Server Manager, click **Tools**, and then click **Group Policy Management**.

b. In the Group Policy Management Console, expand **Forest: Adatum.com**, expand **Domains**, and then expand **Adatum.com**.

c. In the Adatum.com console, right-click **Default Domain Policy**, and then click **Edit**.

d. In the Group Policy Management Editor, expand **Computer Configuration**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, and then click **Public Key Policies**.

e. In the Public Key Policies details pane, right-click **Automatic Certificate Request Settings**, point to **New**, and then click **Automatic Certificate Request**.

f. In the Automatic Certificate Request Setup Wizard, click **Next**.

g. On the **Certificate Template** page, click **Computer**, click **Next**, and then click **Finish**.

h. Close both the Group Policy Management Editor and the Group Policy Management Console.

▶ Task 7: Configure internal resources

1. Start LON-SVR1, and then request a certificate for LON-SVR1 by performing the following steps:

   a. On the Start screen, type **cmd**, and then press Enter.

   b. At the command prompt, type the following command, and then press Enter:

   ```
   mmc
   ```

   c. In the Microsoft Management Console (MMC), click **File**, and then click **Add/Remove Snap-in**.

   d. In the **Add/Remove Snap-in** dialog box, click **Certificates**, click **Add**, click **Computer account**, click **Next**, click **Local computer**, click **Finish**, and then click **OK**.

   e. In the Certificates snap-in console, expand **Certificates (Local Computer)**, expand **Personal**, and then click **Certificates**.

   f. Right-click **Certificates**, point to **All Tasks**, and then click **Request New Certificate**.

   g. Click **Next** twice.

   h. On the **Request Certificates** page, click **Adatum Web Server Certificate**, and then click **More information is required to enroll for this certificate**.

   i. In the **Certificate Properties** dialog box, on the **Subject** tab, under **Subject name**, under **Type**, click **Common name**.

   j. In the **Value** text box, type **nls.adatum.com**, and then click **Add**.

   k. Click **OK**, click **Enroll**, and then click **Finish**.

   l. In the Certificates snap-in, in the details pane, verify that a new certificate with the name **nls.adatum.com** displays with the **Intended Purposes** of **Server Authentication**.

   m. Close the console window. When you are prompted to save settings, click **No**.

2. Change the HTTPS bindings by performing the following steps:

   a. In Server Manager, click **Tools**, and then click **Internet Information Services (IIS) Manager**.

   b. In the Internet Information Services (IIS) Manager console, expand **LON-SVR1**, expand **Sites**, and then click **Default Web Site**.

   c. In the Actions pane, click **Bindings**, in the list of bindings, click **https**, and then click **Edit**.

   d. In the **Edit Site Bindings** dialog box, in the **SSL Certificate** dialog box, click the **nls.adatum.com** certificate, click **OK**, and then click **Close**.

   e. Close the Internet Information Services (IIS) Manager console.

3. Create a shared folder on LON-SVR1 by performing the following steps:

   a. On LON-SVR1, on the taskbar, click the **File Explorer** icon.

   b. In File Explorer, click **This PC**, and then double-click **Local Disk (C:)**.

   c. Right-click an area of free space, point to **New**, and then click **Folder**.

   d. Type **Files**, and then press Enter.

   e. Right-click **Files**, and then click **Properties**.

   f. In the **Files Properties** dialog box, click the **Sharing** tab, and then click **Advanced Sharing**.

   g. In the **Advanced Sharing** dialog box, select the **Share this folder** check box, click **OK**, and then click **Close**.

h. In File Explorer, Double-click **Files**.

i. In the Files pane, right-click and area of free space, point to **New**, and then click **Text Document**.

j. Type **DirectAccess**, and then press Enter.

k. Double-click the **DirectAccess** file.

l. In Notepad, type **This is a corporate file**.

m. Close Notepad, and when prompted, click **Save**.

n. Close File Explorer.

▶ Task 8: Configure a DirectAccess server

1. Obtain required certificates for LON-RTR by performing the following steps:

a. Switch to LON-RTR.

b. Move the mouse to the lower-right corner of the screen, and click **Search**.

c. In the **Search** box, type **cmd**, and then press Enter.

d. At the command prompt, type the following command, and then press Enter:

```
gpupdate /force
```

e. At the command prompt, type the following command, and then press Enter:

```
mmc.exe
```

f. In the MMC, click **File**, and then click **Add/Remove Snap-in**.

g. In the **Add/Remove Snap-in** dialog box, click **Certificates**, click **Add**, click **Computer account**, and then click **Next**.

h. Click **Local computer**, click **Finish**, and then click **OK**.

i. In the Certificates snap-in console, expand **Certificates (Local Computer)**, expand **Personal**, and then click **Certificates**.

j. Right-click **Certificates**, point to **All Tasks**, and then click **Request New Certificate**.

k. Click **Next** twice.

l. On the **Request Certificates** page, click **Adatum Web Server Certificate**, and then click **More information is required to enroll for this certificate**.

m. In the **Certificate Properties** dialog box, on the **Subject** tab, under **Subject name**, under **Type**, click **Common name**.

n. In the **Value** text box, type **131.107.0.2**, and then click **Add**.

o. Click **OK**, click **Enroll**, and then click **Finish**.

p. In the Certificates snap-in details pane, verify that a new certificate with the name **131.107.0.2** displays with the **Intended Purposes** of **Server Authentication**.

q. Right-click the new certificate with the name **131.107.0.2**, and then click **Properties**.

r. In the **Properties** dialog box, in the **Friendly Name** text box, type **IP-HTTPS Certificate**, and then click **OK**.

s. Close the console window. If you are prompted to save settings, click **No**.

2. Create a certificate revocation list (CRL) distribution point on LON-RTR by performing the following steps:

   a. Switch to Server Manager.

   b. In Server Manager, click **Tools**, and then click **Internet Information Services (IIS) Manager**.

   c. If the Internet Information Service Manager message box displays, click **No**.

   d. In the console tree, expand **LON-RTR**, expand **Sites**, click and then right-click **Default Web Site**, and then click **Add Virtual Directory**.

   e. In the **Add Virtual Directory** dialog box, in the **Alias** text box, type **CRLD**. Next to **Physical path**, click the ellipsis (…) button.

   f. In the **Browse for Folder** dialog box, click **Local Disk (C:)**, and then click **Make New Folder**.

   g. Type **CRLDist**, and then press Enter.

   h. In the **Browse for Folder** dialog box, click **OK**.

   i. In the **Add Virtual Directory** dialog box, click **OK**.

   j. In the middle pane of the console, double-click **Directory Browsing**, and in the Actions pane, click **Enable**.

   k. In the console, click the **CRLD** folder.

   l. In the middle pane of the console, double-click the **Configuration Editor** icon.

   m. In the Configuration Editor, in the **Section** drop-down list box, expand **system.webServer**, expand **security**, and then click **requestFiltering**.

   n. In the middle pane of the requestFiltering console, double-click **allowDoubleEscaping** to change the value from False to **True**.

   o. In the actions pane, click **Apply**.

   p. Close Internet Information Services (IIS) Manager.

Question: Why did you make the CRL available on the edge server?

Answer: You made the CRL available on the edge server so that the Internet DirectAccess clients can access the CRL.

3. Share and secure permissions to the CRL distribution point by performing the following steps:

   a. On the desktop, on the taskbar, click the **File Explorer** icon.

   b. In File Explorer, click **This PC**, and then double-click **Local Disk (C:)**.

   c. In the File Explorer details pane, right-click the **CRLDist** folder, and then click **Properties**.

   d. In the **CRLDist Properties** dialog box, click the **Sharing** tab, and then click **Advanced Sharing**.

   e. In the **Advanced Sharing** dialog box, click **Share this folder**.

   f. In the **Share name** text box, add a dollar sign (**$**) to the end of the name so that the share name is now **CRLDist$**.

   g. In the **Advanced Sharing** dialog box, click **Permissions**.

   h. In the **Permissions for CRLDist$** dialog box, click **Add**.

   i. In the **Select Users, Computers, Service Accounts, or Groups** dialog box, click **Object Types**.

   j. In the **Object Types** dialog box, click **Computers**, and then click **OK**.

k. In the **Select Users, Computers, Service Accounts, or Groups** dialog box, in the **Enter the object names to select** text box, type **LON-DC1**, click **Check Names**, and then click **OK**.

l. In the **Permissions for CRLDist$** dialog box, in the **Group or user names** list, click **LON-DC1 (ADATUM\NYC-DC1$)**.

m. In the Permissions for LON-DC1 area, under **Full control**, click **Allow**, and then click **OK**.

n. In the **Advanced Sharing** dialog box, click **OK**.

o. In the **CRLDist Properties** dialog box, click the **Security** tab.

p. On the **Security** tab, click **Edit**.

q. In the **Permissions for CRLDist** dialog box, click **Add**.

r. In the **Select Users, Computers, Service Accounts, or Groups** dialog box, click **Object Types**.

s. In the **Object Types** dialog box, click **Computers**, and then click **OK**.

t. In the **Select Users, Computers, Service Accounts, or Groups** dialog box, in the **Enter the object names to select** text box, type **LON-DC1**, click **Check Names**, and then click **OK**.

u. In the **Permissions for CRLDist** dialog box, in the **Group or user names** list, click **LON-DC1 (ADATUM\LON-DC1$)**.

v. In the Permissions for LON-DC1 area, under **Full control**, click **Allow**, and then click **OK**.

w. In the **CRLDist Properties** dialog box, click **Close**.

x. Close File Explorer.

4. Publish the CRL to LON-RTR by performing the following steps:

Note: These steps make the CRL available on the edge server for Internet-based DirectAccess clients.

a. Switch to LON-DC1.

b. In Server Manager, click **Tools**, and then click **Certification Authority**.

c. In the Certification Authority console, expand **AdatumCA**, right-click **Revoked Certificates**, point to **All Tasks**, and then click **Publish**.

d. In the **Publish CRL** dialog box, click **New CRL**, and then click **OK**.

e. On the desktop, on the taskbar, click the **File Explorer** icon.

f. In the File Explorer address bar, type **\\LON-RTR\CRLDist$**, and then press Enter.

g. In the File Explorer window, verify that the **AdatumCA** files display.

h. Close File Explorer.

5. Complete the DirectAccess Getting Started Wizard on LON-RTR by performing the following steps:

Note: These steps configure LON-RTR as a DirectAccess server.

a. Restart LON-RTR and then sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.

Note: During the following three or four steps, the Enable DirectAccess Wizard might launch. If this happens, click **Cancel**, and proceed with the remaining instructions.

b. On LON-RTR, in Server Manager, click **Tools**, and then click **Routing and Remote Access**.

c. In Routing and Remote Access, right-click **LON-RTR (local)**, and then click **Disable Routing and Remote Access**.

d. In the Routing and Remote Access message box, click **Yes**.

e. Close the Routing and Remote Access console.

f. In Server Manager, on the **Tools** menu, click **Remote Access Management**.

g. In the Remote Access Management console, click **DirectAccess and VPN**.

h. In the results pane, click **Run the Getting Started Wizard**.

i. In the Configure Remote Access Wizard, click **Deploy DirectAccess only**.

j. In **Network Topology**, verify that **Edge** is selected, verify that **131.107.0.2** is the public name used by clients to connect to the remote access server, and then click **Next**.

k. On the **Configure Remote Access** page, click **Finish**.

l. When the configuration completes, click **Close**.

m. In the Remote Access Management console, under **Step 1**, click **Edit**, and then click **Next**.

n. Under **Select one or more security groups**, in the details pane, click **Add**.

o. In the **Select Groups** dialog box, type **DA_Clients**, and then click **OK**.

p. In the groups list, click **Domain Computers**, and then click **Remove**.

q. Click **Next**, and then click **Finish**.

r. In the Remote Access Management console, under **Step 2**, click **Edit**.

s. On the **Network Topology** page, verify that **Edge** is selected, type **131.107.0.2**, and then click **Next**.

t. On the **Network Adapters** page, verify that **CN=131.107.0.2** is used as a certificate to authenticate IP-HTTPS connections, and then click **Next**.

u. On the **Authentication** page, click **Use computer certificates**. Click **Browse**, click **AdatumCA**, click **OK**, and then click **Finish**.

v. In the Remote Access Setup pane, under **Step 3**, click **Edit**.

w. On the **Network Location Server** page, click **The network location server is deployed on a remote web server (recommended)**. In the URL of the network location server, type **https://nls.adatum.com**, and then click **Validate**.

x. Ensure that the URL validates, and then click **Next**.

y. Examine the values, and then click **Next**.

z. In the **DNS Suffix Search List**, click **Next**.

aa. On the **Management** page, click **Finish**.

bb. Under **Step 4**, click **Edit**.

cc. On the **DirectAccess Application Server Setup** page, click **Finish**.

dd. Click **Finish** to apply the changes.

   ee. In the **Remote Access Review** dialog box, click **Apply**.

   ff. Under **Applying Remote Access Setup Wizard Settings**, click **Close**.

6. On LON-RTR, update Group Policy settings by performing the following steps:

   a. Move the mouse pointer to the lower-right corner, and then click **Search**.

   b. In the **Search** box, type **cmd**, and then press Enter.

   c. At the command prompt, type the following commands, pressing Enter at the end of each line:

```
gpupdate /force
Ipconfig
```

7. Verify that the LON-RTR has an IPv6 address for **Tunnel adapter IPHTTPSInterface** that begins with **2002**.

▶ Task 9: Configure DirectAccess Group Policy Object (GPO) settings

1. Switch to LON-DC1.
2. In Server Manager, click **Tools**, and then click **Group Policy Management**.
3. In the Group Policy Management Console, in the navigation pane, expand **Adatum.com**, and then click **DirectAccess Client Settings**.
4. In the details pane, under **WMI Filtering**, click **DirectAccess – Laptop only WMI filter**, and then click **<none>**.
5. In the **Group Policy Management** dialog box, click **Yes**.
6. Close the Group Policy Management Editor and the Group Policy Management Console.
7. Start LON-CL1, and sign in as **Adatum\Administrator** with the password of **Pa$$w0rd**.

Note: This step ensures that LON-CL1 connects to the domain as a member of the DA_Clients security group.

8. On the **Start** screen, type **cmd**, and then press Enter.
9. At the command prompt, type the following command, and then press Enter:

```
gpupdate /force
```

10. At the command prompt, type the following command, and then press Enter:

```
gpresult /R
```

11. Verify that in the list of the Applied Policy objects, under **Computer Settings**, that the **DirectAccess Client Settings** GPO displays.

Note: If the policy is not being applied, run the **gpupdate /force** command again. If the policy is still not being applied, restart the computer. After the computer restarts, sign in as **Adatum\Administrator** and run the **gpresult /R** command again.

▶ Task 10: Verify certificate distribution

1. At the command prompt, type **mmc.exe**, and then press Enter.
2. In the MMC console, click **File**, and then click **Add/Remove Snap-in**.
3. In the **Add/Remove Snap-in** dialog box, click **Certificates**, click **Add**, select **Computer account**, and then click **Next**.
4. Select **Local computer**, click **Finish**, and then click **OK**.
5. In the Certificates snap-in console, click **Certificates (Local Computer)**, expand **Personal**, and then click **Certificates**.
6. In the Certificates details pane, verify that a certificate with the name **LON-CL1.adatum.com** displays with the **Intended Purposes** of **Client Authentication** and **Server Authentication**.
7. Close the console. When you are prompted to save settings, click **No**.

▶ Task 11: Verify IP configuration

1. On LON-CL1, switch to the Desktop.
2. On the desktop, on the taskbar, click the **Internet Explorer** icon.
3. In the Windows Internet Explorer®, in the Address bar, type **http://lon-svr1.adatum.com/**, and then press Enter.
4. Verify that the default IIS 8 web page for LON-SVR1 displays.
5. In the Internet Explorer Address bar, type **https://nls.adatum.com/**, and then press Enter.
6. Verify that the default IIS 8 web page for LON-SVR1 displays.
7. On the taskbar, click the **File Explorer** icon.
8. In the File Explorer address bar, type **\\Lon-SVR1\Files**, and then press Enter.
9. Verify that the **Files** shared folder contents display.
10. Close all open windows except the Command Prompt window.

▶ Task 12: Move LON-CL1, and verify connectivity to intranet resources

1. On LON-CL1, move the mouse pointer to the lower-right corner of the screen, click **Settings**, and then click **Control Panel**.
2. In Control Panel, click **Network and Internet**.
3. In the Network and Internet window, click **Network and Sharing Center**.
4. In the Network and Sharing Center window, click **Change adapter settings**.
5. Right-click **Ethernet**, and then click **Properties**.
6. In the **Ethernet Properties** dialog box, double-click **Internet Protocol Version 4 (TCP/IPv4)**.
7. In the **Internet Protocol Version 4 (TCP/IPv4) Properties** dialog box, click **Use the following IP address**.
8. Provide the following settings, and then click **OK**:
   - IP address: **131.107.0.10**
   - Subnet mask: **255.255.0.0**
   - Default gateway: **131.107.0.2**

9. In the **Ethernet Properties** dialog box, click **OK**.
10. In the Network Connections window, right-click **Ethernet**, and then click **Disable**.
11. In the Network Connections window, right-click **Ethernet**, and then click **Enable**.
12. On your host computer, in Microsoft® Hyper-V Manager®, right-click **20413C-LON-CL1**, and then click **Settings**.
13. Click **Network Adapter**, in the **Virtual Switch** drop-down list box, click **Private Network 2** network, and then click **OK**.
14. On LON-CL1, switch to the desktop, and then click the **Internet Explorer** icon.
15. In the Internet Explorer Address bar, type **http://lon-svr1.adatum.com**, and then press Enter.
16. Verify that the default IIS 8 web page for LON-SVR1 displays.
17. Leave the Internet Explorer window open.
18. On the taskbar, click the **File Explorer** icon.
19. In the File Explorer address bar, type **\\LON-SVR1\Files**, and then press Enter.
20. Verify that a folder window with the contents of the Files shared folder displays.
21. Switch to the Command Prompt window.
22. At the command prompt, type the following command, and then press Enter:

    ```
    ping lon-dc1.adatum.com
    ```

    Verify that you are receiving replies from lon-dc1.adatum.com.

23. Close all open windows.
24. Switch to LON-RTR.
25. Switch to the Remote Access Management console.
26. In the console, click **Remote Client Status**.

Note: Notice that LON-CL1 is connected through IP-HTTPS. In the Connection Details pane, in the bottom-right of the screen, note the use of Kerberos authentication for the machine and the user.

27. Close all open windows.

▶ Task 13: Revert virtual machines

When you are finished with this portion of the lab, before you continue, you must revert all virtual machines to their initial state. To do this, perform the following steps:

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-CL1**, and then click **Revert**.
3. In the **Revert Virtual Machines** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20413C-LON-DC1**, **20413C-LON-SVR1**, and **20413C-LON-RTR**.
5. In the **Virtual Machines** list, right-click **20413C-LON-DC1**, and then click **Connect**.
6. Click **Start**.

7. Log on to LON-DC1 as **Adatum\Administrator** with the password of **Pa$$w0rd**.

8. Repeat steps 5 through 7 for **20413C-LON-SVR1**, **20413C-LON-RTR**, **20413C-LON-CL1**, and **20413C-LON-CL2**.

Results: After completing this exercise, you should have planned and implemented a DirectAccess solution.

Exercise 3: Planning and Implementing a VPN Solution

▶ Task 1: Read the supporting documentation

- Read the documentation provided in the Student Handbook.

▶ Task 2: Update the proposal document with your planned course of action

- Answer the questions in the proposals section of the VPN Strategy document.

Proposals

1. What tunneling protocols will you use?

   To provide the best level of security, you should use either Secure Socket Tunneling Protocol (SSTP) or Layer Two Tunneling Protocol (L2TP). L2TP may provide slightly better security for authentication because both computers and users are authenticated. However, in some cases firewalls may block L2TP VPNs.

   SSTP has similar encryption strength compared to L2TP/IPsec, but it is easier to configure because it requires no computer authentication. In addition, firewalls almost never block SSTP.

2. What authentication or encryption methods do you recommend?

   To obtain the highest level of security, you should use smart cards. All VPN tunneling protocols that are supported by Windows Server 2012 and Windows 8 support the use of Extensible Authentication Protocol – Transport Layer Security (EAP-TLS), which is the authentication method that smart card authentication uses.

3. How many network policies do you envision?

   The following policies are necessary:

   - A single policy for all executives
   - A single policy for Customer Service staff
   - A policy on each group of branch management staff for each regional hub

4. List the network policies and their characteristics.

   The following are the required network policies and their characteristics:

   - Executives. A single network policy with no restrictions for executives.
   - Branch management staff. A network policy for branch management staff at each regional hub site. The policy for each regional hub site will restrict access by using IP filters.
   - Customer service staff. A single network policy for customer service staff that denies remote access.
   - Marketing staff. The marketing staff does not require access to applications or data. In addition, they can be given web-based access to their email instead of using Remote Desktop Services. You can secure web-based access to email with Secure Sockets Layer (SSL). This simplifies client configuration for the marketing staff.

5. In what order will these policies process?

   Only the first network policy with matching conditions is evaluated. Therefore, you must be sure that the appropriate policy is evaluated first, based on the conditions that you have at the organization. Typically, the largest concern is group memberships that overlap. For example, if an executive is a member of both the Executives group and the Customer Service group, then you must ensure that the Executives network policy that allows access is evaluated prior to the Customer Service network policy that denies access.

6. What certificates are required?

   Certificates for each VPN server and each VPN client are required.

▶ Task 3: Examine the suggested proposals in the Lab Answer Key

- Compare your proposals with the ones shown previously.

▶ Task 4: Discuss your proposed solution with the class, as guided by your instructor

- Be prepared to discuss your proposals with the class.

▶ Task 5: Install and configure the Remote Access role

1. Switch to LON-RTR.
2. Sign in as **Adatum\Administrator** with the password **Pa$$w0rd**.
3. In Server Manager, in the details pane, click **Add roles and features**.
4. In the Add Roles and Features Wizard, click **Next**.
5. On the **Select installation type** page, click **Role-based or feature based installation**, and then click **Next**.
6. On the **Select destination server** page, click **Next**.
7. On the **Select server roles** page, select the **Network Policy and Access Services** check box.
8. Click **Add Features**, and then click **Next** twice.
9. On the **Network Policy and Access Services** page, click **Next**.
10. On the **Select role services** page, verify that the **Network Policy Server** check box is selected, and then click **Next**.
11. On the **Confirm installation selections** page, click **Install**.
12. Verify that the installation is successful, and then click **Close**.
13. In Server Manager, click **Tools**, and then click **Network Policy Server**.
14. In Network Policy Manager, in the navigation pane, right-click **NPS (Local)**, and then click **Register server in Active Directory**.
15. In the **Network Policy Server** message box, click **OK**.
16. In the **Network Policy Server** dialog box, click **OK**.
17. Leave the Network Policy Manager open.
18. In Server Manager, click **Tools**, and then click **Routing and Remote Access**.
19. In Routing and Remote Access, right-click **LON-RTR**, and then click **Disable Routing And Remote Access**. Click **Yes**.
20. In Routing and Remote Access, right-click **LON-RTR**, and then click **Configure and Enable routing and remote access**.
21. In the Routing and Remote Access Server Setup Wizard, click **Next**.
22. On the Configuration page, select **Remote access (dial-up or VPN)** and then click **Next**.
23. On the Remote Access page, click **VPN**, and then click **Next**.

24. On the **VPN Connection** page, select **Ethernet 2**, clear the **Enable security** check box, and then click **Next**.
25. On the IP Address Assignment page, click **Next**.
26. On the **Managing Multiple Remote Access Servers** page, click **Next**.
27. Click **Finish**, and then click **OK** twice.

▶ Task 6: Create the required network policies

1. Switch to LON-RTR.
2. Open the Network Policy Server console.
3. In the Network Policy Server console, expand **Policies**, and then click **Network Policies**.
4. In the details pane, right-click the policy at the top of the list, and then click **Disable**.
5. In the details pane, right-click the policy at the bottom of the list, and then click **Disable**.
6. In the navigation pane, right-click **Network Policies**, and then click **New**.
7. In the New Network Policy Wizard, in the **Policy name** text box, type **Adatum VPN Policy**.
8. In the **Type of network access server** list, click **Remote Access Server (VPN-Dial up)**, and then click **Next**.
9. On the **Specify Conditions** page, click **Add**.
10. In the **Select condition** dialog box, click **NAS Port Type**, and then click **Add**.
11. In the **NAS Port Type** dialog box, select the **Virtual (VPN)** check box, click **OK**, and then click **Next**.
12. On the **Specify Access Permission** page, click **Access granted**, and then click **Next**.
13. On the **Configure Authentication Methods** page, click **Next**.
14. On the **Configure Constraints** page, click **Next**.
15. On the **Configure Settings** page, click **Next**.
16. On the **Completing New Network Policy** page, click **Finish**.

▶ Task 7: Create a client VPN

1. Switch to LON-CL2.
2. Sign in as **Adatum\Administrator** with the password of **Pa$$w0rd**.
3. On the **Start** screen, type **Control Panel**, and then in the **Apps** list, click **Control Panel**.
4. In Control Panel, click **Network and Internet**.
5. In the Network and Internet window, click **Network and Sharing Center**.
6. In the Network and Sharing Center window, under **Change your Network Settings**, click **Set up a new connection or network**.
7. On the **Choose a connection option** page, click **Connect to a workplace**, and then click **Next**.
8. On the **How do you want to connect** page, click **Use my Internet connection (VPN)**.
9. Click **I'll set up an Internet connection later**.
10. On the **Type the Internet address to connect to** page, in the **Internet address** text box, type **10.10.0.1**.
11. In the **Destination name** text box, type **Adatum VPN**.

12. Select the **Allow other people to use this connection** check box, and then click **Create**.
13. In the Network and Sharing Center window, click **Change adapter settings**.
14. Right-click the **Adatum VPN** connection, and then click **Properties**.
15. In the **Adatum VPN Properties** dialog box, click the **Security** tab.
16. Under **Authentication**, click **Allow these protocols**, and then click **OK**.

▶ Task 8: Test VPN access

1. In the Network Connections window, right-click the **Adatum VPN** connection, and then click **Connect/Disconnect**.
2. In the **Networks** list on the right, click **Adatum VPN**, and then click **Connect**.
3. In the Network Authentication window, in the **User name** text box, type **Adatum\Administrator**. In the **Password** text box, type **Pa$$w0rd**, and then click **OK**.
4. Wait for the VPN connection to be made, and verify that your connection is successful.

Results: After completing this exercise, you should have planned and implemented a VPN solution.

Exercise 4: Implementing Web Application Proxy

▶ Task 1: Install the Active Directory Federation Services (AD FS) role

1. On LON-DC1, on the desktop, on the taskbar, click the **Windows PowerShell®** icon.
2. At the Windows PowerShell command prompt, run the following command, and then press Enter:

   ```
   Add-KdsRootKey –EffectiveImmediately
   ```

3. Close the Windows PowerShell window.
4. In Server Manager, click **Manage**, and then click **Add Roles and Features**.
5. On the **Before you begin** page, click **Next**.
6. On the **Select installation type** page, click **Next**.
7. On the **Select destination server** page, click **Next**.
8. On the **Select server roles** page, click **Active Directory Federation Services**, and then click **Next**.
9. On the **Select features** page, click **Next**.
10. On the **Active Directory Federation Services** page, click **Next**.
11. On the **Confirm installation selections** page, click **Install**, and then wait for the installation to finish.
12. When the installation completes, click **Close**.
13. In Server Manager, click the yellow notification icon, and then click the **Configure the federation service on this server** link.
14. On the **Welcome** page, ensure that **Create the first federation server in a federation server farm** is selected, and then click **Next**.
15. On the **Connect to Active Directory Domain Services**, click **Next** to use the **ADATUM\Administrator** account.
16. On the **Specify Service Properties** page, under SSL certificate, click **LON-DC1.Adatum.com**. In the **Federation Service Display Name** text box, type **LON-DC1.Adatum.com**, and then click **Next**.
17. On the **Specify Service Account** page, click **Create a Group Managed Service Account**.
18. In the **Account Name** text box, type **ADFS**, and then click **Next**.
19. On the **Specify Configuration Database** page, ensure that **Create a database on this server using Windows Internal Database** is selected, and then click **Next**.
20. On the **Review Options** page, verify that the correct configuration settings are listed, and then click **Next**.
21. On the **Pre-requisite Checks** page, click **Configure**.
22. Wait for the configuration to finish (note that a service principal name registration error may occur), and then click **Close**.
23. On Server Manager, click **Tools**, and then click **Windows PowerShell**.
24. At the Windows PowerShell command prompt, type the following command, and then press Enter:

    ```
    Set-ADFSProperties  –AutoCertificateRollover $False
    ```

25. You must perform this step so that you can modify the certificates that AD FS uses.

26. Close the Windows PowerShell window.
27. In Server Manager, click **Tools**, and then click **AD FS Management**.
28. In the AD FS console, in the left pane, expand **Service**, and then click **Certificates**.
29. Right-click **Certificates**, and then click **Add Token-Signing Certificate**.
30. In the **Select a token-signing certificate** dialog box, click the **LON-DC1.Adatum.com** certificate, and then click **Click here to view certificate properties**.
31. Verify that the certificate purposes include **Proves your identity to a remote computer** and **Ensures the identity of a remote computer**, and then click **OK**.
32. Click **OK** to close the **Windows Security** dialog box.
33. When the **AD FS Management** warning dialog box displays, click **OK**.

Note: Notice that the certificate has a subject of **CN=LON-DC1.Adatum.com**. If no name displays under the **Subject**, or the subject column displays a value other than **CN=LON-DC1.adatum.com** when you add the certificate, delete the certificate, and then add the next certificate in the list.

34. Under **Token-signing**, right-click the newly added certificate, and then click **Set as Primary**. Review the warning message, and then click **Yes**.
35. Select the certificate that has just been superseded, right-click the certificate, and then click **Delete**. Click **Yes** to confirm the deletion.

▶ Task 2: Install the Web Application Proxy role

1. Switch to LON-RTR.
2. On the **Start** screen, click **Server Manager**.
3. In Server manager, on the **Dashboard** page, click **Add roles and features**.
4. In the Add Roles and Features Wizard, click **Next** three times.
5. On the **Select server roles** page, expand **Remote Access**, click **Web Application Proxy**, 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, verify that the installation is successful, and then click **Close**.

▶ Task 3: Configure access to an internal website

1. On LON-RTR, on the **Start** screen, type **cmd**, and then press Enter.
2. At the command prompt, type **mmc**, and then press Enter.
3. In the MMC console, on the **File** menu, click **Add or Remove Snap-In**.
4. In the **Add or Remove Snap-ins** dialog box, click **Certificates**, click **Add**, click **Computer account**, and then click **Next**.
5. Verify that **Local Computer** is selected, click **Finish**, and then click **OK**.
6. Expand **Certificates (Local Computer)** ,then right-click **Personal**, click **All Tasks**, and then click **Request new Certificate**.

7. On the **Before You Begin** page, click **Next**.
8. On the **Select Certificate Enrollment Policy** page, click **Next**.
9. Click **Adatum Web Server**, and then click **More information is required to enroll for this certificate. Click here to configure settings.**
10. In the **Subject Name** group, in the **Type** drop-down list box, click **Common Name**, then in the **Value** text box, type **lon-dc1.adatum.com**, and then click **Add**.
11. In the **Alternative name** group, in the **Type** drop-down list box, click **DNS**, then in the **Value** text box, type **lon-dc1.adatum.com**, and then click **Add**.
12. In the **Alternative name** group, in the **Type** drop-down list box, click **DNS**, then in the **Value** text box, type **enterpriseregistration.adatum.com**, and then click **Add**.
13. In the **Alternative name** group, in the **Type** drop-down list box, click **DNS**, in the **Value** text box, type **lon-svr1.adatum.com**, and then click **Add**.
14. In the **Certificate Properties** dialog box, click **OK**, and then click **Enroll**.
15. Click **Finish** to close the **Certificate Enrollment** dialog box.
16. Switch to LON-SVR1.
17. On the **Start** screen, type **mmc**, and then press Enter.
18. In the MMC console, on the **File** menu, click **Add or Remove Snap-In**.
19. In the **Add or Remove Snap-ins** dialog box, click **Certificates**, click **Add**, click **Computer account**, and then click **Next**.
20. Verify that **Local Computer** is selected, click **Finish**, and then click **OK**.
21. In the MMC console, in the left pane, expand **Certificates (local Computer)**, right-click **Personal**, click **All Tasks**, and then click **Request new Certificate**.
22. On the **Before You Begin** page, click **Next**.
23. On the **Select Certificate Enrollment Policy** page, click **Next**.
24. Click **Adatum Web Server**, and then click **More information is required to enroll for this certificate. Click here to configure settings.**
25. In the **Subject Name** group, in the **Type** drop-down list box, click **Common Name**, and in the **Value** text box, type **lon-svr1.adatum.com**, and then click **Add**.
26. Click **OK** to close the **Certificate Properties** dialog box.
27. Click **Enroll** to proceed with certificate enrollment.
28. Click **Finish** to close the **Certificate Enrollment** dialog box.
29. In Server Manager, on the **Tools** menu, click **Internet Information Services (IIS) Manager**.
30. In the Internet Information Services (IIS) Manager console tree, navigate to **LON-SVR1/Sites**, and then click **Default Web site**.
31. In the Actions pane, click **Bindings**, select **https**, and then click **Edit**.
32. In the **Edit Site Bindings** dialog box, in the **Host name** text box, type **lon-svr1.adatum.com**. In the **SSL Certificate** drop-down list box, click the **lon-svr1.adatum.com** certificate, click **OK**, and then click **Close**.
33. Close the Internet Information Services (IIS) console.
34. Switch to LON-RTR.

35. In Server Manager, on the **Tools** menu, click **Remote Access Management**.
36. In the Remote Access Management console, in the navigation pane, click **Web Application Proxy**.
37. In the middle pane, click **Run the Web Application Proxy Configuration Wizard**.
38. In the Web Application Proxy Configuration Wizard, on the **Welcome** page, click **Next**.
39. On the **Federation Server** page, perform the following steps, and then click **Next**:
    a. In the **Federation service name** text box, type **lon-dc1.adatum.com**.
    b. In the **User name** text box, type **Administrator**, in the **Password** text box, type **Pa$$w0rd**.
40. On the **AD FS Proxy Certificate** page, in the list of certificates currently installed on the Web Application Proxy server, click **lon-dc1.adatum.com**, and then click **Next**.

Note: This is the certificate that Web Application Proxy will use for AD FS proxy functionality.

41. On the **Confirmation** page, review the settings, and then click **Configure**.

Note: If required, you can copy or save the Windows PowerShell command to automate additional installations.

42. On the **Results** page, verify that the configuration is successful, and then click **Close**.
43. On the Web Application Proxy server, in the Remote Access Management console, in the navigation pane, click **Web Application Proxy**, and then in the tasks pane, click **Publish**.
44. In the Publish New Application Wizard, on the **Welcome** page, click **Next**.
45. On the **Preauthentication** page, click **Pass-through**, and then click **Next**.
46. On the **Publishing Settings** page, perform following steps, and then click **Next**:
    a. In the **Name** text box, type **LON-SVR1 Web**.
    b. In the **External URL** text box, type https://lon-svr1.adatum.com.
    c. In the **External certificate** list, click **lon-dc1.adatum.com**.
    d. In the **Backend server URL** text box, ensure that **https://lon-svr1.adatum.com** displays.

Note: Note that this value is entered automatically when you enter the external URL.

47. On the **Confirmation** page, review the settings, and then click **Publish**.

Note: You can copy or save the Windows PowerShell command to set up additional published applications.

48. On the **Results** page, ensure that the application published successfully, and then click **Close**.
49. Switch to LON-SVR1.
50. In Server Manager, on the **Tools** menu, click **Internet Information Services (IIS) Manager**.

51. In the Internet Information Services (IIS) Manager console, expand **LON-SVR1 (ADATUM\Administrator)**.

52. In the Internet Information Services (IIS) Manager, in the console tree, navigate to **Sites**, and then click **Default Web site**.

53. In the Default Web Site Home pane, double-click **Authentication**.

54. In the **Authentication** pane, right-click **Windows Authentication**, and then click **Enable**. Right-click **Anonymous Authentication**, and then click **Disable**.

55. Close the Internet Information Services (IIS) Manager console.

▶ Task 4: Verify access to the internal web site

1. Switch to LON-CL1.

2. On the **Start** screen, type **Control Panel**, and then press Enter.

3. In Control Panel, click **System and Security**, click **System** under **Computer name, domain and workgroup settings**, and then click **Change Settings**.

4. In the **System Properties** dialog box, click **Change**.

5. In the **Computer Name/Domain Changes** dialog box, click **Workgroup**, type **WORKGROUP**, and then click **OK**.

6. In the **Computer Name/Domain Changes** dialog box, click **OK**.

7. If the **Windows Security** dialog box displays, in the **User name** text box, type **Administrator**, in the **Password** text box, type **Pa$$w0rd**, and then click **OK**.

8. In the **Welcome to the WORKGROUP workgroup** dialog box, click **OK**.

9. To restart the computer, click **OK**.

10. To close the **System Properties** dialog box, click **Close**.

11. Click **Restart Now**.

12. On LON-CL1, sign in with user name **Admin** and password **Pa$$w0rd**.

13. On the **Start** screen, type **notepad**, and then click **Notepad**.

14. In the Notepad window, type **131.107.0.2 lon-svr1.adatum.com**.

15. From the **File** menu, click **Save As**.

16. In the **Save As** dialog box, navigate to **Documents**.

17. In the **Save as type** list, click **All files (*.*)**.

18. In the **File name** text box, type **Hosts**, and then click **Save**.

19. On the desktop, on the taskbar, click the **File Explorer** icon.

20. In File Explorer, open the **Documents** folder, right-click the **Hosts** file, and then click **Copy**.

21. In the navigation pane, expand drive **C:**, expand **Windows**, expand **System32**, expand **drivers**, double-click **etc**, and then **Paste** the copied **Hosts** file into the **etc** folder.

22. In the **Replace or Skip Files** dialog box, click **Replace the file in the destination**.

23. In the **Destination Folder Access Denied** dialog box, click **Continue**.

24. On the **Start** screen, click the **Internet Explorer** tile.

25. In the Internet Explorer Address bar, type **https://lon-svr1.adatum.com**, and then press Enter.

26. If Internet Explorer displays a page stating that there is a problem with the certificate used by the page, click **Continue to this website (not recommended)**.
27. In the **Internet Explorer** dialog box, type **Adatum\Bill** for the user name and **Pa$$w0rd** for password, and then click **OK**.
28. Verify that the default IIS 8.0 web page for LON-SVR1 opens.
29. If you are unable to connect to https://lon-svr1.adatum.com, perform the following steps:

    a. On LON-CL1, on the **Start** screen, type **cmd**, and then press Enter.

    b. At the command prompt, type the following command, and then press Enter:

    ```
    regedit
    ```

    c. In the **User Account Control** dialog box, click **Yes**.

    d. In the Registry Editor window, in the navigation pane, expand **HKLM**, expand **Software**, expand **Policies**, expand **Microsoft**, expand **Windows NT**, expand **DNSClient**, and then expand **DNSPolicyConfig**.

Note: Notice the three entries starting with DA.

    e. In the Registry Editor window, in the navigation pane, right-click each of the three entries starting with **DA**, click **Delete**, and in the **Confirm Key Delete** dialog box, click **Yes**.

    f. Close the Registry Editor window.

    g. Restart LON-CL1 and perform steps 24 through 28 to verify connectivity to default IIS 8.0 web page.

Results: After completing this exercise, you should have implemented a Web Application Proxy solution.

▶ Task: To prepare for the next module

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

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20413C-LON-CL1**, and then click **Revert**.
3. In the **Revert Virtual Machines** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for the following machines: 20413C-LON-DC1, 20413C-LON-CL2, 20413C-LON-RTR, and 20413C-LON-SVR1.