

OFFICIAL MICROSOFT LEARNING PRODUCT

20412A

Configuring Advanced Windows Server®
2012 Services

MCT USE ONLY. STUDENT USE PROHIBITED

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

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

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

© 2012 Microsoft Corporation. All rights reserved.

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

Product Number: 20412A

Part Number: X18-48644

Released: 09/2012

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

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

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

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

1. DEFINITIONS.

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

- j. "Microsoft Partner Network Member" or "MPN Member" means a silver or gold-level Microsoft Partner Network program member in good standing.
- k. "Personal Device" means one (1) device, workstation or other digital electronic device that you personally own or control that meets or exceeds the hardware level specified for the particular MOC Course.
- l. "Private Training Session" means the instructor-led training classes provided by MPN Members for corporate customers to teach a predefined learning objective. These classes are not advertised or promoted to the general public and class attendance is restricted to individuals employed by or contracted by the corporate customer.
- m. "Trainer Content" means the trainer version of the MOC Course and additional content designated solely for trainers to use to teach a training session using a MOC Course. Trainer Content may include Microsoft PowerPoint presentations, instructor notes, lab setup guide, demonstration guides, beta feedback form and trainer preparation guide for the MOC Course. To clarify, Trainer Content does not include virtual hard disks or virtual machines.

2. **INSTALLATION AND USE RIGHTS.** The Licensed Content is licensed not sold. The Licensed Content is licensed on a one copy per user basis, such that you must acquire a license for each individual that accesses or uses the Licensed Content.

2.1 Below are four separate sets of installation and use rights. Only one set of rights apply to you.

a. **If you are a Authorized Learning Center:**

- i. If the Licensed Content is in digital format for each license you acquire you may either:
 - 1. install one (1) copy of the Licensed Content in the form provided to you on a dedicated, secure server located on your premises where the Authorized Training Session is held for access and use by one (1) End User attending the Authorized Training Session, or by one (1) MCT teaching the Authorized Training Session, **or**
 - 2. install one (1) copy of the Licensed Content in the form provided to you on one (1) Classroom Device for access and use by one (1) End User attending the Authorized Training Session, or by one (1) MCT teaching the Authorized Training Session.
- ii. You agree that:
 - 1. you will acquire a license for each End User and MCT that accesses the Licensed Content,
 - 2. each End User and MCT will be presented with a copy of this agreement and each individual will agree that their use of the Licensed Content will be subject to these license terms prior to their accessing the Licensed Content. Each individual will be required to denote their acceptance of the EULA in a manner that is enforceable under local law prior to their accessing the Licensed Content,
 - 3. for all Authorized Training Sessions, you will only use qualified MCTs who hold the applicable competency to teach the particular MOC Course that is the subject of the training session,
 - 4. you will not alter or remove any copyright or other protective notices contained in the Licensed Content,

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

b. If you are a MPN Member.

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

c. If you are an End User:

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

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

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

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

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

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

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

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

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

survive this agreement.

- c. **Term.** If you are an Authorized Training Center, MCT or MPN, you agree to cease using all copies of the beta version of the Licensed Content upon (i) the date which Microsoft informs you is the end date for using the beta version, or (ii) sixty (60) days after the commercial release of the Licensed Content, whichever is earliest (“**beta term**”). Upon expiration or termination of the beta term, you will irretrievably delete and destroy all copies of same in the possession or under your control.
4. **INTERNET-BASED SERVICES.** Microsoft may provide Internet-based services with the Licensed Content, which may change or be canceled at any time.
 - a. **Consent for Internet-Based Services.** The Licensed Content may connect to computer systems over an Internet-based wireless network. In some cases, you will not receive a separate notice when they connect. Using the Licensed Content operates as your consent to the transmission of standard device information (including but not limited to technical information about your device, system and application software, and peripherals) for internet-based services.
 - b. **Misuse of Internet-based Services.** You may not use any Internet-based service in any way that could harm it or impair anyone else’s use of it. You may not use the service to try to gain unauthorized access to any service, data, account or network by any means.
 5. **SCOPE OF LICENSE.** The Licensed Content is licensed, not sold. This agreement only gives you some rights to use the Licensed Content. Microsoft reserves all other rights. Unless applicable law gives you more rights despite this limitation, you may use the Licensed Content only as expressly permitted in this agreement. In doing so, you must comply with any technical limitations in the Licensed Content that only allows you to use it in certain ways. Except as expressly permitted in this agreement, you may not:
 - install more copies of the Licensed Content on devices than the number of licenses you acquired;
 - allow more individuals to access the Licensed Content than the number of licenses you acquired;
 - publicly display, or make the Licensed Content available for others to access or use;
 - install, sell, publish, transmit, encumber, pledge, lend, copy, adapt, link to, post, rent, lease or lend, make available or distribute the Licensed Content to any third party, except as expressly permitted by this Agreement.
 - reverse engineer, decompile, remove or otherwise thwart any protections or disassemble the Licensed Content except and only to the extent that applicable law expressly permits, despite this limitation;
 - access or use any Licensed Content for which you are not providing a training session to End Users using the Licensed Content;
 - access or use any Licensed Content that you have not been authorized by Microsoft to access and use; or
 - transfer the Licensed Content, in whole or in part, or assign this agreement to any third party.
 6. **RESERVATION OF RIGHTS AND OWNERSHIP.** Microsoft reserves all rights not expressly granted to you in this agreement. The Licensed Content is protected by copyright and other intellectual property laws and treaties. Microsoft or its suppliers own the title, copyright, and other intellectual property rights in the Licensed Content. You may not remove or obscure any copyright, trademark or patent notices that appear on the Licensed Content or any components thereof, as delivered to you.

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

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

This limitation applies to

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

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

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

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

EXONÉRATION DE GARANTIE. Le contenu sous licence visé par une licence est offert « tel quel ». Toute utilisation de ce contenu sous licence est à votre seule risque et péril. Microsoft n'accorde aucune autre garantie expresse. Vous pouvez bénéficier de droits additionnels en vertu du droit local sur la protection 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 December 2011

Welcome!

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

- **Microsoft Certified Trainers and Instructors**—Your instructor is a technical and instructional expert who meets ongoing certification requirements. And, if instructors are delivering training at one of our Certified Partners for Learning Solutions, they are also evaluated throughout the year by students and by Microsoft.
- **Certification Exam Benefits**—After training, consider taking a Microsoft Certification exam. Microsoft Certifications validate your skills on Microsoft technologies and can help differentiate you when finding a job or boosting your career. In fact, independent research by IDC concluded that 75% of managers believe certifications are important to team performance¹. Ask your instructor about Microsoft Certification exam promotions and discounts that may be available to you.
- **Customer Satisfaction Guarantee**—Our Certified Partners for Learning Solutions offer a satisfaction guarantee and we hold them accountable for it. At the end of class, please complete an evaluation of today's experience. We value your feedback!

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

Sincerely,

Microsoft Learning
www.microsoft.com/learning

Microsoft | Learning

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

Acknowledgments

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

Stan Reimer – Content Developer

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

Damir Dizdarevic – Subject Matter Expert/Content Developer

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

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. Orin 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, Microsoft System Center, and security topics. Orin founded and runs the Melbourne System Center Users Group.

Vladimir Meloski – Content Developer

Vladimir Meloski is an MCT, an MVP on Exchange Server, and consultant providing unified communications and infrastructure solutions based on Microsoft Exchange Server, Microsoft 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 MOC courses.

Nick Portlock – Author

Nick Portlock has been an MCT for 15 years. He is a self-employed IT trainer, consultant and author. Last year, Nick taught in more than 20 countries. He specializes in Active Directory, Group Policy, and Domain Name System, and has consulted with a variety of companies over the last decade. Nick has reviewed more than 100 Microsoft courses, and is a member of the Windows® 7 STEP program.

Gary Dunlop – Subject Matter Expert

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

Ulf B. Simon-Weidner – Technical Reviewer

Ulf B. Simon-Weidner is a senior consultant with a European provider for infrastructure solutions in Germany. He also is an independent author, consultant, speaker and trainer. Ulf has been repeatedly awarded MVP for Windows Server Directory Services for the past decade, and has been an MCT for more than 10 years. Throughout his professional career, Ulf has had several consulting engagements with major European or Global corporations. He also published multiple books and magazine articles about Active Directory, Windows Server, Windows client operating systems, and security. Ulf is a frequently visiting speaker for conferences including Microsoft TechEd North America and Europe, or The Experts Conference. Ulf provides his technical and from-the-field experience in multiple Windows Server coursewares as a technical reviewer.

Contents

Module 1: Implementing Advanced Network Services

Lesson 1: Configuring Advanced DHCP Features	1-2
Lesson 2: Configuring Advanced DNS Settings	1-11
Lesson 3: Implementing IPAM	1-21
Lab: Implementing Advanced Network Services	1-31

Module 2: Implementing Advanced File Services

Lesson 1: Configuring iSCSI Storage	2-2
Lesson 2: Configuring BranchCache	2-9
Lesson 3: Optimizing Storage Usage	2-16
Lab A: Implementing Advanced File Services	2-22
Lab B: Implementing BranchCache	2-28

Module 3: Implementing Dynamic Access Control

Lesson 1: Overview of Dynamic Access Control	3-2
Lesson 2: Planning for Dynamic Access Control	3-8
Lesson 3: Deploying Dynamic Access Control	3-13
Lab: Implementing Dynamic Access Control	3-22

Module 4: Implementing Network Load Balancing

Lesson 1: Overview of NLB	4-2
Lesson 2: Configuring an NLB Cluster	4-5
Lesson 3: Planning an NLB Implementation	4-10
Lab: Implementing Network Load Balancing	4-16

Module 5: Implementing Failover Clustering

Lesson 1: Overview of Failover Clustering	5-2
Lesson 2: Implementing a Failover Cluster	5-14
Lesson 3: Configuring Highly Available Applications and Services on a Failover Cluster	5-20
Lesson 4: Maintaining a Failover Cluster	5-25
Lesson 5: Implementing a Multi-Site Failover Cluster	5-30
Lab: Implementing Failover Clustering	5-36

Module 6: Implementing Failover Clustering with Hyper-V

Lesson 1: Overview of Integrating Hyper-V with Failover Clustering	6-2
Lesson 2: Implementing Hyper-V Virtual Machines on Failover Clusters	6-7
Lesson 3: Implementing Hyper-V Virtual Machine Movement	6-15
Lesson 4: Managing Hyper-V Virtual Environments by Using VMM	6-21

Lab: Implementing Failover Clustering with Hyper-V	6-31
Module 7: Implementing Disaster Recovery	
Lesson 1: Overview of Disaster Recovery	7-2
Lesson 2: Implementing Windows Server Backup	7-7
Lesson 3: Implementing Server and Data Recovery	7-16
Lab: Implementing Windows Server Backup and Restore	7-20
Module 8: Implementing Distributed Active Directory Domain Services Deployments	
Lesson 1: Overview of Distributed AD DS Deployments	8-2
Lesson 2: Deploying a Distributed AD DS Environment	8-9
Lesson 3: Configuring AD DS Trusts	8-18
Lab: Implementing Complex AD DS Deployments	8-23
Module 9: Implementing Active Directory Domain Services Sites and Replication	
Lesson 1: Overview of AD DS Replication	9-2
Lesson 2: Configuring AD DS Sites	9-10
Lesson 3: Configuring and Monitoring AD DS Replication	9-16
Lab: Implementing AD DS Sites and Replication	9-22
Module 10: Implementing Active Directory Certificate Services	
Lesson 1: PKI Overview	10-2
Lesson 2: Deploying CAs	10-10
Lesson 3: Deploying and Managing Certificate Templates	10-16
Lesson 4: Implementing Certificate Distribution and Revocation	10-21
Lesson 5: Managing Certificate Recovery	10-29
Lab: Implementing Active Directory Certificate Services	10-33
Module 11: Implementing Active Directory Rights Management Services	
Lesson 1: AD RMS Overview	11-2
Lesson 2: Deploying and Managing an AD RMS Infrastructure	11-7
Lesson 3: Configuring AD RMS Content Protection	11-13
Lesson 4: Configuring External Access to AD RMS	11-19
Lab: Configuring AD RMS	11-24
Module 12: Implementing Active Directory Federation Services	
Lesson 1: Overview of AD FS	12-2
Lesson 2: Deploying AD FS	12-11
Lesson 3: Implementing AD FS for a Single Organization	12-17
Lesson 4: Deploying AD FS in a B2B Federation Scenario	12-23
Lab: Implementing AD FS	12-28

Lab Answer Keys

Module 1 Lab: Implementing Advanced Network Services	L1-1
Module 2 Lab A: Implementing Advanced File Services	L2-11
Module 2 Lab B: Implementing BranchCache	L2-18
Module 3 Lab: Implementing Dynamic Access Control	L3-25
Module 4 Lab: Implementing Network Load Balancing	L4-35
Module 5 Lab: Implementing Failover Clustering	L5-41
Module 6 Lab: Implementing Failover Clustering with Hyper-V	L6-49
Module 7 Lab: Implementing Windows Server Backup and Restore	L7-55
Module 8 Lab: Implementing Complex AD DS Deployments	L8-61
Module 9 Lab: Implementing AD DS Sites and Replication	L9-67
Module 10 Lab: Implementing Active Directory Certificate Services	L10-71
Module 11 Lab: Configuring AD RMS	L11-85
Module 12 Lab: Implementing AD FS	L12-95

MCT USE ONLY. STUDENT USE PROHIBITED

About This Course

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

Course Description



Note: This first release ('A') MOC version of course 20412A has been developed on prerelease software (Release Candidate (RC)). Microsoft Learning will release a 'B' version of this course after the RTM version of the software is available.

This course will provide you with the knowledge and skills you need to provision advanced services in a Windows Server® 2012 enterprise environment. This course will teach you how to configure and manage high availability features, file and storage solutions, and network services in Windows Server 2012. You will also learn about configuring the Active Directory® Domain Services (AD DS) infrastructure, and implementing backups and disaster recovery.

Audience

This course is intended for IT Professionals who have real-world hands-on experience implementing, managing and maintaining a Windows Server 2012 infrastructure in an existing Enterprise environment, and wish to acquire the skills and knowledge necessary to carry out advanced management and provisioning of services within that Windows Server 2012 environment.

The secondary audience for this course will be candidates aspiring to acquire the Microsoft Certified Systems Administrator (MCSA) credential either in its own right or in order to proceed in acquiring the Microsoft Certified System Engineer (MCSE) credentials, for which this is a prerequisite. IT professionals seeking certification in the *70-412: Configuring Advanced Windows Server® 2012 Services* exam also may take this course.

Student Prerequisites

This course requires that you meet the following prerequisites:

- At least two years hands-on experience working in a Windows Server 2008 or Windows Server 2012 environment
- Equivalent knowledge of *20410A: Installing and Configuring Windows Server® 2012* course
 - Installing and configuring Windows Server 2012 into existing enterprise environments, or as standalone installations
 - Configuring local storage
 - Configuring roles and features
 - Configuring file and print services
 - Configuring Windows Server 2012 servers for local and remote administration
 - Configuring IPv4 and IPv6 addresses
 - Configuring Domain Name System (DNS) and Dynamic Host Configuration Protocol (DHCP) services
 - Installing domain controllers

- Creating and configuring users, groups, computers and organizational units (OUs)
- Creating and managing group policies
- Configuring local security policies
- Configuring Windows Firewall
- Configuring Windows Server 2012 Hyper-V®
- Equivalent knowledge of *20411A: Administering Windows Server® 2012* course
 - Deploying and managing Windows Server images
 - Installing and configuring Update Services
 - Monitoring the Windows Server 2012 environment
 - Installing and configuring Distributed Files System (DFS)
 - Installing and configuring File Server Resource Manager (FSRM)
 - Configuring file and disk access, and audit policies
 - Configuring DNS security and integration with AD DS
 - Maintaining network integrity by configuring Network Access using Network Policy Server (NPS) and Network Access Protection (NAP)
 - Configuring Remote Access using virtual private networks (VPNs) and Windows® 7 DirectAccess
 - Configuring Domain Controllers
 - Managing and maintaining the Active Directory environment
 - Managing and maintaining the Windows Server 2012 domain environment using Group Policy

Course Objectives

After completing this course, students will be able to:

- Configure advanced features for DHCP and DNS, and configure IP Address Management (IPAM).
- Configure file services to meet advanced business requirements.
- Configure Dynamic Access Control (DAC) to manage and audit access to shared files.
- Provide high availability and load balancing for web-based applications by implementing Network Load Balancing (NLB).
- Provide high availability for network services and applications by implementing failover clustering.
- Deploy and manage Hyper-V virtual machines in a failover cluster.
- Implement a backup and disaster recovery solution based on business and technical requirements.
- Plan and implement an AD DS deployment that includes multiple domains and forests.
- Plan and implement an AD DS deployment that includes multiple locations.
- Implement an Active Directory Certificate Services (AD CS) deployment.
- Implement an Active Directory Rights Management Services (AD RMS) deployment.
- Implement an Active Directory Federation Services (AD FS) deployment.

Course Outline

The course outline is as follows:

Module 1, "Implementing Advanced Network Services" describes how to configure advanced DHCP features and DNS settings, and implement IPAM, which is a new Windows Server 2012 feature.

Module 2, "Implementing Advanced File Services" describes how to configure Internet Small Computer System Interface (iSCSI) storage and Windows BranchCache®. The module also describes how to implement Windows Server 2012 features that optimize storage utilization.

Module 3, "Implementing Dynamic Access Control" describes DAC, which is a new Windows Server 2012 feature. It also explains how to plan for a DAC implementation, and how to configure DAC.

Module 4, "Implementing Network Load Balancing" describes the features and working of network load balancing (NLB). It also explains how to configure an NLB cluster and plan an NLB implementation.

Module 5, "Implementing Failover Clustering" describes failover clustering features in Windows Server 2012. The module also describes how to implement and maintain failover clusters, and how to configure highly available applications and services on a failover cluster.

Module 6, "Implementing Failover Clustering with Hyper-V" describes options to make virtual machines highly available, and covers the implementation of Hyper-V virtual machines on failover clusters and Hyper-V virtual machine movement.

Module 7, "Implementing Disaster Recovery" describes disaster recovery, server and data recovery, and the planning and implementation of a backup solution for Windows Server 2012.

Module 8, "Implementing Distributed Active Directory Domain Services Deployments" provides an overview of distributed AD DS deployments and the process of implementation for the same. It also describes how to configure AD DS trusts, and implement complex AD DS deployments.

Module 9, "Implementing Active Directory Domain Services Sites and Replication" describes how replication works in AD DS in a Windows Server 2012 AD DS environment, and how to configure AD DS sites to optimize AD DS network traffic. It also shows how to configure and monitor AD DS replication.

Module 10, "Implementing Active Directory Certificate Services" provides an overview of Public Key Infrastructure (PKI), and describes how to deploy certification authorities (CAs) and certificate templates. It also covers certificate distribution and revocation, and management of certificate recovery.

Module 11, "Implementing Active Directory Rights Management Services" describes AD RMS and how you can use it to achieve content protection. It also explains how to deploy and manage an AD RMS infrastructure, and configure AD RMS content protection and external access to AD RMS.

Module 12, "Implementing Active Directory Federation Services" describes the identity federation business scenarios, and how you can use AD FS to address the scenarios. It also describes how to deploy AD FS, and how to implement it for a single organization, and in a business-to-business (B2B) scenario.

Exam/Course Mapping

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

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

Exam 70-412: Configuring Advanced Windows Server 2012 Services				
Exam Objective Domain		Course Content		
Configure and Manage High Availability (16%)		Module	Lesson	Lab
Configure Network Load Balancing (NLB).	This objective may include but is not limited to: Install NLB nodes; configure NLB prerequisites; configure affinity; configure port rules; configure cluster operation mode; upgrade an NLB cluster	Mod 4	Lesson 1/2/3	Mod 4 Ex 1/2/3
Configure failover clustering.	This objective may include but is not limited to: Configure Quorum; configure cluster networking; restore single node or cluster configuration; configure cluster storage; implement Cluster Aware Updating; upgrade a cluster	Mod 5	Lesson 2/5	Mod 5 Ex 2/4
Manage failover clustering roles.	This objective may include but is not limited to: Configure role-specific settings including continuously available shares; configure VM monitoring; configure failover and preference settings	Mod 5	Lesson 1/4	Mod 5 Ex 1
Manage Virtual Machine (VM) movement.	This objective may include but is not limited to: Perform Live Migration; perform quick migration; perform storage migration; import, export, and copy VMs; migrate from other platforms (P2V and V2V)	Mod 5	Lesson 3/4	Mod 5 Ex 3
Configure File and Storage Solutions (15%)				
Configure advanced file services.	This objective may include but is not limited to: Configure NFS data store; configure BranchCache; configure File Classification Infrastructure (FCI) using File Server Resource Manager (FSRM); configure file access auditing	Mod 2	Lesson 2/3	Mod 2 Ex 2/3

(continued)


Exam 70-412: Configuring Advanced Windows Server 2012 Services				
Exam Objective Domain		Course Content		
Configure File and Storage Solutions (15%)				
Implement Dynamic Access Control (DAC).	This objective may include but is not limited to: Configure user and device claim types; implement policy changes and staging; perform access-denied remediation; configure file classification	Mod 3	Lesson 1/2/3	Mod 3 Ex 1/2/3/4/5/6
Configure and optimize storage.	This objective may include but is not limited to: Configure iSCSI Target and Initiator; configure Internet Storage Name server (iSNS); implement thin provisioning and trim; manage server free space using Features on Demand	Mod 2	Lesson 1/3	Mod 2 Ex 1
Implement Business Continuity and Disaster Recovery (18%)				
Configure and manage backups.	This objective may include but is not limited to: Configure Windows Server backups; configure Windows Online backups; configure role-specific backups; manage VSS settings using VSSAdmin; create System Restore snapshots	Mod 7	Lesson 2/3	Mod 7 Ex 1/2/3/4
Recover servers.	This objective may include but is not limited to: Restore from backups; perform a Bare Metal Restore (BMR); recover servers using Windows Recovery Environment (Win RE) and safe mode; apply System Restore snapshots; configure the Boot Configuration Data (BCD) store	Mod 7	Lesson 2/3	Mod 7 Ex 1/2/3/4
Configure site-level fault tolerance.	This objective may include but is not limited to: Configure Hyper-V Replica including Hyper-V Replica Broker and VMs; configure multi-site clustering including network settings, Quorum, and failover settings	Mod 6	Lessons 1/3	Mod 6 Ex 1
		Mod 5	Lesson 1	
Configure Network Services (17%)				
Implement an advanced Dynamic Host Configuration Protocol (DHCP) solution.	This objective may include but is not limited to: Create and configure superscopes and multicast scopes; implement DHCPv6; configure high availability for DHCP including DHCP failover and split scopes; configure DHCP Name Protection	Mod 1	Lesson 1	Mod 1 Ex 1

(continued)

Exam 70-412: Configuring Advanced Windows Server 2012 Services				
Exam Objective Domain		Course Content		
Configure Network Services (17%)				
Implement an advanced DNS solution.	This objective may include but is not limited to: Configure security for DNS including DNSSEC, DNS Socket Pool, and cache locking; configure DNS logging; configure delegated administration; configure recursion; configure netmask ordering; configure a GlobalNames zone	Mod 1	Lesson 2	Mod 1 Ex 2
Deploy and manage IPAM.	This objective may include but is not limited to: Configure IPAM manually or by using Group Policy; configure server discovery; create and manage IP blocks and ranges; monitor utilization of IP address space; migrate to IPAM; delegate IPAM administration; manage IPAM collections	Mod 1	Lesson 3	Mod 1 Ex 3
Configure the Active Directory Infrastructure (18%)				
Configure a forest or a domain	This objective may include but is not limited to: Implement multi-domain and multi-forest Active Directory environments including interoperability with previous versions of Active Directory; upgrade existing domains and forests including environment preparation and functional levels; configure multiple user principal name (UPN) suffixes	Mod 8	Lesson 1/2	Mod 8 Ex 1
Configure trusts.	This objective may include but is not limited to: Configure external, forest, shortcut, and realm trusts; configure trust authentication; configure SID filtering; configure name suffix routing	Mod 8	Lesson 3	Mod 8 Ex 2
Configure sites.	This objective may include but is not limited to: Configure sites and subnets; create and configure site links; manage site coverage; manage registration of SRV records; move domain controllers between sites	Mod 9	Lesson 2/3	Mod Ex 1/2
Manage Active Directory and SYSVOL replication.	This objective may include but is not limited to: Configure replication to Read-Only Domain Controllers (RODCs); configure Password Replication Policy (PRP) for RODCs; monitor and manage replication; upgrade SYSVOL replication to Distributed File System Replication (DFSR)	Mod 9	Lesson 1/3	Mod 9 Ex 3

(continued)

Exam 70-412: Configuring Advanced Windows Server 2012 Services				
Exam Objective Domain		Course Content		
Configure Identity and Access Solutions (15%)				
Implement Active Directory Federation Services 2.1 (AD FSv2.1).	This objective may include but is not limited to: Implement claims-based authentication including Relying Party Trusts; configure Claims Provider Trust rules; configure attribute stores including Active Directory Lightweight Directory Services (AD LDS); manage AD FS certificates; configure AD FS proxy; integrate with Cloud Services	Mod 12	Lesson 1/2/3/4	Mod 12 Ex 1/2/3/4
Install and configure Active Directory Certificate Services (AD CS).	This objective may include but is not limited to: Install an Enterprise Certificate Authority (CA); configure CRL distribution points; install and configure Online Responder; implement administrative role separation; configure CA backup and recovery	Mod 10	Lesson 1/2/3/4/5	Mod 10 Ex 1/2/3/4/5/6
Manage certificates.	This objective may include but is not limited to: Manage certificate templates; implement and manage certificate deployment, validation, and revocation; manage certificate renewal; manage certificate enrollment and renewal to computers and users using Group Policies; configure and manage key archival and recovery	Mod 10	Lesson 3/4/5	Mod 10 EX 3/4/5/6
Install and configure Active Directory Rights Management Services (AD RMS).	This objective may include but is not limited to: Install a licensing or certificate AD RMS server; manage AD RMS Service Connection Point (SCP); manage AD RMS client deployment; manage Trusted User Domains; manage Trusted Publishing Domains; manage Federated Identity support; manage RMS templates; configure Exclusion Policies	Mod 11	Lesson 1/2/3/4	Mod 11 Ex 1/2/3/4

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

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

- Real world, hands-on experience Implementing, Managing and Configuring Active Directory and Networking infrastructure, working in a Windows Server 2008, Windows Server 2008 R2 or Windows Server 2012 Enterprise environment.
- Additional study outside of the content in this handbook

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

<http://www.microsoft.com/learning/en/us/exam.aspx?ID=70-412&locale=en-us#tab3>

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

<http://www.microsoft.com/learning/en/us/exam.aspx?ID=70-412&locale=en-us#tab1>

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

Course Materials

The following materials are included with your kit:

- **Course Handbook:** a succinct classroom learning guide that provides the critical technical information in a crisp, tightly-focused format, which is essential for an effective in-class learning experience.
 - **Lessons:** guide you through the learning objectives and provide the key points that are critical to the success of the in-class learning experience.
 - **Labs:** provide a real-world, hands-on platform for you to apply the knowledge and skills learned in the module.
 - **Module Reviews and Takeaways:** provide on-the-job reference material to boost knowledge and skills retention.
 - **Lab Answer Keys:** provide step-by-step lab solution guidance.



Course Companion Content: searchable, easy-to-browse digital content with integrated premium online resources that supplement the Course Handbook.

- **Modules:** include companion content, such as questions and answers, detailed demo steps and additional reading links, for each lesson. Additionally, they include Lab Review questions and answers and Module Reviews and Takeaways sections, which contain the review questions and answers, best practices, common issues and troubleshooting tips with answers, and real-world issues and scenarios with answers.
- **Resources:** include well-categorized additional resources that give you immediate access to the most current premium content on TechNet, MSDN®, or Microsoft® Press®.



Note: For this version of the Courseware on Prerelease Software (specify RC0/Beta etc.), Companion Content is not available. However, the Companion Content will be published when the next (B) version of this course is released, and students who have taken this course will be able to download the Companion Content at that time from the <http://www.microsoft.com/learning/companionmoc> site. Please check with your instructor when the 'B' version of this course is scheduled to release to learn when you can access Companion Content for this course.



Student Course files: includes the Allfiles.exe, a self-extracting executable file that contains all required files for the labs and demonstrations.



Note: For this version of the Courseware on Prerelease Software (specify RC0/Beta etc.), Allfiles.exe file is not available. However, this file will be published when the next (B) version of this course is released, and students who have taken this course will be able to download the Allfiles.exe at that time from the <http://www.microsoft.com/learning/companionmoc> site.

- **Course evaluation:** at the end of the course, you will have the opportunity to complete an online evaluation to provide feedback on the course, training facility, and instructor.
- To provide additional comments or feedback on the course, send an email to support@mscourseware.com. To inquire about the Microsoft Certification Program, send an email to mcphep@microsoft.com.

Virtual Machine Environment

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

Virtual Machine Configuration

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



Important: At the end of each lab, you must close the virtual machine and must not save any changes. To close a virtual machine (VM) without saving the changes, perform the following steps:

1. On the virtual machine, on the **Action** menu, click **Close**.
2. In the **Close** dialog box, in the **What do you want the virtual machine to do?** list, click **Turn off** and delete changes, and then click **OK**.

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

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

(continued)

Virtual machine	Role
20412A-LON-HOST1	Windows Server 2012 Member server in the Adatum.com domain
20412A-LON-HOST2	Windows Server 2012 Member server in the Adatum.com domain
20412A-TOR-DC1	Windows Server 2012 Member server in the Adatum.com domain

Software Configuration

The following software is installed on the VMs:

- Windows Server 2012 Datacenter Edition, Release Candidate
- Windows 8, Release Preview
- Office 2010, SP1

Classroom Setup

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

Course Hardware Level

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

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

*Striped

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

MCT USE ONLY. STUDENT USE PROHIBITED

Module 1

Implementing Advanced Network Services

Contents:

Module Overview	1-1
Lesson 1: Configuring Advanced DHCP Features	1-2
Lesson 2: Configuring Advanced DNS Settings	1-11
Lesson 3: Implementing IPAM	1-21
Lab: Implementing Advanced Network Services	1-31
Module Review and Takeaways	1-36

Module Overview

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

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

Objectives

After completing this module you will be able to:

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

Lesson 1

Configuring Advanced DHCP Features

DHCP plays an important role in the Windows Server 2012 operating system infrastructure. It is the primary means of distributing important network configuration information to network clients, and it provides configuration information to other network-enabled services, including Windows Deployment Services (WDS) and Network Access Protection (NAP). To support a Windows Server-based network infrastructure, it is important that you understand the DHCP server role. Windows Server 2012 improves the functionality of DHCP by providing failover capabilities.

Lesson Objectives

After completing this lesson you will be able to:

- Describe DHCP components.
- Explain how to configure DHCP interaction with DNS.
- Describe super scopes and multicast scopes.
- Explain how DHCP works with IPv6.
- Describe DHCP name protection.
- Describe DHCP failover.

Overview of DHCP Components

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

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

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


Component	Description
DHCP server service	After installing the DHCP role, the DHCP server is implemented as a service. This service can distribute IP addresses and other network configuration information to clients who request it.
DHCP scopes	The DHCP administrator configures the range of IP addresses and related information allotted to the server for distribution to requesting clients. Each scope can only be associated with a single IP subnet. A scope must

DHCP components consist of:

- The DHCP server service
- DHCP options
- DHCP console
- DHCP scopes
- DHCP database

When you use DHCP:

- Clients request IP configuration through a broadcast
- IP addresses are leased to clients for a configurable period, and are regularly renewed
- DHCP servers must be authorized in AD DS

Component	Description
	<p>consist of:</p> <ul style="list-style-type: none"> • A name and description • A range of addresses that can be distributed • A subnet mask <p>A scope can also define:</p> <ul style="list-style-type: none"> • IP addresses that should be excluded from distribution • The duration of the IP address lease • DHCP options <p>You can configure a single DHCP server with multiple scopes, but the server must be either connected directly to each subnet that it serves, or have a DHCP relay agent in place. Scopes also provide the primary way for the server to manage and distribute any related configuration parameters (DHCP options) to clients on the network.</p>
DHCP options	<p>When you assign the IP address inform, you can simultaneously assign many other network configuration parameters. The most common DHCP options include:</p> <ul style="list-style-type: none"> • Default Gateway IP address • DNS server IP address • DNS domain suffix • Windows Internet Name Service (WINS) server IP address <p>You can apply the options at different levels. They can be applied:</p> <ul style="list-style-type: none"> • Globally to all scopes • Specifically to particular scopes • To specific clients based on a Class ID value • To clients that have specific IP address reservations configured <p> Note: IPv6 scopes are slightly different, and will be discussed later in this lesson.</p>
DHCP database	<p>The DHCP database contains configuration data about the DHCP server, and stores information about the IP addresses that have been distributed. By default, the DHCP database files are stored in the %systemroot%\System32\Dhcp folder.</p>
DHCP console	<p>The DHCP console is the main administrative tool for managing all aspects of the DHCP server. This management console is installed automatically on any server that has the DHCP role installed. However, you can also install it on a remote server or Windows 8 client by using the Remote Server Administration Tools (RSAT) and by connecting to the DHCP server for remote management.</p>

How Clients Acquire IP Addresses

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

configuration on behalf of the requesting client on the remote subnet, and then forwards that configuration to the client.

DHCP Leases

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

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

DHCP Server Authorization

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

Configuring DHCP Interaction With DNS

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

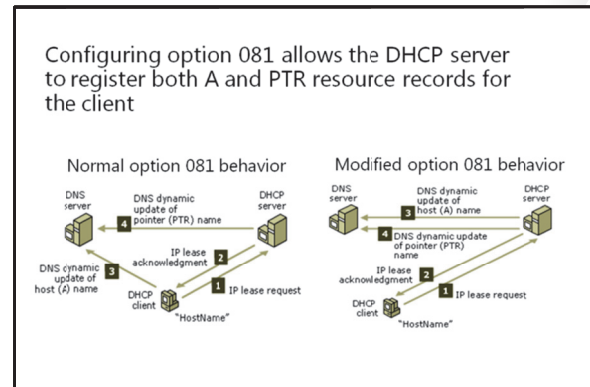
Configuring DHCP Option 081

You can configure DHCP option 081 to control the way that resource records are updated in the DNS database. This option permits the client to provide its fully qualified domain name (FQDN) and instructions to the DHCP server about how it would like the server to process DNS dynamic updates on its behalf. You configure option 081 on the **DNS** tab of the Properties window for the protocol node, or per scope in the DHCP console. You can also configure DHCP to perform updates on behalf of its clients to any DNS servers that support dynamic updates.

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

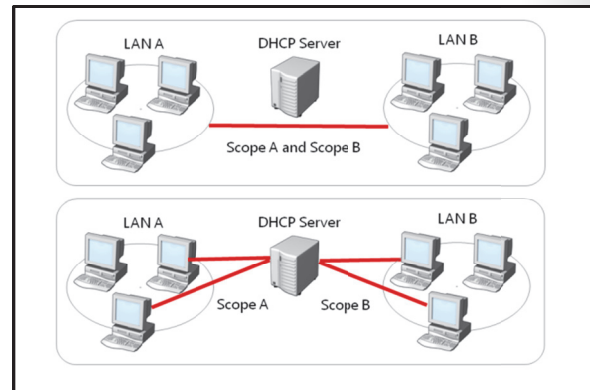
- The DHCP server dynamically updates DNS address host (A) resource records and pointer (PTR) resource records only if requested by the DHCP clients. By default, the client requests that the DHCP server register the DNS PTR resource record, while the client registers its own DNS A resource record.
-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

You can modify this option so that it instructs the DHCP server to always dynamically update DNS A resource records and PTR resource records no matter what the client requests. In this way, the DHCP server becomes the owner of the resource record because the DHCP server performed the registration of the resource records. Once the DHCP server becomes the owner of the client computer's A and PTR resource records, only that DHCP server can update the DNS resource records for the client computer based on the duration and renewal of the DHCP lease.



Configuring Advanced DHCP Scope Designs

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



Benefits of Superscopes

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

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

Multicast Scopes

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

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

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

DHCP Integration With IPv6

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

DHCPv6 supports stateful and stateless configurations

DHCPv6 also supports scopes that you can configure with the following properties:

- Name and description
- Preference
- Valid and Preferred lifetimes
- Prefix
- Exclusions
- DHCP options

configuration settings.



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

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

Stateful and Stateless Configuration

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

- Stateful configuration. Occurs when the DHCPv6 server assigns the IPv6 address to the client along with additional DHCP data.
- Stateless configuration. Occurs when the subnet router assigns IPv6 automatically, and the DHCPv6 server only assigns other IPv6 configuration settings.

DHCPv6 Scopes for IPv6

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

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

Configuring an IPv6 Scope

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

1. In the DHCP console, right-click the IPv6 node, and then click **New Scope**.

2. Configure a scope prefix and preference.
3. Define the starting and ending IP addresses, and any exclusions.
4. Configure the **Preferred** and **Valid** lifetime properties.
5. Activate the scope to enable it.

What Is DHCP Name Protection?

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

DHCP Name Protection:

- Prevents Windows operating systems from having their DNS name registrations overwritten by non-Windows operating systems that have the same name
- Uses a DHCID resource record to track the machines that originally requested the DNS names
- Is configurable at the network adapter level and at the scope level

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

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

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

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

To enable DHCP Name Protection for a scope:

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

What Is DHCP Failover?

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

DHCP Failover

DHCP clients renew their leases on their IP addresses at regular, configurable intervals. When the DHCP service fails, the leases time out and clients no longer have IP addresses. In the past, DHCP failover was not possible because DHCP servers were independent and unaware of each other. Therefore, configuring two separate DHCP servers to distribute the same pool of addresses could lead to duplicate addresses. Additionally, providing redundant DHCP services required you 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 and only for IPv4 scopes and subnets.

Configuring DHCP Failover

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



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

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

Mode	Characteristics
Hot Standby	In this mode, one server is the primary server and the other is the secondary server. The primary server actively assigns IP configurations for the scope or subnet. The secondary DHCP server only assumes this role if the primary server becomes unavailable. A DHCP server can simultaneously act as the primary for one scope or subnet, and be the secondary for another. Administrators must configure a percentage of the scope addresses to be assigned to the standby server. These addresses are supplied during the Maximum Client Lead Time (MCLT) interval if the primary server is down. The default MCLT value is 5 percent of the scope. The secondary server takes control of the whole IP range after the MCLT

DHCP failover:

- Requires failover relationships to have unique names
- Supports the Hot Standby mode and the Load Sharing mode

When you use DHCP failover:

- The MCLT determines when a failover partner assumes control of the subnet or scope
- The auto state switchover interval determines when a failover partner is considered to be down
- Message authentication can validate the failover messages
- Firewall rules are autoconfigured during DHCP installation

Mode	Characteristics
	interval has passed. Hot Standby mode is best suited to deployments in which a disaster recovery site is located at a different location. That way the DHCP server will not service clients unless there is a main server outage.
Load Sharing	This is the default mode. In this mode both servers 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.

MCLT

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

Auto State Switchover Interval

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

Message Authentication

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

Firewall Considerations

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

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

Demonstration: Configuring DHCP Failover

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

Demonstration Steps

Configure a DHCP failover relationship

1. Log on to **LON-SVR1** as **Adatum\Administrator**. Note that the server is authorized, but that no scopes are configured.
2. Switch to LON-DC1. In Server Manager, click **Tools**, and then on the drop-down list, click **DHCP**.
3. In the DHCP console, launch the Configure Failover Wizard.
4. Configure failover replication with the following settings:
 - Partner server: **172.16.0.21**

- Relationship Name: **Adatum**
 - Maximum Client Lead Time: **15 minutes**
 - Mode: **Load balance**
 - Load Balance Percentage: **50%**
 - State Switchover Interval: **60 minutes**
 - Message authentication shared secret: **Pa\$\$w0rd**
5. Complete the Configure Failover Wizard.
 6. Switch back to LON-SVR1, and note that the IPv4 node is active and the Adatum scope is configured.

Lesson 2

Configuring Advanced DNS Settings

In TCP/IP networks of any size, certain services are essential. DNS is one of the most critical network services for any network, because many other applications and services—including AD DS—rely on DNS to resolve resource names to IP addresses. Without DNS, user authentications fail, and network-based resources and applications may become inaccessible. For this reasons, you need to manage and protect DNS. This lesson discusses management techniques and options for optimizing DNS resolution. Windows Server 2012 implements DNSSEC to protect DNS responses. Windows Server 2012 also supports global name zones to provide single-label name resolution.

Lesson Objectives

After completing this lesson you will be able to:

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

Managing DNS Services

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

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

Delegating Administration of DNS

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

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

Configuring DNS Logging

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

- Starting and stopping of the DNS service.

To manage DNS services:

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

Backup methods for the DNS database depend on how the database is deployed:

- Backup Active Directory integrated zones through system state backups or by using `dnscmd`
- Nonintegrated primary zone are single files that you can copy or back up

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

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

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

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

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

Aging and Scavenging

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

Aging and scavenging is disabled by default. You can enable aging and scavenging in the **Advanced** properties of the DNS server, or you can enable it for selected zones in the zone's Properties window.

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



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

Backing Up the DNS Database

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

Backing Up Active Directory Integrated Zones

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

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

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

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

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

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

Backing Up Primary Zones

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

Optimizing DNS Name Resolution

In a typical DNS query event, a client computer attempts to resolve a FQDN to an IP address. For example, if a user tries to go to the FQDN www.microsoft.com, the client computer will perform a recursive query to the DNS server that it is configured to discover the IP address associated with that FQDN. The local DNS server must then respond with an authoritative response. If the local DNS server has no copy of the DNS namespace for which it was queried, it will respond with an authoritative answer to the client computer. If the local DNS server does not have that information, it will perform recursion. *Recursion* refers to the process of having the local DNS server itself make a recursive query to another DNS server until it finds the authoritative answer and returns that answer to the client that made the original request. By default, this server will be one of the servers on the Internet that is listed as a root hint. When the local DNS server receives a response, it will return that information to the original requesting client computer.

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

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

- Forwarding
- Conditional forwarding
- Stub zones
- Netmask ordering

Forwarding

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

Conditional Forwarding

You also can use conditional forwarders to forward queries according to specific domain names. A *conditional forwarder* is a DNS server on a network that forwards DNS queries according to the query's DNS domain name. For example, you can configure a DNS server to forward all queries that it receives for names ending with `corp.adatum.com` to the IP address of a specific DNS server or to the IP addresses of multiple DNS servers. This can be useful when you have multiple DNS namespaces in a forest. For example, suppose `Contoso.com` and `Adatum.com` are merged. Rather than each domain having to host a complete replica of the other domain's DNS database, you could create conditional forwarders so that they point to each other's specific DNS servers for resolution of internal DNS names.

Stub Zones

A *stub zone* is a copy of a zone that contains only those resource records necessary to identify that zone's authoritative DNS servers. A stub zone resolves names between separate DNS namespaces, which might be necessary when you want a DNS server that is hosting a parent zone to remain aware of all the authoritative DNS servers for one of its child zones. A stub zone that is hosted on a parent domain DNS server will receive a list of all new DNS servers for the child zone, when it requests an update from the stub zone's master server. By using this method, the DNS server that is hosting the parent zone maintains a current list of the authoritative DNS servers for the child zone as they are added and removed.

A stub zone consists of the following:

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

Stub zones have the following characteristics:

- Stub zones are created using the New Zone Wizard.
- Stub zones can be stored in AD DS.
- Stub zones can be replicated either in the domain only, or throughout the entire forest.
- Stub zone master servers are one or more DNS servers that are responsible for the initial copy of the zone information, and are usually the DNS server that is hosting the primary zone for the delegated domain name.

Netmask Ordering

Netmask ordering returns addresses for type A (address records) DNS queries that prioritize resources on the client computer's local subnet to the client. In other words, addresses of hosts that are on the same subnet as the requesting client will have a higher priority in the DNS response to the client computer.

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

What Is the GlobalName Zone?

The DNS Server Service in Windows Server 2012 provides the GlobalName zone, which you can use to contain single-label names that are unique across an entire forest. This eliminates the need to use the NetBIOS-based WINS to provide support for single-label names. GlobalName zones provide single-label name resolution for large enterprise networks that do not deploy WINS and that have multiple DNS domain environments. GlobalName zones are created manually and do not support dynamic record registration.

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

You use a GlobalName zone to maintain a list of DNS search suffixes for resolving names among multiple DNS domain environments. For example, if an organization supports two DNS domains, such as adatum.com and contoso.com, users in the adatum.com DNS domain need to use a FQDN such as data.contoso.com to locate the servers in contoso.com. Otherwise, the domain administrator needs to add a DNS search suffix for contoso.com on all the systems in the adatum.com domain. If the clients just search for the server name "data," then the search would fail.

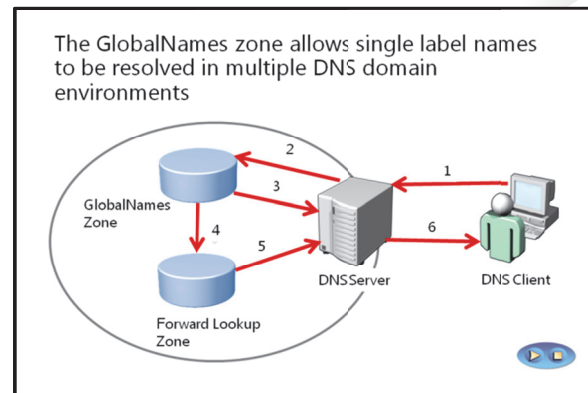
Global names are based on creating alias (CNAME) resource records in a special forward lookup zone that uses single names to point to FQDNs. For example, GlobalName zones would enable clients in both the adatum.com domain and the contoso.com domain to use a single label name, such as data, to locate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

Creating a GlobalName Zone

To create a GlobalName zone, do the following:

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

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



Options for Implementing DNS Security

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

- DNS cache locking
- DNS socket pool
- DNSSEC

DNS Cache Locking

Cache locking is a security feature available with Windows Server 2012 that allows you to control

when information in the DNS cache can be overwritten. When a recursive DNS server responds to a query, it caches the results so that it can respond quickly if it receives another query requesting the same information. The period of time the DNS server keeps information in its cache is determined by the Time to Live (TTL) value for a resource record. Information in the cache can be overwritten before the TTL expires if updated information about that resource record is received. If an attacker successfully overwrites information in the cache, the attacker might be able to redirect your network traffic to a malicious site. When you enable cache locking, the DNS server prohibits cached records from being overwritten for the duration of the TTL value.

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

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

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

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

3. Restart the DNS service to apply the changes.

DNS Socket Pool

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

The default size of the DNS socket pool is 2,500. When you configure the DNS socket pool, you can choose a size value from 0 to 10,000. The larger the value, the greater the protection you will have against DNS spoofing attacks. If the DNS server is running Windows Server 2012, you can also configure a DNS socket pool exclusion list.

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

1. Launch an elevated command prompt.

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

2. Run the following command:

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

3. Restart the DNS service to apply the changes

DNSSEC

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

How DNSSEC Works

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

DNSSEC functions as follows:

- 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 associated with digital signatures
- Resolvers use trust anchors to retrieve public keys and build trust chains
- DNSSEC requires trust anchors to be configured on all DNS servers participating in DNSSEC
- DNSSEC uses the NRPT, which contains rules that control the requesting client computer behavior for sending queries and handling responses

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

Trust Anchors

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

Name Resolution Policy Table

The Name Resolution Policy Table (NRPT) contains rules that control the DNS client behavior for sending DNS queries and processing the responses from those queries. For example, a DNSSEC rule prompts the client computer to check for validation of the response for a particular DNS domain suffix. As a best practice, 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:

1. Install Windows Server 2012 and assign the DNS role to the server. Typically, a domain controller also acts as the DNS server. However, this is not a requirement.
2. Sign the DNS zone by using the DNSSEC Configuration Wizard, which is located in the DNS console.

3. Configure trust anchor distribution points.
4. Configure the NRPT on the client computers.

Assigning the DNS Server Role

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

Signing the Zone

The following signing options are available:

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



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

Configuring Trust Anchor Distribution Points

If the zone is Active Directory–integrated, you should select to distribute the trust anchors to all the servers in the forest. If trust anchors are required on computers that are not joined to the domain—for example, a DNS server in the perimeter network (also known as DMZ, demilitarized zone, and screened subnet—then you should enable automated key rollover.



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

Configuring NRPT on Client Computers

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

New DNSSEC Features for Windows Server 2012

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

DNSSEC Zone Signing Wizard

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

- KSK options
- ZSK options
- Trust anchor distribution options
- Signing and polling parameters

DNSSEC enhancements for Windows Server 2012 include:

- Simplified DNSSEC implementation
- A DNSSEC Zone Signing Wizard that steps you through the process of signing and configuring signing parameters for zones
- The following new resource records:
 - DNSKEY
 - DS
 - RRSIG
 - NSEC

New Resource Records

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

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

Resource record	Purpose
DNSKEY	This record publishes the public key for the zone. It checks the authority of a response against the private key held by the DNS server. These keys require periodic replacement through key rollovers. Windows Server 2012 supports automated key rollovers.
DS	This record is a delegation record that contains the hash of the public key of a child zone. This record is signed by the parent zone's private key. If a child zone of a signed parent is also signed, the DS records from the child must be manually added to the parent so that a chain of trust can be created.
RRSIG	This record holds a signature for a set of DNS records. It is used to check the authority of a response.
NSEC	When the DNS response has no data to provide to the client, this record authenticates that the host does not exist.

Other New Enhancements

Other enhancements for Windows Server 2012 include:

- Support for DNS dynamic updates in DNSSEC signed zones.

- Automated trust anchor distribution through AD DS.
- Windows PowerShell–based command-line interface for management and scripting.

Demonstration: Configuring DNSSEC

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

Demonstration Steps

Configure DNSSEC

1. Log on to **LON-DC1** as **Adatum\Administrator**.
2. Start the DNS console.
3. Use the DNSSEC Zone Signing Wizard to sign the Adatum.com zone. Accept all default settings.
4. Verify that the DNSKEY resource records were created in the Trust Points zone.
5. Use the Group Policy Management Console to configure NRPT. Create a rule that enables DNSSEC for the Adatum.com suffix and that requires DNS client computers to verify that the name and address data is validated.

Lesson 3

Implementing IPAM

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

Lesson Objectives

After completing this lesson you will be able to:

- Describe IPAM.
- Describe IPAM architecture.
- Describe the requirements for IPAM implementations.
- Manage IP addressing using IPAM.
- Install and configure IPAM.
- Describe considerations for implementing IPAM.

What Is IPAM?

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

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

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

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

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

Characteristics of IPAM

Characteristics of IPAM include:

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

Benefits of IPAM

IPAM benefits include:

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



Note: IPAM does not support management and configuration of non-Microsoft network elements.

IPAM Architecture

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

IPAM architecture consists of:

- Four main modules
 - IPAM discovery
 - IPAM address space management
 - Multiserver management and monitoring
 - Operational auditing and IP address tracking
 - A server component and a client component
 - Local security groups that IPAM creates during installation
- You can deploy IPAM in the following topologies:
- Distributed
 - Centralized
 - Hybrid

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

The IPAM server can only manage one Active Directory forest. IPAM is deployed in one of three topologies:

- Distributed. An IPAM server is deployed to every site in the forest.
- Centralized. Only one IPAM server is deployed in the forest.
- Hybrid. A central IPAM server is deployed together with a dedicated IPAM server in each site.



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

IPAM has two main components:

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

IPAM Security Groups

The following table describes the local security groups that are created automatically when you install IPAM.

Group	Description
IPAM Users	Members of this group can view all information that is located in server discovery, IP address space, and server management. They can view IPAM and DHCP server operational events, but they cannot view IP address tracking information.
IPAM MSM	IPAM multi-server management (MSM) administrators have IPAM users privileges, and can perform IPAM common management tasks and server

Group	Description
Administrators	management tasks.
IPAM ASM Administrators	IPAM ASM administrators have IPAM users privileges, and can perform IPAM common management tasks and IP address space tasks.
IPAM IP Audit Administrators	Members of this group have IPAM users privileges, can perform IPAM common management tasks, and can view IP address tracking information.
IPAM Administrators	IPAM Administrators can view all IPAM data and perform all IPAM tasks.

Requirements for IPAM Implementation

To ensure a successful IPAM implementation, you 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.
- Log on 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 controller and NPS servers for IPAM's IP address tracking and auditing feature.

Prerequisites	Hardware and Software Requirements
<ul style="list-style-type: none"> • IPAM server must belong to the domain • IPAM server cannot be a domain controller • IPv6 must be enabled in order to manage IPv6 • Log on with a domain account • You must be in the correct IPAM local security group • Logging account logon events must be enabled for IP address tracking and auditing 	<ul style="list-style-type: none"> • CPU – dual core 2.0 GHz or higher • Windows Server 2012 operating system • 4 GB of RAM • 80 GB free disk space

IPAM Hardware and Software Requirements

The IPAM hardware and software requirements are as follows:

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

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

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

Managing IP Addressing Using IPAM

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

You can view and manage the IP address space 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 using the following views:

- DNS and DHCP servers
- DHCP scopes
- DNS zone monitoring
- Server groups

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

IPAM provides a number of ways to filter the view of the IP address space. You can customize how you view and manage the IP address space 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. Conceptually, an IP block is an IP subnet marked by a start and an end IP address, and it is typically assigned to an organization by various Regional Internet Registries (RIRs). Network administrators use IP address blocks to create and allocate IP address ranges to DHCP. They can add, import, edit, and delete IP address blocks. IPAM automatically maps IP address ranges to the appropriate IP address block based on the boundaries of the range. You can add and import IP address blocks in the IPAM console.

IP Address Ranges

IP address ranges are the next hierarchical level of IP address space entities after IP address blocks. Conceptually, an IP address range is an IP subnet marked by a start and end IP address, and it typically corresponds to a DHCP scope, or a static IPv4 or IPv6 address range or address pool that is used to assign addresses to hosts. An IP address range is uniquely identifiable by the value of the mandatory **Managed By Service** and **Service Instance** options, which help IPAM manage and maintain overlapping or duplicate IP address ranges from the same console. You can add or import IP address ranges from within the IPAM console.

IP Addresses

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

IP Address Inventory

In 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 defined by the **Device Type** option within the IP addresses view. These groups allow you to customize the way your address space displays for managing and tracking IP usage. You can add or import IP addresses from within the IPAM console. For example, you could add the IP addresses for printers or routers, assign IP address the appropriate device type of printer or router, and then view your IP inventory filtered by the device type you assigned.

IP Address Range Groups

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

Monitoring and Managing

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

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

Demonstration: Installing and Configuring IPAM

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

Demonstration Steps

Install and Configure IPAM

1. Log on to **LON-SVR2** as **Adatum\Administrator**.
2. In Server Manager, add the IPAM feature and all required supporting features.
3. In the IPAM Overview pane, provision the IPAM server using Group Policy.
4. Enter **IPAM** as the Group Policy Object (GPO) name prefix, and provision IPAM.
5. In the IPAM Overview pane, configure server discovery for the Adatum domain.
6. In the IPAM Overview pane, start the server discovery process.

7. In the IPAM Overview pane, add the servers to be managed.
8. Verify that IPAM access is currently blocked.
9. 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
```

10. Set the manageability status to **Managed**.
11. Switch to LON-DC1.
12. Force the update of Group Policy.
13. Switch back to LON-SVR2 and refresh the IPv4 view.
14. In the IPAM Overview pane, retrieve data from the managed server.

IPAM Management and Monitoring

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

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

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

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

Monitoring DHCP and DNS

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

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

For DHCP servers, the server view allows you to track various server settings, server options, the number of scopes, and the number of active leases that are configured on the server. For DNS servers, this view

allows you to track all zones that are configured on the server, along with details of the zone type. The view also allows you to see the total number of zones that are configured on the server, and the overall zone health status as derived from the zone status of individual zones on the server.

DHCP Server Management

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

- Edit DHCP server properties
- Edit DHCP server options
- Create DHCP scopes
- Configure predefined options and values
- Configure the user class across multiple servers simultaneously
- Create and edit new and existing user classes across multiple servers simultaneously
- Configure the vendor class across multiple servers simultaneously
- Start the management console for a selected DHCP server
- Retrieve server data from multiple servers

DNS Server Management

You can start the DNS management console for any managed DNS server from a central console in the IPAM server and 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. The IPAM configuration events console gathers all of the configuration events. These configuration event catalogs allows you to view, query, and generate reports of the consolidated configuration changes, along with details specific to each record.

Considerations for Implementing IPAM

IPAM is an agentless technology that uses Windows remote management protocols to manage, monitor, and collect data from distributed servers in the environment. As such, you should be aware of some implementation considerations.

Installation Considerations

Although IPAM is relatively simple to install, there are certain considerations:

- IPAM should not be installed on a domain controller, DHCP server, or DNS server.
- The installation wizard in Server Manager automatically installs the features required to support IPAM. There are no extra steps required of the administrator.

Considerations for IPAM implementation include:

- Installation considerations
- Functional considerations
- Administrative considerations
- Migrating existing IP data into IPAM

- The IPAM client is installed automatically on Windows Server 2012 along with the IPAM server, but you can uninstall the client separately.
- You can uninstall IPAM by using Server Manager. All dependencies, local security groups, and scheduled tasks will be deleted. The IPAM database will be detached from the Windows Internal Database.

Functional Considerations

Consider the following IPAM functional specifications:

- IPAM does not support multiple forest topologies.
- IPAM can only use Windows Internal Database; it cannot use any other type of database.
- The IPAM server must collect DHCP lease information to enable address tracking. Ensure that the DHCP audit log file size is configured so that it is large enough to contain audit events for the entire day.
- For domain controllers and network policy servers, enable the required events for logging. You can use Group Policy security settings to perform this task.

Administrative Considerations

Domain and enterprise administrators have full access to IPAM administration. You can delegate administrative duties to other users or groups by using the IPAM security groups. The installation process creates local security groups (which have no members by default) on the IPAM server. The local security groups provide the permissions that are required for administering and using the multiple services that IPAM employs.

IPAM installation automatically creates the local user groups listed in the following table.

Group	Description
IPAM Users	Members of this group can view all information in IPAM server inventory, IP address space, and IPAM server management consoles. They can view IPAM and DHCP server operational events, but they cannot view IP address tracking information.
IPAM MSM Administrators	Members of this group have all the privileges of the IPAM Users group, and they can perform IPAM monitoring and management tasks.
IPAM ASM Administrators	Members of this group have all the privileges of IPAM Users group, and they can perform IPAM IP address space tasks.
IPAM IP Audit Administrators	Members of this group have all the privileges of IPAM Users group, and they can view IP address tracking information.
IPAM Administrators	Members of this group can view all IPAM information and perform all IPAM tasks.

Migrating Existing IP Data Into IPAM

Many organizations use Microsoft Office Excel spreadsheets to document the IP address space allocation for static addresses and network devices. Because these spreadsheets must be updated manually, they are prone to error. You can migrate the existing data from these spreadsheets into IPAM by converting the spreadsheets to .csv files, and then importing the information into IPAM.

Lab: Implementing Advanced Network Services

Scenario

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

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

Objectives

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

Lab Setup

20412A-LON-DC1

20412A-LON-SVR1

20412A-LON-SVR2

20412A-LON-CL1

Estimated time: **60 minutes**

Virtual Machine(s)	20412A-LON-DC1 20412A-LON-SVR1 20412A-LON-SVR2 20412A-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 **20412A-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on using the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
5. Repeat steps 2-4 for **20412A-LON-SVR1** and **20412A-LON-SVR2**. Do not start **20412A-LON-CL1** until directed to do so.

Exercise 1: Configuring Advanced DHCP Settings

Scenario

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

The main tasks for this exercise are as follows:

1. Configure a superscope
2. Configure DHCP name protection
3. Configure and verify DHCP failover

► Task 1: Configure a superscope

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

► Task 2: Configure DHCP name protection

- Switch to the DHCP console on LON-DC1, and enable **DHCP Name Protection** for the IPv4 node.

► Task 3: Configure and verify DHCP failover

1. On LON-SVR1, start the DHCP console, and observe the current state of DHCP. Note that the server is authorized, but no scopes are configured.
2. On LON-DC1, in the DHCP console, launch the Configure Failover Wizard.
3. Configure failover replication with the following settings:
 - Partner server: **172.16.0.21**
 - Relationship Name: **Adatum**
 - Maximum Client Lead Time: **15 minutes**
 - Mode: Load balance
 - Load Balance Percentage: **50%**
 - State Switchover Interval: **60 minutes**

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

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

Exercise 2: Configuring Advanced DNS Settings

Scenario

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

The main tasks for this exercise are as follows:

1. Configure DNSSEC.
2. Configure the DNS socket pool.
3. Configure DNS cache locking.
4. Configure a GlobalName Zone.

► Task 1: Configure DNSSEC

1. On LON-DC1, start the DNS Manager.
2. Use the DNSSEC Zone Signing Wizard to sign the Adatum.com zone. Accept all the default settings.
3. Verify that the DNSKEY resource records have been created in the Trust Points zone.
4. Minimize the DNS console.
5. Use the Group Policy Management Console to configure NRPT. Create a rule that enables DNSSEC for the Adatum.com suffix, and that requires DNS clients to verify that the name and address data were validated.

► Task 2: Configure the DNS socket pool

1. On LON-DC1, start a command prompt with elevated credentials.

2. Run the following command to view the current size of the socket pool.

```
dnscmd /info /socketpoolsize
```

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

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

4. Restart the DNS service.
5. Run **dnscmd** to confirm the new socket pool size.

► Task 3: Configure DNS cache locking

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

```
dnscmd /info /CacheLockingPercent
```

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

```
dnscmd /config /CacheLockingPercent 75
```

3. Restart the DNS service.
4. Run **dnscmd** to confirm the new cache lock value.

► Task 4: Configure a GlobalName Zone

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

```
Dnscmd LON-DC1 /ZoneAdd Contoso.com /DsPrimary /DP /forest
```

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

```
dnscmd lon-dc1 /config /enablglobalnamesupport 1
```

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

```
Dnscmd LON-DC1 /ZoneAdd GlobalNames /DsPrimary /DP /forest
```

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

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

Exercise 3: Configuring IP Address Management

Scenario

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

The main tasks for this exercise are as follows:

1. Install the IPAM feature.
2. Configure IPAM-related GPOs.
3. Configure IP management server discovery.
4. Configure managed servers.
5. Configure and verify a new DHCP scope with IPAM..
6. Configure IP address blocks, record IP addresses, and create DHCP reservations and DNS records

► **Task 1: Install the IPAM feature**

- On LON-SVR2, install the **IP Address Management (IPAM) Server** feature.

► **Task 2: Configure IPAM-related GPOs**

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

► **Task 3: Configure IP management server discovery**

1. In the IPAM Overview pane, configure server discovery for the Adatum domain.
2. In the IPAM Overview pane, start the server discovery process.

► **Task 4: Configure managed servers**

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

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

3. Set the manageability status to **Managed**.
4. Switch to LON-DC1, and force the update of Group Policy using the `gpupdate /force`.
5. Return to LON-SVR2 and refresh the server access status for LON-DC1 and the IPv4 console view. It may take up to 10 minutes for the status to change. If necessary, repeat both refresh tasks as needed until a green check mark displays next to LON-DC1 and the IPAM Access Status shows Unblocked.
6. In the IPAM Overview pane, retrieve data from the managed server.

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

1. On LON-SVR2, use IPAM to create a new DHCP scope with the following parameters:
 - Scope start address: **10.0.0.50**
 - Scope end address: **10.0.0.100**
 - Subnet mask: **255.0.0.0**
 - Default gateway: **10.0.0.1**
2. On LON-DC1, verify the scope in the DHCP MMC.

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

1. On LON-SVR2, add an IP address block in the IPAM console with the following parameters:
 - Network ID: **172.16.0.0**
 - Prefix length: **16**
 - Description: **Head Office**
2. Add IP addresses for the network router by adding to the IP Address Inventory with the following parameters:
 - IP address: **172.16.0.1**
 - MAC address: **112233445566**
 - Device type: **Routers**
 - Description: **Head Office Router**
3. Use the IPAM console to create a DHCP reservation as follows:
 - IP address: **172.16.0.10**
 - MAC address: **223344556677**
 - Device type: **Host**
 - Reservation server name: **LON-DC1.Adatum.com**
 - Reservation name: **Webserver**
 - Reservation type: **Both**
4. Use the IPAM console to create the DNS host record as follows:
 - Device name: **Webserver**
 - Forward lookup zone: **Adatum.com**
 - Forward lookup primary server: **LON-DC1.adatum.com**
5. Right-click the **IPv4** entry and create the DHCP reservation and create the DNS Host record.
6. On LON-DC1, open the DHCP console and confirm that the reservation was created in the 172.16.0.0 scope.
7. On LON-DC1, open the DNS Manager console. Confirm that the DNS host record was created.

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

► **To prepare for the next module**

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

Module Review and Takeaways

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

Common Issues and Troubleshooting Tips

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

Real-world Issues and Scenarios

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

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

What are some possible causes of the incorrect configurations?

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

Best Practice

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

Tools

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

Module 2

Implementing Advanced File Services

Contents:

Module Overview	2-1
Lesson 1: Configuring iSCSI Storage	2-2
Lesson 2: Configuring BranchCache	2-9
Lesson 3: Optimizing Storage Usage	2-16
Lab A: Implementing Advanced File Services	2-22
Lab B: Implementing BranchCache	2-28
Module Review and Takeaways	2-33

Module Overview

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

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

Objectives

After completing this module, you will be able to:

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

Lesson 1

Configuring iSCSI Storage

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

Lesson Objectives

After completing this lesson, you will be able to:

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

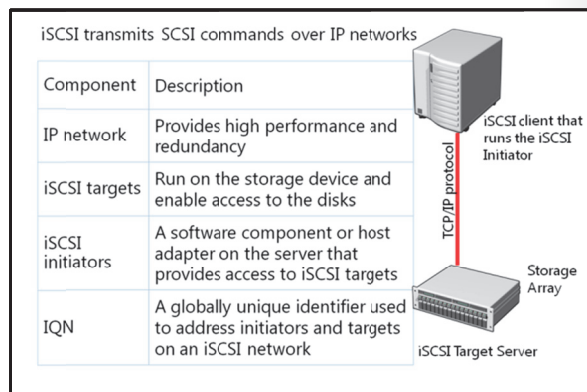
What Is iSCSI?

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

iSCSI relies on standard Ethernet networking architecture. Specialized hardware such as host bus adapters (HBA) or network switches are optional. iSCSI uses TCP/IP (typically, TCP port 3260). This means that iSCSI simply enables two hosts to negotiate tasks (for example, session establishment, flow control, and packet size,) and then exchange SCSI commands by using an existing Ethernet network. By doing this, iSCSI uses a popular, high performance, local storage bus subsystem architecture, and emulates it over LANs and WANs, 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 the performance of an iSCSI SAN deployment by operating it on a dedicated network or subnet, as best practices recommend.



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



An iSCSI SAN deployment includes the following:

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

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

iSCSI Target Server and iSCSI Initiator

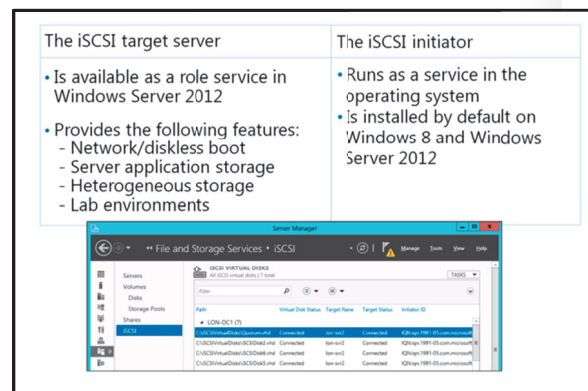
The iSCSI target server and the iSCSI initiator are described below.

iSCSI Target Server

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

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

- Network/diskless boot. By using boot-capable network adapters or a software loader, you can use iSCSI targets to deploy diskless servers quickly. By using differencing virtual disks, you can save up to 90 percent of the storage space for the operating system images. This is ideal for large deployments of identical operating system images, such as a Hyper-V® server farm, or high-performance computing (HPC) clusters.
- Server application storage. Some applications such as Hyper-V and Microsoft® Exchange Server require block storage. The iSCSI target server can provide these applications with continuously available block storage. Because the storage is remotely accessible, it can also combine block storage for central or branch office locations.



- Heterogeneous storage. iSCSI target server supports iSCSI initiators that are not based on the Windows® operating system, so you can share storage on Windows servers in mixed environments.
- Lab environments. The iSCSI target server role enables your Windows Server 2012 computers to be network-accessible block storage devices. This is useful in situations in which you want to test applications before deploying them on SAN storage.

iSCSI target servers that provide block storage utilize your existing Ethernet network; no additional hardware is required. If high availability is an important criterion, consider setting up a high availability cluster. With a high availability cluster, you will need shared storage for the cluster—either hardware Fibre Channel storage, or a Serial Attached SCSI (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 installed by default since Windows Server 2008 and Windows Vista. To connect your computer to an iSCSI target, you simply start the Microsoft iSCSI Initiator Service and configure it.

The new features in Windows Server 2012 include:

- Authentication. You can enable Challenge Handshake Authentication Protocol (CHAP) to authenticate initiator connections, or you can enable reverse CHAP to allow the initiator to authenticate the iSCSI target.
- Query initiator computer for ID. This is only supported with Windows 8 or Windows Server 2012.



Additional Reading: For more information about the introduction of iSCSI targets in Windows Server 2012, refer to:

<http://blogs.technet.com/b/filecab/archive/2012/05/21/introduction-of-iscsi-target-in-windows-server-2012.aspx>

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

Options for Implementing High Availability for iSCSI

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

Configuring iSCSI for High Availability

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

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

There are two technologies for implementing iSCSI for high availability:

- 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 iSCSI target server, you can use MPIO to provide failover redundancy during network outages

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:

- If you have multiple network interface cards (NICs) in your iSCSI initiator and iSCSI target server, you can use MPIO to provide failover redundancy during network outages.
- Requires a device-specific module (DSM) if you want to connect to a third-party SAN device that is connected to the iSCSI initiator. The Windows operating system includes a default MPIO DSM that is installed as the MPIO feature within Server Manager.
- 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

Because iSCSI is a protocol that provides access to storage devices over a TCP/IP network, it is crucial that you secure your iSCSI solution to protect it from malicious users or attacks. You can mitigate risks to your iSCSI solution by providing security at various infrastructure layers. The term *defense-in-depth* is often used to describe the use of multiple security technologies at different points throughout your organization.

Defense-in-Depth security strategy includes:

- Policies, procedures, and awareness. As a security best practice, security policy measures need to operate within the context of organizational policies. For example, consider enforcing a strong user password policy throughout the organization, but having an even stronger administrator password policy for accessing iSCSI storage devices and computers that have iSCSI management software installed.
- Physical security. If any unauthorized person can gain physical access to iSCSI storage devices or a computer on your network, then most other security measures are not useful. You must ensure that iSCSI storage devices, the computers that manage them, and the servers to which they are connected are physically secure, and that access is granted to authorized personnel only.
- Perimeter. Perimeter networks mark the boundary between public and private networks. Providing firewalls and reverse proxy servers in the perimeter network enables you to provide more secure corporate services across the public network, and to prevent possible attacks on the iSCSI storage devices from the Internet.
- Networks. Once you connect iSCSI storage devices to a network, they are susceptible to a number of threats. These threats include eavesdropping, spoofing, denial of service, and replay attacks. You should use authentication such as CHAP, to protect communication between iSCSI initiators and iSCSI

Use the defense-in-depth approach to secure iSCSI solutions:

Data	BitLocker, ACLs, EFS, backup/restore procedures
Application	Application hardening, antivirus
Host	Hardening, authentication, update management
Internal network	Network segments, IPsec
Perimeter	Firewalls
Physical security	Guards, locks, tracking devices
Policies, procedures, and awareness	Security documents, user education

targets. You might also consider implementing Internet Protocol security (IPsec) for encrypting the traffic between iSCSI initiators and iSCSI targets. Isolating iSCSI traffic to its own virtual LAN (VLAN) also strengthens security by not allowing malicious users that are connected on corporate VLAN network to attack iSCSI storage devices that are connected to a different VLAN. You should also protect network equipment such as routers and switches that are used by iSCSI storage devices, from unauthorized access.

- **Host.** The next layer of defense is the protection layer for the host computers that are connected to iSCSI storage devices. You must maintain secure computers by using the latest security updates. You should consistently use the Windows Update feature in Windows operating systems to keep your operating system up-to-date. You also have to configure security policies such as password complexity, configure the host firewall, and install antivirus software.
- **Application.** Applications are only as secure as their latest security update. For applications that run on your servers but do not integrate in Windows Update, you should regularly check for security updates issued by the application vendor. You should also update the iSCSI management software according to vendor recommendations and best practices.
- **Data.** This is the final layer of security. To help protect your network, ensure that you are using file user permissions properly. Do this by using BitLocker, Access Control Lists (ACLs), implementing the encryption of confidential data with Encrypting File System (EFS), and performing regular backups of data.

Demonstration: Configuring an iSCSI Target

In this demonstration, you will see how to:

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

Demonstration Steps

Add the iSCSI Target Server Role Service

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

Create two iSCSI virtual disks and an iSCSI target

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**. Create a virtual disk with the following settings:
 - Name: **iSCSIDisk1**
 - Disk size: **5 GB**
 - iSCSI target: **New**
 - Target name: **LON-SVR2**
 - Access servers: **172.16.0.22**
3. On the **View results** page, wait until creation completes, and then close the **View Results** page.

4. In the iSCSI VIRTUAL DISKS pane, click **TASKS**, and then in the **TASKS** drop-down list, click **New iSCSI Virtual Disk**. Create a virtual disk that has these settings:
 - Name: **iSCSIDisk2**
 - Disk size: **5 GB**
 - iSCSI target: **LON-SVR2**
5. On the **View results** page, wait until creation completes, and then close the **View Results** page.

Demonstration: Connecting to the iSCSI Storage

In this demonstration, you will see how to:

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

Demonstration Steps

Connect to the iSCSI target

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

Verify the presence of the iSCSI drive

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

Considerations for Implementing iSCSI Storage

When designing your iSCSI storage solution, consider 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.)

Consider the following when designing your iSCSI storage solution:

- 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
- The iSCSI storage solution team must contain 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

- Read the vendor-specific best practices for different types of deployments and applications that will use iSCSI storage solution, such as Exchange Server and Microsoft SQL Server®.
- IT personnel who will 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

Configuring BranchCache

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

Lesson Objectives

After completing this lesson, you will be able to:

- Describe how BranchCache works.
- Describe BranchCache requirements.
- Explain how to configure BranchCache server settings.
- Explain how to configure BranchCache client settings.
- Explain how to configure BranchCache.
- Explain how to monitor BranchCache.

How Does BranchCache Work?


The BranchCache feature introduced with Windows Server 2008 R2 and Windows 7 reduces the network use on WAN connections between branch offices and headquarters by locally caching frequently used files on computers in the branch office.

BranchCache improves the performance of applications that use one of the following protocols:

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

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

BranchCache improves the responsiveness of common network applications that access intranet servers across slow WAN links. Because BranchCache does not require additional infrastructure, you can improve the performance of remote networks by deploying Windows 7 or Windows 8 to client computers, and by deploying Windows Server 2008 R2 and Windows Server 2012 to servers, and then enabling the BranchCache feature.



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

You can configure BranchCache to use hosted cache mode or distributed cache mode:

- **Hosted cache.** This mode operates by deploying a computer that is running Windows Server 2008 R2 or newer versions as a hosted cache server in the branch office. Client computers are locating the host computer so that they can retrieve content from the hosted cache when available. If the content is not available in the hosted cache, the content is retrieved from the content server by using a WAN link and then provided to the hosted cache so that the successive client requests can get it from there.
- **Distributed cache.** For smaller remote offices, you can configure BranchCache in the distributed cache mode without requiring a server. In this mode, local client computers running Windows 7 or Windows 8 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 provide content to other requesting clients.



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

BranchCache functionality in Windows Server 2012 has improved in the following ways:

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

How Client Computers Retrieve Data by Using BranchCache

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

1. The client computer that is running Windows 8 connects to a content server that is running Windows Server 2012 in the head office, and requests content similar to how it would retrieve content without using BranchCache.
2. The content server in the head office authenticates the user and verifies that the user is authorized to access the data.
3. Instead of sending the content itself, the content server in the head office returns identifiers or hashes of the requested content to the client computer. The content server sends that data over the same connection that the content would have typically been sent.
4. Using retrieved identifiers, the client computer does the following:
 - If you configure it to use distributed cache, the client computer multicasts on the local subnet to find other client computers that have already downloaded the content.
 - If you configure it to use hosted cache, the client computer searches for the content on the configured hosted cache.

5. If the content is available in the branch office, either on one or more clients or on the hosted cache, the client computer retrieves the data from the branch office. The client computer also ensures that the data is updated and has not been tampered with or corrupted.
6. If the content is not available in the remote office, then the client computer retrieves the content directly from the server across the WAN link. The client computer then either makes it available on the local network to other requesting client computers (distributed cache mode) or sends it to the hosted cache, where it is made available to other client computers.

BranchCache Requirements

BranchCache optimizes traffic flow between head offices and branch offices. Windows Server 2008 R2, Windows Server 2012, and client computers running Windows 7 and Windows 8 can benefit from using BranchCache. (Earlier versions of Windows operating systems do not benefit from this feature.) You can use BranchCache to cache only the content that is stored on file servers or web servers that are running Windows Server 2008 R2 or Windows Server 2012.

Requirements for using BranchCache	Requirements for the modes
<ul style="list-style-type: none"> • Install the BranchCache feature or the BranchCache for Network Files feature on the server that is hosting the content • Configure client computers by using either Group Policy or the netsh command 	<ul style="list-style-type: none"> • In distributed cache mode, the content cache at a branch office is distributed between client computers. • In hosted cache mode, the content cache at the branch office is hosted on one or more server computers that are hosted cache servers.

Requirements for Using BranchCache

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

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

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

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

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

BranchCache is supported on the full installation and Server Core installation of Windows Server 2012. By default, BranchCache is not installed on Windows Server 2012.

Requirements for Distributed Cache Mode and Hosted Cache Mode

In the Distributed Cache mode, BranchCache works across a single subnet only. If client computers are configured to use the Distributed Cache mode, any client computer can use a multicast protocol called WS-Discovery to search locally for the computer that has already downloaded and cached the content. You should configure the client firewall to enable incoming traffic, HTTP, and WS-Discovery.

In clients, however, they will search for a hosted cache server, and if one is discovered, clients automatically self-configure as hosted cache mode clients. In the Hosted Cache mode, the client computers automatically self-configure as hosted cache mode clients, and they will search for the host server so that they can retrieve content from the Hosted Cache. Furthermore, you can use Group Policy so

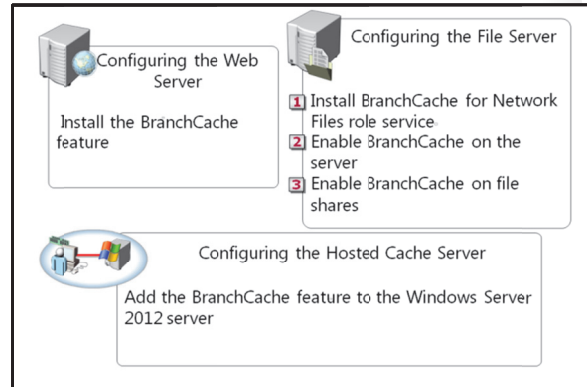
that you can use the FQDN of the hosted cache servers or enable automatic Hosted Cache discovery by Service Connection Points (SCPs). You must configure a firewall to enable incoming HTTP traffic from the Hosted Cache server.

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

Configuring BranchCache Server Settings

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

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



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

Configuring BranchCache Client Settings

You do not have to install the BranchCache feature on client computers, because BranchCache is already included if the client is running Windows 7 or Windows 8. However, BranchCache is disabled by default on client computers. To enable and configure BranchCache, you must perform the following steps:

1. Enable BranchCache.
2. Enable the Distributed Cache mode or the Hosted Cache mode. Windows 8 clients can use either mode dynamically.
3. Configure the client firewall to enable BranchCache protocols.

To enable and configure BranchCache, do the following:

1. Enable BranchCache
2. Enable distributed cache mode or hosted cache mode
3. Configure the client firewall

You can modify BranchCache settings and perform additional configuration tasks, such as:

- Setting the cache size
- Setting the location of the hosted cache server
- Clearing the cache
- Creating and replicating a shared key for using in a server cluster

Enabling BranchCache

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

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

1. Open the Group Policy Management Console.
2. Create a GPO that will be linked to the organizational unit where client computers are located.
3. In a GPO, browse to **Computer Configuration, Policies, Administrative Templates: Policy definitions (ADMX files) retrieved from the local computer, Network**, and then click **BranchCache**.
4. Enable **Turn on BranchCache** setting in GPO.

Enabling the Distributed Cache Mode or Hosted Cache Mode

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

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

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

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

1. Use the following netsh syntax for the Distributed Cache mode:

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

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

```
netsh branchcache set service mode=hostedclient location=<Hosted Cache server>
```

Configuring the Client Firewall to Enable BranchCache Protocols

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

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

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

Additional Configuration Tasks for BranchCache

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

- Setting the cache size.
- Setting the location of the Hosted Cache server.
- Clearing the cache.
- Creating and replicating a shared key for using in a server cluster.

Demonstration: Configuring BranchCache

In this demonstration, you will see how to:

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

Demonstration Steps

Add BranchCache for the Network Files role service

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

Enable BranchCache for the server

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

Enable BranchCache for a file share

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

Monitoring BranchCache


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

You can also use the following tools to monitor BranchCache:

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

The BranchCache monitoring tools include:

- The netsh branchcache show status command
- Event Viewer
- Performance monitor counters



Lesson 3

Optimizing Storage Usage

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

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

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the File Server Resource Manager.
- Describe file classification.
- Describe classification rules.
- Explain how to configure file classification.
- Describe storage optimization options in Windows Server 2012.
- Explain how to configure Data deduplication.

What Is File Server Resource Manager?

You can use the File Server Resource Manager (FSRM) to manage and classify data that is stored on file servers. FSRM includes the following features:

- File classification infrastructure. This feature automates the data classification process. You can dynamically apply access policies to files based on their classification. Example policies include Dynamic Access Control for restricting access to files, file encryption, and file expiration. You can classify files automatically by using file classification rules, or manually by modifying the properties of a selected file or folder. We can modify file properties automatically based on the application, type or content of the file, or by manually setting options on the server that will trigger file classification.
- File management tasks. You can use this feature to apply a conditional policy or action to files, based on their classification. The conditions of a file management task include the file location, the classification properties, the date the file was created, the last modified date of the file, and the last time that the file was accessed. The actions that a file management task can take include the ability to expire files, encrypt files, and run a custom command.

You can use the FSRM to manage and classify data that is stored on file servers

File Server Resource Manager includes the following features:	The new features in File Server Resource Manager include:
<ul style="list-style-type: none"> • File classification infrastructure • File management tasks • Quota management • File screening management • Storage reports 	<ul style="list-style-type: none"> • Dynamic Access Control • Manual classification • Access Denied Assistance • File management tasks • Automatic classification

- Quota management. You can use this feature to limit the space that is allowed for a volume or folder. You can apply quotas automatically to new folders that are created on a volume. You can also define quota templates that you can apply to new volumes or folders.
- File screening management. You can use this feature to control the types of files that users can store on a file server. You can limit the extension that can be stored on your file shares. For example, you can create a file screen that disallows files with an .mp3 extension from being stored in personal shared folders on a file server.
- Storage reports. You can use this feature to identify trends in disk usage, and identify how your data is classified. You can also monitor attempts by users to save unauthorized files.

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

The following FSRM features are new with Windows Server 2012:

- Integration with Dynamic Access Control. Dynamic Access Control can use a file classification infrastructure to help you centrally control and audit access to files on your file servers.
- Manual classification. Manual classification enables users to classify files and folders manually without the need to create automatic classification rules.
- Access Denied Assistance. You can use Access Denied Assistance to customize the access denied error message that displays for users in Windows 8 Consumer Preview when they do not have access to a file or a folder.
- File management tasks. The updates to file management tasks include Active Directory® Domain Services (AD DS) and Active Directory Rights Management Services (AD RMS) file management tasks, continuous file management tasks, and dynamic namespace for file management tasks.
- Automatic classification. The updates to automatic classification increase the level of control you have over how data is classified on your file servers, including continuous classification, Windows PowerShell for custom classification, updates to the existing content classifier, and dynamic namespace for classification rules.



Additional Reading: For more information about FSRM, see:
<http://technet.microsoft.com/en-us/library/hh831746.aspx>

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

What Is File Classification?

File classifications enable administrators to configure automatic procedures for defining a desired property on a file, based on conditions specified in classification rules. For example, you can set the **Confidentiality** property to **High** on all documents whose content contains the word “**secret**.”

In Windows Server 2008 R2 and Windows Server 2012, classification management and file management tasks enable administrators to manage groups of files based on various file and folder attributes. You can automate file and folder maintenance tasks, such as cleaning up stale data, or protecting sensitive information. For this reason, file and folder maintenance tasks are more efficient as compared to maintaining the file system by navigating through its hierarchical view.

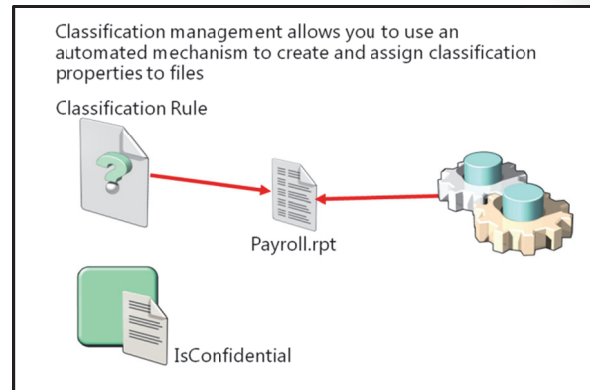
Classification management is designed to ease the burden and management of data that is spread across the organization. You can classify files in a variety of ways. In most scenarios, classification is performed manually. The file classification infrastructure in Windows Server 2012 enables organizations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

You can use file classification to perform the following actions:

1. Define classification properties and values, which you can assign to files by running classification rules.
2. Create, update, and run classification rules. Each rule assigns a single predefined property and value to files within a specified directory, based on installed classification plug-ins.
3. When running a classification rule, reevaluate files that are already classified. You can choose to overwrite existing classification values, or add the value to properties that support multiple values. You can also use classification rules to declassify files that are not in the classification criterion anymore.

What Are Classification Rules?

The file classification infrastructure uses classification rules to scan files automatically, and then to classify them according to the contents of a file. You configure file classifications in the File Server Resource Manager console. Classification properties are defined centrally in AD DS so that these definitions can be shared across file servers within the organization. You can create classification rules that scan files for a standard string, or for a string that matches a pattern (regular expression). When a configured classification parameter is found in a file, that file is classified as configured in the classification rule.



The file classification infrastructure scan files automatically, and then classifies them according to the contents of a file

When planning for file classification implementation, do following:

1. Identify classifications
2. Determine classification method
3. Determine schedule
4. Perform review

When planning for file classifications, you should do following:

1. Identify which classification or classifications that you want to apply on documents.
2. Determine the method you to want to use to identify documents for classification.
3. Determine the schedule for automatic classifications.
4. Establish a review of classification success.

After you have a defined the classifications, you can plan the Dynamic Access Control implementation by defining conditional expressions that enable you to control access to highly confidential documents based on particular user attributes.

Demonstration: Configuring File Classification

In this demonstration, you will see how to:

- Create a classification property.
- Create a classification rule.

Demonstration Steps

Create a classification property

1. On LON-SVR1, the Server Manager console should open automatically. From Server Manager, start the **File Server Resource Manager**.
2. In File Server Resource Manager, create a local property with the following settings:
 - o Name: Corporate Documentation
 - o Property Type: **Yes/No**
3. In File Server Resource Manager, create a classification rule with the following settings:
 - o General tab, Rule name: **Corporate Documents Rule**, and ensure that the rule is enabled.
 - o Scope tab: E:\Labfiles\Corporate Documentation
 - o Classification tab, Classification method: Folder Classifier
 - o Property-Choose a property to assign to files: **Corporate Documentation**
 - o Property-Specify a value: **Yes**.
 - o Evaluation type tab, Re-evaluate existing property values, and Aggregate the values.
4. Run the classification with all rules, and select Wait for classification to complete.
5. Review the **Automatic classification report** that displays in Windows Internet Explorer®, and ensure that report lists the same number of files classified as in the **Corporate Documentation** folder.

Options for Storage Optimization in Windows Server 2012

Windows Server 2012 includes new options for storage optimization that provide you with an efficient way to deploy, administer, and secure your storage solutions. Storage optimization features include:

- File access auditing. File access auditing in Windows Server 2012 creates an audit event whenever files are accessed by users. 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. Features on Demand enables you to save on disk space by allowing you to remove role and feature files from the operating system. If these roles and features need to be installed on the server, the installation files will be retrieved from remote locations, installation media, or Windows Update. You can remove feature files from both physical and virtual computers, Windows image (.wim) files, and offline virtual hard disks (VHDs).
- 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 the server workloads is maintained by throttling the CPU and memory resources that are consumed. Using Data deduplication jobs, you can schedule when Data deduplication should run, specify the resources to deduplicate, and fine-tune file selection. When combined with BranchCache, the same optimization techniques are applied to data that is transferred over the WAN to a branch office. This results in faster file download times, and reduced bandwidth consumption.
- NFS Data Store. The Network file system (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, so that the client can still connect to the NFS server.

Storage optimization features include:

- File access auditing
- Features on Demand
- Data deduplication
- NFS data stores

Demonstration: Configuring Data Deduplication

In this demonstration, you will see how to:

- Add the Data deduplication role service.
- Enable Data deduplication.

Demonstration Steps

Add the Data deduplication role service

1. Log on to **LON-SVR1** as **Adatum\Administrator** using the password **Pa\$\$w0rd**.
2. Open Server Manager.

3. In the Add Roles and Features Wizard, install the following roles and features to the local server, and accept the default values:
 - **File And Storage Services (Installed)\File and iSCSI Services\Data Deduplication**

Enable Data deduplication

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

Lab A: Implementing Advanced File Services

Scenario

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

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

Objectives

- Configure iSCSI storage.
- Configure the file classification infrastructure.

Lab Setup

Estimated Time: 75 minutes

20412A-LON-DC1

20412A-LON-SVR1

20412A-LON-SVR2

Estimated time: **75 minutes**

Virtual Machine(s)	20412A-LON-DC1 20412A-LON-SVR1 20412A-LON-SVR2
User Name	Adatum\Administrator
Password	Pa\$\$w0rd

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

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, click **20412A-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on using the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
5. Repeat steps 2-4 for **20412A-LON-SVR1** and **20412A-LON-SVR2**.

Exercise 1: Configuring iSCSI Storage

Scenario

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

The main tasks for this exercise are as follows:

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

► Task 1: Install the iSCSI target feature

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

► Task 2: Configure the iSCSI targets

1. On LON-DC1, in Server Manager, in the navigation pane, click **File and Storage Services**, and then click **iSCSI**.
2. Create a virtual disk with the following settings:
 - Storage location: **C:**
 - Disk name: **iSCSIDisk1**
 - Size: **5 GB**
 - iSCSI target: **New**
 - Target name: **lon-dc1**
 - Access servers: **172.16.0.22** and **131.107.0.2**
3. On the **View results** page, wait until the creation completes, and then click **Close**.
4. Create a New iSCSI virtual disk with these settings:
 - Storage location: **C:**
 - Disk name: **iSCSIDisk2**
 - Size: **5 GB**
 - iSCSI target: **lon-dc1**
5. Create a New iSCSI virtual disk with these settings:
 - Storage location: **C:**
 - Disk name: **iSCSIDisk3**
 - Size: **5 GB**
 - iSCSI target: **lon-dc1**

6. Create a New iSCSI virtual disk with these settings:
 - o Storage location: **C:**
 - o Disk name: **iSCSIDisk4**
 - o Size: **5 GB**
 - o iSCSI target: **lon-dc1**
7. Create a New iSCSI virtual disk with these settings:
 - o Storage location: **C:**
 - o Disk name: **iSCSIDisk5**
 - o Size: **5 GB**
 - o iSCSI target: **lon-dc1**

► **Task 3: Configure MPIO**

1. On LON-SVR2, from Server Manager, open the **Routing and Remote access** console.
2. On the Enable DirectAccess Wizard, click **Cancel**.
3. Right-click **LON-SVR2** and then click **Disable Routing and Remote Access**. Click **Yes** and then close the Routing and Remote Access console.
4. In Server Manager, start the Add Roles and Features Wizard and install the **Multipath I/O** feature.
5. In Server Manager, on the **Tools** menu, open **iSCSI Initiator**, and configure the following:
 - o Enable the **iSCSI Initiator** service
 - o **Quick Connect** to target: **LON-DC1**
6. In Server Manager, on the **Tools** menu, open **MPIO**, and configure the following:
 - o Enable **Add support for iSCSI devices** on Discover Multi-paths
7. After the computer restarts, log on to **LON-SVR2**, with username of **Adatum\Administrator** and password of **Pa\$\$w0rd**.
8. In Server Manager, on the **Tools** menu, click **MPIO** and verify that **Device Hardware ID MSFT2005iSCSIBusType_0x9** is added to the list.

► **Task 4: Connect to and configure the iSCSI targets**

1. On LON-SVR2, in Server Manager, on the **Tools** menu, open **iSCSI Initiator**.
2. In the **iSCSI Initiator Properties** dialog box, perform the following steps:
 - o **Disconnect** all **Targets**.
 - o **Connect** and **Enable multi-path**.
 - o Set **Advanced** options as follows:
 - Local Adapter: **Microsoft iSCSI Initiator**
 - Initiator IP: **172.16.0.22**
 - Target Portal IP: **172.16.0.10 / 3260**
 - o **Connect** to another target, **enable multi-path**, and configure the following **Advanced** settings:
 - Local Adapter: **Microsoft iSCSI Initiator**
 - Initiator IP: **131.107.0.2**

- Target Portal IP: **131.107.0.1 / 3260**
3. In the **Targets** list, open **Devices** for **iqn.1991-05.com.microsoft:lon-dc1-lon-dc1-target**, access the MPIO information, and then verify that in **Load balance policy, Round Robin** is selected. Verify that two paths are listed by looking at the IP addresses of both network adapters.

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

Exercise 2: Configuring the File Classification Infrastructure

Scenario

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

The main tasks for this exercise are as follows:

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

► Task 1: Create a classification property for corporate documentation

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

► Task 2: Create a classification rule for corporate documentation

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


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

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

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

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

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

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

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

► **Task 5: Verify that corporate documents are expired**

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

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

► **To prepare for the next lab**

When you finish the lab, revert 20417A-LON-SVR2. To do this, complete the following steps.

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

Keep all other virtual machines running for the next lab.

Lab B: Implementing BranchCache

Scenario

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

Objectives

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

Lab Setup

Estimated Time: 30 Minutes

20412A-LON-DC1

20412A-LON-SVR1

20412A-LON-CL1

20412A-LON-CL2

Estimated time: **30 minutes**

Virtual Machine(s)	20412A-LON-DC1 20412A-LON-SVR1 20412A-LON-CL1 20412A-LON-CL2
User Name	Adatum\Administrator
Password	Pa\$\$w0rd

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

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, click **20412A-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:
 - o User name: **Adatum\Administrator**
 - o Password: **Pa\$\$w0rd**
5. Repeat steps 2-4 for **20412A-LON-SVR1**, **20412A-LON-CL1**, and **20412A-LON-CL2**.

Exercise 1: Configuring the Main Office Servers for BranchCache

Scenario

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

The main tasks for this exercise are as follows:

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

► **Task 1: Configure LON-DC1 to use BranchCache**

1. Switch to LON-DC1.
2. Open Server Manager, and install the **BranchCache for network files** role service.
3. Open the Local Group Policy Editor (**gpedit.msc**).
4. Navigate to and open **Computer Configuration/Administrative Templates/Network/Lanman Server/Hash Publication for BranchCache**.
5. Enable the BranchCache setting, and then select **Allow hash publication only for shared folders on which BranchCache is enabled**.

► **Task 2: Simulate a slow link to the branch office**

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

► **Task 3: Enable a file share for BranchCache**

1. In a Windows Explorer window, create a new folder named **C:\Share**.
2. Share this folder with the following properties:
 - Share name: **Share**
 - Permissions: default
 - Caching: **Enable BranchCache**
3. Copy **C:\Windows\System32\write.exe** to the **C:\Share** folder.

► **Task 4: Configure client firewall rules for BranchCache**

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

- d. Rule type: **predefined**
 - e. Use **BranchCache – Peer Discovery (Uses WSD)**
 - f. Action: **Allow**
6. Close the Group Policy Management Editor and Group Policy Management console.

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

Exercise 2: Configuring the Branch Office Servers for BranchCache

Scenario

The next step you must perform is to configure a file server for the BranchCache feature. You will install the BranchCache feature, request a certificate, and then link it to BranchCache.

The main tasks for this exercise are as follows:

1. Install the BranchCache feature on LON-SVR1.
2. Start the BranchCache host server.

► Task 1: Install the BranchCache feature on LON-SVR1

- On LON-SVR1, from Server Manager, add the **BranchCache for Network Files** role service and the **BranchCache** feature.

► Task 2: Start the BranchCache host server

1. On LON-DC1, open Active Directory Users and Computers, and create a new organizational unit (OU) called **BranchCacheHost**.
2. Move **LON-SVR1** into the **BranchCacheHost** OU.
3. Open Group Policy Management, and block GPO inheritance on the **BranchCacheHost** OU.
4. Restart **LON-SVR1** and log on as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
5. On LON-SVR1, open Windows PowerShell, and run the following cmdlet:

```
Enable-BCHostedServer
-RegisterSCP
```

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

```
Get-BCStatus
```

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

Exercise 3: Configuring Client Computers for BranchCache

Scenario

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

The main task for this exercise is as follows:

1. Configure client computers to use BranchCache in Hosted Cache mode

► **Task 1: Configure client computers to use BranchCache in Hosted Cache mode**

1. On LON-DC1, open Server Manager, and then open Group Policy Management.
2. Edit the Default Domain Policy.
3. In the Group Policy Management Editor, browse to **Computer Configuration\Policies\Administrative Templates\Network\BranchCache**, and configure the following:
 - Turn on BranchCache: **Enabled**
 - Enable Automatic Hosted Cache Discovery by Service Connection Point: Enabled
 - Configure BranchCache for network files: Enabled
 - **Type the maximum round trip network latency (milliseconds) after which caching begins:**
0
4. Start **20412A-LON-CL1**, open a command prompt window, and refresh the Group Policy settings using the command **gpupdate /force**.
5. At the command prompt, type **netsh branchcache show status all**, and then press Enter.
6. Start the **20412A-LON-CL2**, open the command prompt window, and refresh the Group Policy settings using the command **gpupdate /force**.
7. At the command prompt, type **netsh branchcache show status all**, and then press Enter.

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

Exercise 4: Monitoring BranchCache

Scenario

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

The main tasks for this exercise are as follows:

1. Configure Performance Monitor on LON-SVR1.
2. View performance statistics on LON-CL1.
3. View performance statistics on LON-CL2.
4. Test BranchCache in the Hosted Cache mode.

► **Task 1: Configure Performance Monitor on LON-SVR1**

1. On LON-SVR1, open Performance Monitor.
2. In the Performance Monitor console, in the navigation pane, under **Monitoring Tools**, click **Performance Monitor**.
3. Remove existing counters, change to report view, and then add the **BranchCache** object counters to the report.

► **Task 2: View performance statistics on LON-CL1**

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

► **Task 3: View performance statistics on LON-CL2**

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

► **Task 4: Test BranchCache in the Hosted Cache mode**

1. Switch to LON-CL1.
2. Open `\\LON-DC1.adatum.com\share`, and copy the executable file to the local desktop. This could take several minutes because of the simulated slow link.
3. Read the performance statistics on LON-CL1. This file was retrieved from LON-DC1 (Retrieval: Bytes from Server). After the file was cached locally, it was passed up to the hosted cache. (Retrieval: Bytes Served).
4. Switch to LON-CL2.
5. Open `\\LON-DC1.adatum.com\share`, and copy the executable file to the local desktop. This should not take as long, because the file is cached.
6. Read the performance statistics on LON-CL2. This file was obtained from the hosted cache (Retrieval: Bytes from Cache).
7. Read the performance statistics on LON-SVR1. This server has offered cached data to clients (Hosted Cache: Client file segment offers made).

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

► **To prepare for the next module**

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

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20412A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20412A-LON-SVR1**, **20412A-LON-CL1**, and **20412A-LON-CL2**.

Module Review and Takeaways

Question: How does BranchCache differ from the Distributed File System?

Question: Why would you want to implement BranchCache in Hosted Cache mode instead of Distributed Cache mode?

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

Question: Why would you implement a file classification infrastructure?

Real-world Issues and Scenarios

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

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

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

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

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

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

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



Best Practice:

- When considering an iSCSI storage solution for your organization, spend most of the time on the design process. The design process is crucial because it allows you to optimize the solution for all

technologies that will be using iSCSI storage, such as file services, Exchange Server, and SQL Server. The design should also accommodate future growth of your organizations' business data. Successful design processes guarantee a successful deployment of the solution that will meet your organization's business requirements.

- When planning for BranchCache deployment, ensure that you work closely with your network administrators so that you can optimize network traffic across the WAN.
- When planning for file classifications, ensure that you start with your organization's business requirements. Identify the classifications that you will apply to documents, and then define a method that you will use to identify documents for classification. Before you deploy the file classification infrastructure, create a test environment and test the scenarios to ensure that your solution will result in a successful deployment and that your organizations' business requirements will be met.

Tools

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

Module 3

Implementing Dynamic Access Control

Contents:

Module Overview	3-1
Lesson 1: Overview of Dynamic Access Control	3-2
Lesson 2: Planning for Dynamic Access Control	3-8
Lesson 3: Deploying Dynamic Access Control	3-13
Lab: Implementing Dynamic Access Control	3-22
Module Review and Takeaways	3-30

Module Overview

The Windows Server® 2012 operating system introduces a new feature for enhancing access control for file-based and folder-based resources. This feature, called Dynamic Access Control, extends regular NTFS file system–based access control by enabling administrators to use claims, resource properties, policies, and conditional expressions to manage access. In this module, you will learn about Dynamic Access Control, and how to plan for it and implement it.

Objectives

After completing this module, you will be able to:

- Describe Dynamic Access Control and its components.
- Plan for Dynamic Access Control implementation.
- Deploy Dynamic Access Control.

Lesson 1

Overview of Dynamic Access Control

Dynamic Access Control is a new Windows Server 2012 feature that you can use for access management. Dynamic Access Control offers a new way of securing and controlling access to resources. Before you implement this feature, you should understand how it works and which components it uses. This lesson presents an overview of Dynamic Access Control.

Lesson Objectives

After completing this lesson, you will be able to:

- Define Dynamic Access Control.
- Describe the foundation technologies for Dynamic Access Control.
- Compare Dynamic Access Control with alternative and similar technologies such as NTFS permissions and Active Directory Rights Management Services (AD RMS).
- Define identity.
- Define claim and claim types.
- Define a central access policy.

What Is Dynamic Access Control?

Typically, most of an organization's data is stored on file servers. Therefore, IT administrators must provide proper security and access control to file server resources. In previous versions of Windows Server, IT administrators controlled most access to file server resources by using NTFS permissions and access control lists.

Dynamic Access Control in Windows Server 2012 is a new access control mechanism for file system resources. It enables administrators to define central file access policies that can apply to every file server in the organization. Dynamic Access Control implements a safety net over file servers, and over any existing share and NTFS permissions. It also ensures that regardless of how the share and NTFS permissions might change, this central overriding policy is still enforced. Dynamic Access Control combines multiple criteria into the access decision; this is something that NTFS permissions cannot do.

Dynamic Access Control provides a flexible way to apply and manage access and auditing to domain-based file servers. Dynamic Access Control uses claims in the authentication token, resource properties on the resource, and conditional expressions within permission and auditing entries. With this combination of features, you can now grant access to files and folders based on Active Directory® Domain Services (AD DS) attributes.

Dynamic Access Control provides:

- Data identification. You can use automatic and manual file classification to tag data in file servers across the organization.

Dynamic Access Control provides:

- A safety net over all file server-based resources
- Data identification
- Access control to files
- File access auditing
- Optional RMS protection integration

- Access control to files. Central access policies enable organizations to define, for example, who can access health information within an organization.
- Auditing of file access. You can use central audit policies for compliance reporting and forensic analysis. For example, you can identify who accessed highly sensitive information.
- Optional RMS protection integration. You can use Rights Management Services (RMS) encryption for sensitive Microsoft® Office documents. For example, you can configure RMS to encrypt all documents containing Health Insurance Portability and Accountability Act (HIPAA) information.

Dynamic Access Control is designed for four main end-to-end scenarios:

- Central access policy for access to files. Enables organizations to set safety net policies that reflect business and regulatory compliance.
- Auditing for compliance and analysis. Enables targeted auditing across file servers for compliance reporting and forensic analysis.
- Protecting sensitive information. Identifies and protects sensitive information within a Windows Server 2012 environment, and when it leaves the Windows Server 2012 environment.
- Access denied remediation. Improves the access-denied experience to reduce help desk load and incident time for troubleshooting.

Foundation Technologies for Dynamic Access Control

Dynamic Access Control combines many Windows Server 2012 technologies to provide a flexible and granular authorization and auditing experience. Dynamic Access Control uses the following technologies:

- Network protocols, such as TCP/IP, Remote Procedure Call (RPC), Server Message Block (SMB), and Lightweight Directory Access Protocol (LDAP), for network communications between hosts, and interaction with file system and directory lookups.
- Domain Name System (DNS) for host name resolution.
- AD DS and its dependent technologies for enterprise network management.
- The Kerberos version 5 protocol, including FAST Search and Compound Identity for secure authentication.
- Windows Security (local security authority (LSA), Net Logon service) for secure logon transactions.
- File classifications for file categorization.
- Auditing for secure monitoring and accountability.

Dynamic Access Control relies on many technologies in Windows Server 2012 such as:

- Network protocols
- DNS
- AD DS
- Kerberos
- File classifications
- Auditing

Several components and technologies are updated in Windows Server 2012 to support Dynamic Access Control. The most important updates are:

- A new Windows authorization and audit engine that can process conditional expressions and central policies.
- Kerberos authentication support for user claims and device claims.

- Improvements to the file classification infrastructure.
- Optional RMS extensibility support so that partners can provide solutions for encrypting files that are not Microsoft Office files.

Dynamic Access Control vs. Alternative Permissions Technologies

Dynamic Access Control controls access to file-based resources. It does not overlap with older, well-known technologies that provide similar functionality. Instead, Dynamic Access Control extends the functionality of older technologies for controlling file-based resource access.

In previous versions of Windows Server, the basic mechanism for file and folder access control was NTFS permissions. By using NTFS permissions and their access control lists (ACLs), administrators can control access to resources based on user name security identifiers (SIDs) or group membership SIDs, and the level of access such as Read-only, Change, and Full Control. However, once you provide someone with, for example, Read Only access to a document, you cannot prevent that person from copying the content of that document into a new document or printing the document.

By implementing AD RMS, you can establish an additional level of file control. Unlike, NTFS permissions, which are not application-aware, AD RMS sets a policy that can control document access inside the application that the user uses to open it. By implementing AD RMS, you enable users to protect documents within applications.

However, you cannot set conditional access to files by using NTFS and AD RMS. For example, you cannot set NTFS permissions so that users can access documents if they are members of specific groups and have their **EmployeeType** attributes set to **Full Time Employee (FTE)**. Additionally, you cannot set permissions so that only users who have a department attribute populated with the same value as the department attribute for the resource can access the content. However, you can use conditional expressions to accomplish these tasks. You can use Dynamic Access Control to count attribute values on users or resource objects when providing or denying access.

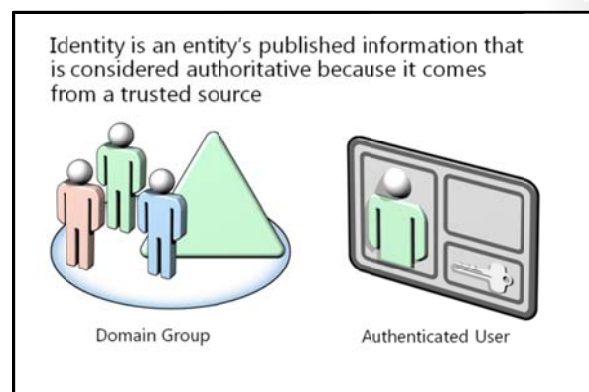
Various permissions technologies exist:

- NTFS permissions and ACLs provide access control that is based on user's SID or group membership SID
- AD RMS provides greater protection for documents by controlling how applications can use them
- Dynamic Access Control provides access control based on values of specific attributes

What Is Identity?

Identity is usually defined as a set of data that uniquely describes a person or a thing (sometimes referred to as *subject* or *entity*) and contains information about the subject's relationships to other entities. Identity is usually verified by using a trusted source of information.

For example, when you go to the airport, you show your passport. Your passport contains your name, address, date of birth, and photograph. Each item of personal information is a claim that is made about you by the country issuing your passport. Your country ensures that the



information that is published in a passport is accurate for the passport owner. Because you usually use the passport outside of your country of residence, other countries must also trust the information in your passport. They must trust the organization that issued your passport and consider it reliable. Based on that trust, other countries grant you access to their territories (which can be considered resources). Therefore, in this example, to access resources in other countries, each person is required to have a document (passport) that is issued by a reliable and trusted source and that contains critical claims that describe the person.

The Windows Server operating system uses a similar concept of identity. An administrator creates a user account in AD DS for a person. The domain controller publishes user account information, such as a security identifier and group membership attributes. When a user accesses a resource, Windows Server creates an authorization token.

To continue the foreign travel analogy, you are the user, and the authorization token is the passport. Each unique piece of information in the authorization token is a claim made about your user account. Domain controllers publish these claims. Domain-joined computers and domain users trust domain controllers to provide authoritative information.

We can then say that identity, with respect to authentication and authorization, is information published about an entity from a trusted source. Furthermore, the information is considered authoritative because the source is trusted.

Earlier versions of Windows Server used the security identifier (SID) to represent the identity of a user or computer. Users authenticate to the domain with a specific user name and password. The unique logon name translates into the SID. The domain controller validates the password and publishes the SID of the security principal and the SIDs of all the groups within which the principal is a member. The domain controller claims the user's SID is valid and should be used as the identity of the user. All domain members trust their domain controller; therefore, the response is treated as authoritative.

Identity is not limited to the user's SID. Applications can use any information about the user as a form of identity, if the application trusts that the source of the information is authoritative. For example, many applications implement role based access control (RBAC). RBAC limits access to resources based on whether the user is a member of a specific role. Microsoft SharePoint® Server is a good example of software that implements role-based security. Windows Server 2012 can also take advantage of these options to extend and enhance the way identity is determined for a security principal.

What Is a Claim?

Windows Server 2008 and Windows Server 2003 use claims in Active Directory Federation Services (AD FS). In this context, claims are statements made about users—for example, name, identity, key, group, privilege, or capability—which are understood by the partners in an AD FS federation. AD FS also provides AD DS–based claims, and the ability to convert the data from these claims into Security Assertions Markup Language (SAML) format. In previous versions of AD FS, the only attributes that could be retrieved from AD DS and incorporated directly into a claim was SID information for user and group accounts. All other claim information was defined within and referenced from a separate database, known as an *attribute store*. Windows Server 2012 now allows you to read and use any attribute directly from AD DS. You do not need to use a separate AD FS attribute store to hold this type of information for Active Directory–based computer or user accounts.

Claims are statements made by AD DS about specific users and computer objects in AD DS

AD DS in Windows Server 2012 supports:

- User claims
- Device claims

By definition, a *claim* is something that AD DS states about a specific object (usually a user or computer). A claim provides information from the trusted source about an entity. Some examples of claims are the SID of a user or computer, the department classification of a file, and the health state of a computer. All these claims state something about a specific object.

An entity normally contains more than one claim. When configuring resource access, any combination of claims can be used to authorize access to resources.

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

User Claim

A *user claim* is information that is provided by a Windows Server 2012 domain controller about a user. Windows Server 2012 domain controllers can use most AD DS user attributes as claim information. This provides administrators with a wide range of possibilities for configuring and using claims for access control.

Device Claim

A *device claim*—which is often called a computer claim—is information that is provided by a Windows Server 2012 domain controller about a device that is represented by a computer account in AD DS. As with user claims, device claims can use most of the AD DS attributes that are applicable to computer objects.

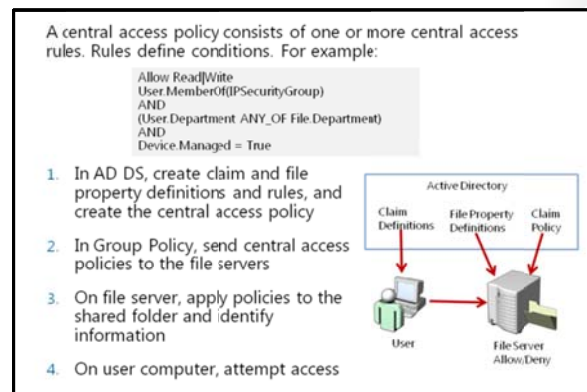
What Is a Central Access Policy?

One of the fundamental components of Dynamic Access Control is the central access policy. This Windows Server 2012 feature enables administrators to create policies that they can apply to one or more file servers. You create policies in the Active Directory Administrative Center, which then stores them in AD DS, and you then apply them by using Group Policy. The central access policy contains one or more central access policy rules. Each rule contains settings that determine applicability and permissions.

Before you create a central access policy, you must create at least one central access rule. Central access rules define all parameters and conditions that control access to specific resources.

A central access rule has three configurable parts:

- Name. For each central access rule, you should configure a descriptive name.
- Target resource. This is a condition that defines to which data the policy applies. You define a condition by specifying an attribute and its value. For example, a particular central policy rule might apply to any data that you classify as sensitive. You can also apply the rule to all resources to which the central access policy applies.



- **Permissions.** This is a list of one or more access control entries (ACEs) that define who can access data. For example, you can specify Full Control Access to a user with attribute **EmployeeType** set to **FTE** (full-time employee). This is the key component of each central access rule. You can combine and group conditions that you place in the central access rule. You can set permissions either to **proposed** (for staging purposes) or **current**.

After you configure one or more central access rules, you then add these rules to the central access policy, which is then applied to the resources.

The central access policy enhances, but does not replace, the local access policies or discretionary access control lists (DACLS) that are applied to files and folders on a specific server. For example, if a DACL on a file allows access to a specific user, but a central access policy that is applied to the file restricts access to the same user, the user will not be able to obtain access to the file. Likewise, if the central access policy allows access but the DACL does not allow access, then the user will not be able to access the file.

Before implementing the central access policy, perform these steps:

1. Create a claim, and then connect it to users or computer objects by using attributes.
2. Create file property definitions.
3. Create one or more central access rules.
4. Create a Central Access Policy object and define its settings.
5. Use Group Policy to deploy the policy to file servers. By doing this, you make file servers aware that a central access policy exists in AD DS.
6. On the file server, apply that policy to a specific shared folder.

Lesson 2

Planning for Dynamic Access Control

Dynamic Access Control requires detailed planning prior to implementation. You should identify reasons for implementing Dynamic Access Control, and plan for central access policy, file classifications, auditing, and access-denied assistance. In this lesson, you will learn about planning Dynamic Access Control.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe reasons for implementing Dynamic Access Control.
- Plan for central access policy.
- Plan for file classifications.
- Plan for file access auditing.
- Plan for access-denied assistance.

Reasons for Implementing Dynamic Access Control

Before you implement Dynamic Access Control, you should clearly identify the reasons for using this feature. Dynamic Access Control should be well designed before you implement it. An improperly planned implementation can cause some users to be denied access to required data, and other users to be granted access to restricted data.

The most common reason to implement Dynamic Access Control is to extend functionality of an existing access control model. Most organizations use NTFS and share permissions to implement access control for file and folder resources. In most cases, NTFS is sufficient, but in some scenarios, it does not work. For example, you cannot use NTFS ACL to protect a resource on a file server, which means that a user must simultaneously be a member of two groups to access the resource. You must use Dynamic Access Control instead of traditional methods for implementing access control when you want to use more specific information for authorization purposes. NTFS and share permissions use only user or group objects.

The most common reasons for implementing Dynamic Access Control are:

- An inability to achieve the desired results with NTFS
- A requirement for access control based on attributes

Planning for Central Access Policy

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

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

1. Identify the resources that you want to protect. If all these resources are on one file server or in just one folder, then you might not have to implement a central access policy. Instead, you can configure conditional access on the folder's ACL. However, if resources are distributed across several servers or folders, then you may benefit from deploying a central access policy. Data that might require additional protection may include payroll records, medical history data, employee personal information, and company customer lists. You can also use targeting within central access rules to identify resources to which you want to apply central access policy.
2. Define the authorization policies. These policies are usually defined from your business requirements. Some examples are:
 - o All documents that have property confidentiality set to high must be available only to managers.
 - o Marketing documents from each country should be writable only by marketing people from the same country.
 - o Only full time employees should be able to access technical documentation from previous projects.
3. Translate the authorization policies that you require into expressions. In the case of Dynamic Access Control, expressions are attributes that are associated with both the resources (files and folders) and the user or device that seeks access to the resources. These expressions state additional identification requirements that must be met to access protected data. Values that are associated with any expressions on the resource obligate the user or device to produce the same value.
4. Next, you should break down the expressions that you have created, and determine what claim types, resource properties, and device claims that you must create to deploy your policies. In other words, you must identify the attributes for access filtering.

When planning a central access policy, you should:

- Identify the business case
- Identify the resources to be protected
- Define the authorization policies as defined by your business requirements
- Translate the authorization policies into conditional expressions
- Define claim types, resource properties, and rules



Note: You must use user claims to deploy central access policies. You can use security groups to represent user identities.

Planning File Classifications

When planning your Dynamic Access Control implementation, you should include file classifications. Although file classifications are not mandatory for Dynamic Access Control, they can enhance the automation of the entire process. For example, if you require that all documents that are classified Confidentiality: High be accessible to top management only, regardless of the server on which the documents exist, you should ask yourself how you identify these documents, and how to classify them appropriately.

The file classification infrastructure uses classification rules to scan files automatically, and then classify them according to the contents of the file. **Classification** and **Resource** properties are defined centrally in AD DS so that these definitions can be shared across file servers in the organization. You can create classification rules that scan files for a standard string or for a string that matches a pattern (regular expression). When a configured classification parameter is found in a file, that file is classified as configured in the classification rule.

When planning for file classifications, do the following:

- Identify which classification or classifications that you want to apply to documents.
- Determine the method that you will use to identify documents for classification.
- Define the schedule for automatic classifications.
- Establish periodic reviews to determine the success of the classifications.

You configure file classifications in the File Server Resource Manager console.

Once you have defined the classifications, you can plan the implementation of Dynamic Access Control by defining conditional expressions which will enable you to control access to high confidential documents based on particular user attributes.

Planning File Access Auditing

In Windows Server 2008 R2 and Windows Server 2012, you can use advanced audit policies to implement detailed and more precise file system auditing. In Windows Server 2012, you can also implement auditing together with Dynamic Access Control to utilize the new Windows security auditing capabilities. By using conditional expressions, you can configure auditing so that it only occurs in specific cases. For example, you may want to audit attempts to open shared folders by users in countries other than the country where the shared folder is located. You achieve this by implementing proposed permissions in the central access rules.

With Global Object Access auditing, administrators can define computer system access control lists (SACLs) according to the object type for either the file system or registry. The specified SACL is then applied automatically to every object of that type. You can use a Global Object Access Audit policy to

When planning for file classification, you should:

- Identify the classifications
- Determine the method you will use to classify the files
- Define the schedule
- Perform reviews

File access auditing:

- Tracks changes to user and machine attributes
- Retrieves more information from user logon events
- Provides more information from object access auditing
- Tracks changes to central access policies, central access rules, and claims
- Tracks changes to file attributes

enforce the Object Access Audit policy for a computer, file share, or registry without configuring and propagating conventional SACLs. Configuring and propagating a SACL is a complex administrative task that is difficult to verify, particularly if you must verify to an auditor that a security policy is being enforced.



Note: Auditors can verify that every resource in the system is protected by an audit policy by viewing the contents of the Global Object Access Auditing policy setting.

Resource SACLs are also useful for diagnostic scenarios. For example, setting a Global Object Access Audit policy to log all activity for a specific user and enabling the Access Failures audit policies in a resource such as a file system or registry can help administrators quickly identify which object in a system is denying a user access.

Before you implement auditing you should prepare an audit plan. In the auditing plan, you should identify resources, users, and activities that you want to track. You can implement auditing for several scenarios, such as:

- Tracking changes to user and machine attributes. As with files, users and machine objects can have attributes, and changes to these can affect whether users can access files. Therefore, tracking changes to user or machine attributes can be valuable. Users and machine objects are stored in AD DS, which means that you can track their attributes using Directory Service Access auditing.
- Obtaining more information from user logon events. In Windows Server 2012, a user logon event (4624) contains information about the attributes of the file that was accessed. You can view this additional information by using audit log management tools to correlate user logon events with object access events, and by enabling event filtering based on both file attributes and user attributes.
- Providing more information from object access auditing. In Windows Server 2008 R2 and Windows Server 2012, file access events 4656 and 4663 now contain information about the attributes of the file that was accessed. Event log filtering tools can use this additional information to help you identify the most relevant audit events.
- Tracking changes to central access policies, central access rules, and claims. Because these objects are stored in AD DS, you can audit them just as you would any other securable object in AD DS by using Directory Service Access auditing.
- Tracking changes to file attributes. File attributes determine which central access policy applies to the file. A change to the file attributes can potentially affect the access restrictions on the file. You can track changes to file attributes on any machine by configuring Authorization Policy Change auditing and Object Access auditing for file systems. Event 4911 is introduced in Windows Server 2012 to differentiate this event from other Authorization policy change events.

Planning Access Denied Assistance

Access Denied Assistance helps end users determine why they cannot access a resource. It also allows IT staff to properly diagnose a problem, and then direct the resolution. Windows Server 2012 enables you to customize messages about denied access and provide users with the ability to request access without contacting the help desk or IT team. In combination with Dynamic Access Control, Access Denied Assistance can inform the file administrator of user and resource claims, enabling the administrator to make educated decisions about how to adjust policies or fix user attributes (for example, if the department is listed as HR instead of Human Resources).

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

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

When planning for Access Denied Assistance, consider:

- The message that users will view
- The email text that users will use to request access
- The recipients for access request email messages
- The target operating systems

Lesson 3

Deploying Dynamic Access Control

To deploy Dynamic Access Control, you must perform several steps and configure several objects. In this lesson, you will learn about implementing and configuring Dynamic Access Control.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the prerequisites for implementing Dynamic Access Control.
- Enable support in AD DS for Dynamic Access Control.
- Implement claims and resource property objects.
- Implement central access rules and policies.
- Implement file access auditing.
- Implement Access Denied assistance.
- Implement file classifications.
- Implement Dynamic Access Control.

Prerequisites for Implementing Dynamic Access Control

Before implementing Dynamic Access Control, you must ensure that your servers meet certain prerequisites. Claims-based authorization requires the following infrastructure:

- Windows Server 2012 must be installed on the file server that will host the resources that Dynamic Access Control will be protecting. The file server that will host the share must be a Windows Server 2012 file server so that it can read claims and device authorization data from a Kerberos v5 ticket, translate those SIDs and claims from the ticket into an authentication token, and then compare the authorization data in the token against conditional expressions in the security descriptor.
- At least one Windows Server 2012 domain controller must be accessible by the Windows client computer in the user's domain. The new authorization and auditing mechanism requires extensions to AD DS. These extensions build the Windows claim dictionary, which is where Windows operating systems store claims for an Active Directory forest. Claims authorization also relies on the Kerberos Key Distribution Center (KDC). The Windows Server 2012 KDC contains the Kerberos enhancements that are required to transport claims within a Kerberos ticket and compound identity. Windows Server 2012 KDC also includes an enhancement to support Kerberos armoring. Kerberos armoring is an implementation of Flexible Authentication Secure Tunneling. It provides a protected channel between the Kerberos client and the KDC.
- Windows Server 2012 domain controllers must be in each domain if you are using claims across a forest trust.

Dynamic Access Control is a feature that is specific to Windows Server 2012

To deploy Dynamic Access Control, you must have these technologies:

- Windows Server 2012 domain controller
- Windows Server 2012 file server
- Windows 8 Desktop (for device claims)

- You must have a Windows 8 client if you are using device claims. Older Windows operating systems do not support device claims.

Although a Windows Server 2012 domain controller is required, there is no requirement for having a Windows Server 2012 domain and forest functional level, unless you want to use claims across a forest trust. This means that you can also have domain controllers on Windows Server 2008 and Windows Server 2008 R2 with the forest functional level located on Windows Server 2008.



Note: Implementing Dynamic Access Control in an environment with multiple forests has additional setup requirements.

Enabling Support in AD DS for Dynamic Access Control

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

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

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

To configure the Support Dynamic Access Control and Kerberos armoring policy setting, choose one of the four listed options:

- **Do not support Dynamic Access Control and Kerberos armoring**
- **Support Dynamic Access Control and Kerberos armoring**
- **Always provide claims and FAST RFC behavior**
- **Also fail unarmored authentication requests**

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

The Support Dynamic Access Control and Kerberos armoring policy setting configures Dynamic Access Control and Kerberos armoring in a mix-mode environment—when there is a mixture of Windows Server 2012 domain controllers and domain controllers running earlier versions of Windows Server.

You use the remaining policy settings when all the domain controllers are Windows Server 2012 domain controllers and the domain functional level is configured to Windows Server 2012. The Always provide claims and FAST RFC behavior policy setting and the Also fail unarmored authentication requests policy

Use Group Policy to enable support for Dynamic Access Control. Do the following:

1. Link the GPO that contains the Dynamic Access Control settings to the Domain Controllers OU
2. Navigate to the KDC node in the Group Policy Object Editor to access the Dynamic Access Control settings
3. Choose one of the following options:
 - Do not support Dynamic Access Control and Kerberos armoring
 - Support Dynamic Access Control and Kerberos armoring
 - Always provide claims and FAST RFC behavior
 - Also fail unarmored authentication requests

setting enable Dynamic Access Control and Kerberos armoring for the domain. However, the latter policy setting requires all Kerberos authentication service and ticket-granting service (TGS) communication to use Kerberos armoring. Windows Server 2012 domain controllers read this configuration while other domain controllers ignore this setting.

Implementing Claims and Resource Property Objects


After you enable support for Dynamic Access Control in AD DS, you must next create and configure claims and resource property objects.

Creating and Configuring Claim Types

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

To create attribute-based claims in Active Directory Administrative Center, navigate to the Dynamic Access Control node, and then open the Claim Types container. By default, no claim types are defined here. In the Actions pane, you can click **Create Claim Type** to view the list of attributes. These attributes are used to source values for claims. When you create a claim, you associate the claim with the specific attribute. The value of that attribute is populated as a claim value. Therefore, it is crucial that the information contained within the Active Directory attributes that are used to source claim types contain accurate information, or remain blank.

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

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

Creating and Configuring Resource Properties

Although resource properties are at the core of Dynamic Access Control, you should implement them after you have defined user and device claims. Remember that if a claim does not match the specified resource property value, then access to the data might not be allowed. Therefore, reversing the order of

Conditional expressions can include both claims and resource property objects

Claims	<ul style="list-style-type: none"> • Created for users and computers • Have attributes as a source • Created by using Active Directory Administrative Center or PowerShell
Resource property objects	<ul style="list-style-type: none"> • Created for resources • Have properties as a source • Created by using Active Directory Administrative Center or PowerShell

implementation would risk inadvertently blocking users from data that they otherwise should be able to access.

When you use claims to control access to files and folders, you must also provide additional information for those resources. You do this by configuring the resource property objects. You manage resource properties in the Resource Properties container, which is located in the Dynamic Access Control node in the Active Directory Administrative Center. You can create your own resource properties, or you can use one of preconfigured properties, such as **Project**, **Department**, and **Folder Usage**. All predefined resource property objects are disabled by default. If you want to use any of them, you should first enable them. If you want to create your own resource property object, you can specify the property type, and the allowed or suggested values.

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

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

- If you want to include only specific folders, you can use the Advanced Security Settings Editor to create conditional expressions directly in the security descriptor.
- To include several (or all) file servers, you can create central access policy rules, and then link those rules to the Central Policy objects. You can then use Group Policy to deploy the Central Policy objects to file servers, and then configure the share to use the Central Policy object. However, using central access policies is the most efficient and preferred method for securing files and folders. This is discussed further in the next topic.
- You can use file classifications to include certain files with a common set of properties across various folders or files.

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



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

Implementing Central Access Rules and Policies

Central access policies enable you manage and deploy consistent authorization throughout the organization by using central access rules and Central Access Policy objects.

Central access policies act as security nets that an organization applies across its servers. You use Group Policy to deploy the policies, and you apply the policies to all Windows Server 2012 file servers that will use Dynamic Access Control. A central access policy enables you to deploy a consistent configuration to several file servers.

The main component of a central access policy is central access rule. Central Access Policy objects represent a collection of central access rules. Before you create a central access policy, you should create a central access rule because policies are comprised of rules.

A central access rule contains multiple criteria that the Windows operating system uses when evaluating access. For example, a central access rule can use conditional expressions to target specific files and folders. Each central access rule has multiple permission entry lists that you use to manage the rule's current or proposed permission entries. You can also return the rule's current permission entry list to its last known list of permission entries. Each central access rule can be a member of one or more Central Access Policy objects.

Configuring Central Access Rules

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

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

1. Provide a name and description for the rule. You should also choose to protect the rule against accidental deletion.
2. Configure the target resources. Use the **Target Resources** section to create a scope for the access rule. You create the scope by using resource properties within one or more conditional expressions. To simplify the process, you can keep the default value (**All resources**) and apply resource filtering. You can join the conditional expressions by using logical operators, such as **AND** and **OR**. Additionally, you can group conditional expressions together to combine the results of two or more joined conditional expressions. The **Target Resources** section displays the currently configured conditional expression that is being used to control the rule's applicability.
3. Configure permissions with either of the following options:
 - **Use following permissions as proposed permissions.** Use this option to add the permissions entries in the permissions list to the list of proposed permissions entries for the newly created central access rule. You can combine the proposed permissions list with file system auditing to model the effective access that users have to the resource, without having to change the permissions entries in the current permissions list. Proposed permissions generate a special audit event to the event log that describes the proposed effective access for the users.



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

Central access policies enable you to manage and deploy consistent authorization throughout the enterprise

The main component of a central access policy is the central access rule, which specifies:

- Target resources
- Permissions
- Conditions

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

Implementing File Access Auditing

The Global Object Access Auditing feature in Windows 8 and Windows Server 2012 enables you to configure object access auditing for every file and folder in a computer's file system. You use this feature to centrally manage and configure Windows operating systems to monitor every file and folder on the computer. To enable object access auditing in previous versions of Windows Server, you had to configure this option in basic audit policies (in GPOs), and turn on auditing for a specific security principal in the object's SACL. Sometimes this approach did not easily reconcile

with company policies—such as “Log all administrative write activity on servers containing financial information—because you can turn on object access audit logging at the object level, but not at the server level. The new audit category in Windows Server 2008 R2 and Windows Server 2012 enables administrators to manage object access auditing using a much wider scope.

Dynamic Access Control enables you to create targeted audit policies using resource properties, and expressions based on user and computer claims. For example, you could create an audit policy to track all **Read** and **Write** operations on High Confidential files performed by employees who do not have a **High Security Clearance** attribute populated with the appropriate value. You can author expression-based audit policies directly on a file or folder, or centrally via Group Policy using Global Object Access Auditing. By using this approach, you do not prevent unauthorized access; instead, you register attempts to access the content by unauthorized people.

You configure Global Object Access Auditing when you enable object access auditing and global object access auditing. Enabling this feature turns on auditing for the computer that applies the policy setting. However, enabling auditing alone does not always generate auditing events. The resource—in this instance files and folders—must contain audit entries.


You should configure Global Object Access Auditing for your enterprise by using the security policy of a domain-based GPO. The two security policy settings that are required to enable global object access auditing are located in the following locations:

- Computer Configuration\Windows Settings\Security Settings\Advanced Audit Policy\Audit Policies\Object Access\Audit File System
- Computer Configuration\Windows Settings\Security Settings\Advanced Audit Policy\Audit Policy\Global Object Access Auditing\File System

Global Object Access Auditing includes a subcategory for file system and registry.

The Global Object Access Auditing policy:

- Centrally manages and configures the Windows operating system to monitor every file and folder on the server
- Integrates with Dynamic Access Control
- Provides new audit policy categories in Group Policy

 **Note:** If both a file or folder SACL and a Global Object Access Auditing policy (or a single registry setting SACL and a Global Object Access Auditing policy) are configured on a computer, the effective SACL is derived by combining the file or folder SACL and the Global Object Access Auditing policy. This means that an audit event is generated if an activity matches either the file or folder SACL or the Global Object Access Auditing policy.

Implementing Access Denied Assistance

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

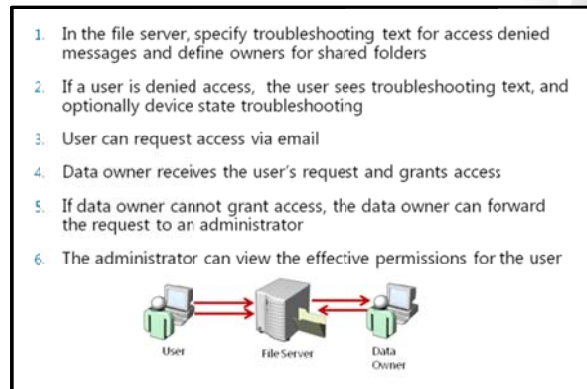
When users receive this error, they usually try to contact the administrator to obtain access. However, administrators usually do not approve access to resources, so they redirect the users to someone else for approval.

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

Access Denied Remediation

The Access Denied Assistance feature provides three ways for troubleshooting issues with access denied errors:

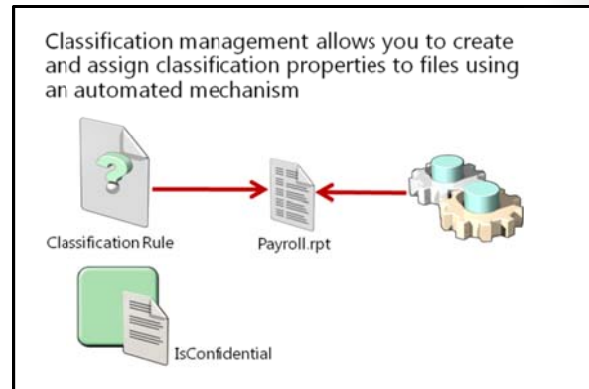
- Self-remediation. Administrators can create customized access denied messages that are authored by the server administrator. By using the information in these messages, users can try to self-remediate access denied cases. The message can also include URLs that direct users to self-remediation websites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



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

Implementing File Classifications

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



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

Classification management is designed to ease the burden and management of data that is spread out in the organization. You can classify files in a variety of ways. In most scenarios, you classify files manually. The file classification infrastructure in Windows Server 2008 R2 enables organizations to convert these manual processes into automated policies. Administrators can specify file management policies based on a file's classification, and then apply corporate requirements for managing data based on a business value.

You can use file classification to perform the following actions:

- Define classification properties and values, which you can assign to files by running classification rules.
- Create, update, and run classification rules. Each rule assigns a single predefined property and value to files within a specified directory based on installed classification plug-ins.
- When running a classification rule, reevaluate files that are already classified. You can choose to overwrite existing classification values or add the value to properties that support multiple values. You can also use this to declassify files that are no longer in the classification criteria.

Implementing Central Access Policy Changes

After you implement Dynamic Access Control, you might have to make some changes. For example, you might have to update conditional expressions, or you might want to change claims. You must carefully plan any change to Dynamic Access Control components.

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

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

Dynamic Access Control allows you to test a central access policy update by staging it

Windows Server 2012 staging:

- Is implemented by deploying proposed permissions
- Compares the proposed permissions against the current permissions
- Causes audit logs events to appear in the security log on the file server

Lab: Implementing Dynamic Access Control

Scenario

The Research team at A. Datum Corporation is involved in confidential work that provides a great deal of value to the business. Additionally, other groups at A. Datum, such as the Executive department, frequently store files containing business-critical information on the company file servers. The security department in the organization wants to ensure that these confidential files are only accessible to properly authorized personnel, and that all access to these files is audited.

As one of the senior network administrators at A. Datum, you are responsible for addressing these security requirements by implementing Dynamic Access Control on the file servers. You will work closely with the business groups and the security department to identify which files need to be secured, and who should have access to these files. You will then implement Dynamic Access Control based on the company requirements.

Objectives

- Plan the Dynamic Access Control implementation.
- Configure user and device claims.
- Configure resource property definitions.
- Configure central access rules and central access policies.
- Validate and remediate Dynamic Access Control.
- Implement new resource policies.

Lab Setup

20412A-LON-DC1

20412A-LON-SVR1

20412A-LON-CL1

20412A-LON-CL2

Estimated time: **90 minutes**

Virtual Machine(s)	20412A-LON-DC1 20412A-LON-SVR1 20412A-LON-CL1 20412A-LON-CL2
User Name	Adatum\Administrator
Password	Pa\$\$w0rd

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

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V® Manager, click **20412A-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on using the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**

5. Repeat steps 2-4 for **20412A-LON-SVR1**.
6. Do not start **20412A-LON-CL1** and **20412A-LON-CL2** until directed to do so.

Exercise 1: Planning the Dynamic Access Control Implementation

Scenario

A. Datum Corporation must ensure that documents used by the Research team and the Executive department are secured. Most of the files used by these departments are currently stored in shared folders dedicated to these departments, but confidential documents sometimes appear in other shared folders. Only members of the Research team should be able to access Research team folders, and only Executive department managers should be able to access highly confidential documents.

The security department is also concerned that managers are accessing files using their home computers, which may not be highly secure. Therefore, you must create a plan for securing the documents regardless of where they are located, and ensure that the documents can only be accessed from authorized computers. Authorized computers for managers are members of the security group ManagersWks.

The support department reports that a high number of calls are generated by users who cannot access resources. You must implement a feature that helps users understand error messages better, and that will enable them to request access automatically.

The main tasks for this exercise are as follows:

1. Plan the Dynamic Access Control deployment.
2. Prepare AD DS to support Dynamic Access Control.

► Task 1: Plan the Dynamic Access Control deployment

- Based on the scenario, describe how you will design Dynamic Access Control to fulfill the requirements for access control.

► Task 2: Prepare AD DS to support Dynamic Access Control

1. On the LON-DC1, in Server Manager, open **Active Directory Users and Computers**.
2. Create a new OU named **Test**.
3. Move **LON-CL1**, **LON-CL2**, and **LON-SVR1** computer objects into the Test OU.
4. On LON-DC1, from Server Manager, open the **Group Policy Management Console**.
5. Remove the **Block Inheritance** setting that is applied to the Managers OU. This is to remove the block inheritance setting used in a later module in the course.
6. Edit the **Default Domain Controllers Policy** GPO.
7. In the Group Policy Management Editor, under **Computer Configuration**, expand **Policies**, expand **Administrative Templates**, expand **System**, and then click **KDC**.
8. Enable the **KDC support for claims, compound authentication and Kerberos armoring** policy setting.
9. In the **Options** section, click **Supported**.
10. On LON-DC1, refresh Group Policy.
11. Open Active Directory Users and Computers, and in the **Users** container, create a security group named **ManagersWKS**.
12. Add **LON-CL1** to the **ManagersWKS** group.

13. Verify that user **Aidan Delaney** is a member of **Managers** department, and that **Allie Bellew** is the member of the **Research** department. Department entries should be filled into the appropriate attribute for each user profile.

Results: After completing this exercise, you will have planned for Dynamic Access Control deployment, and you will have prepared AD DS for Dynamic Access Control implementation.

Exercise 2: Configuring User and Device Claims

Scenario

The first step in implementing Dynamic Access Control is to configure the claims for the users and devices that access the files. In this exercise, you will review the default claims, and then create new claims based on the department and computer description attributes. For users, you will create a claim for a department attribute. For computers, you will create a claim for a description attribute.

The main tasks for this exercise are as follows:

1. Review the default claim types.
2. Configure claims for users.
3. Configure claims for devices.

► Task 1: Review the default claim types

1. On LON-DC1, in Server Manager, open the **Active Directory Administrative Center**.
2. In the Active Directory Administrative Center, click the **Dynamic Access Control** node.
3. Open the **Claim Types** container, and verify that there are no default claims defined.
4. Open the **Resource Properties** container, and note that all properties are disabled by default.
5. Open the **Resource Property Lists** container, and then open the properties of the **Global Resource Property List**.
6. In the **Resource Properties** section, review available resource properties, and then click **Cancel**.

► Task 2: Configure claims for users

1. In the Active Directory Administrative Center, in the navigation pane, click **Dynamic Access Control**.
2. Open the **Claim Types** container, and create a new claim type for users and computers using the following settings:
 - Source Attribute: **Department**
 - Display name: **Company Department**

► Task 3: Configure claims for devices

1. In the Active Directory Administrative Center, in the Tasks pane click **New**, and then select **Claim Type**.
2. Create a new claim type for computers using the following settings:
 - Source Attribute: **description**
 - Display name: **description**

Results: After completing this exercise, you will have reviewed the default claim types, configured claims for users, and configured claims for devices.

Exercise 3: Configuring Resource Property Definitions

Scenario

The second step in implementing Dynamic Access Control is to configure the resource property lists and resource property definitions. After you do this, you should make a new classification rule that classifies all files containing the word "secret". These files should be assigned a value of **High** for the **Confidentiality** attribute. You should also assign the department property to the folder that belongs to the Research team.

The main tasks for this exercise are as follows:

1. Configure resource property definitions.
2. Classify files.
3. Assign properties to a folder.

► Task 1: Configure resource property definitions

1. In the Active Directory Administrative Center, click **Dynamic Access Control**, and then open the **Resource Properties** container.
2. Enable the **Department** and **Confidentiality** Resource properties.
3. Open **Properties** for **Department**.
4. Add **Research** as suggested value.
5. Open the **Global Resource Property List**, ensure that **Department** and **Confidentiality** are included in the list, and then click **Cancel**.
6. Close the Active Directory Administrative Center.

► Task 2: Classify files

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

6. Verify values for Confidentiality. Doc1.txt and Doc2.txt should have confidentiality set to **High**.

► **Task 3: Assign properties to a folder**

1. On LON-SVR1, open Windows Explorer.
2. Browse to **C:\Research**, and open its properties.
3. On the **Classification** tab, set the **Department** value to **Research**.

Results: After completing this exercise, you will have configured resource properties for files, classified files, and assigned properties to a folder.

Exercise 4: Configuring Central Access Rules and Central Access Policies

Scenario

Now that you have configured your resource property definitions, you need to configure the central access rules and policies that will link the claims and property definitions.

The main tasks for this exercise are as follows:

1. Configure central access rules.
2. Create a central access policy.
3. Publish a central access policy by using Group Policy.
4. Apply the central access policy to resources.
5. Configure access denied remediation settings.

► **Task 1: Configure central access rules**

1. On LON-DC1, in Server Manager, click **Tools**, and then click **Active Directory Administrative Center**.
2. Click **Dynamic Access Control**, and then open the **Central Access Rules** container.
3. Create a new Central Access Rule with the following values:
 - Name: **Department Match**
 - Target Resource: **use condition Resource-Department-Equals-Value-Research**
 - Permissions: **Remove Administrators, and then add Authenticated Users, Modify, with condition User-Company Department-Equals-Resource-Department**
4. Create another Central Access Rule with the following values:
 - Name: **Access Confidential Docs**
 - Target Resource: **use condition Resource-Confidentiality-Equals-Value-High**
 - Permissions : Set first condition to: **User-Group-Member of each-Value-Managers**
 - Permissions: Set second condition to: **Device-Group-Member of each-Value-ManagersWKS**

► **Task 2: Create a central access policy**

1. On LON-DC1, in the Active Directory Administrative Center, create a new **Central Access Policy** with following values:

Name: **Protect confidential docs**

Rules included: **Access Confidential Docs**

2. Create another Central Access Policy with following values:

Name: **Department Match**

Rules included: **Department Match**

3. Close the Active Directory Administrative Center.

► **Task 3: Publish a central access policy by using Group Policy**

1. On LON-DC1, from the Server Manager, open the **Group Policy Management Console**.
2. Create new GPO named **DAC Policy**, and in the **Adatum.com** domain, link it to **Test OU**.
3. Edit the DAC Policy, browse to **Configuration/Policies/Windows Settings/Security Settings/File System**, and then right-click **Central Access Policy**.
4. Click **Manage Central Access Policies**.
5. Click both **Department Match** and **Protect confidential docs**, click **Add**, and then click **OK**.
6. Close both the Group Policy Management Editor and the Group Policy Management Console.

► **Task 4: Apply the central access policy to resources**

1. On LON-SVR1, start Windows PowerShell.
2. Refresh Group Policy on LON-SVR1.
3. Open Windows Explorer, and browse to the **C:\Docs** folder.
4. Apply the **Protect confidential docs** central policy to the **C:\Docs** folder.
5. Browse to the **C:\Research** folder.
6. Apply the **Department Match** Central Policy to the **C:\Research** folder.

► **Task 5: Configure access denied remediation settings**

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

Results: After completing this exercise, you will have configured central access rules and central access policies for Dynamic Access Control.

Exercise 5: Validating and Remediating Dynamic Access Control

Scenario

To ensure that the Dynamic Access Control settings are configured correctly, you need to test various access scenarios. You will test both approved users and devices, and unapproved users and devices. You will also validate the access remediation configuration.

The main task for this exercise is as follows:

1. Validate Dynamic Access Control functionality.

► Task 1: Validate Dynamic Access Control functionality

1. Start and then log on to **LON-CL1** as **Adatum\April** with the password **Pa\$\$w0rd**.
2. Click the **Desktop** tile, and then open Windows Explorer.
3. Browse to **\\LON-SVR1\Docs**, and verify that you can only open Doc3.
4. Try to access **\\LON-SVR1\Research**. You should be unable to access it.
5. Log off LON-CL1.
6. Log on to **LON-CL1** as **Adatum\Allie** with the password **Pa\$\$w0rd**.
7. Open Windows Explorer, and try to access **\\LON-SVR1\Research**. You should be able to access it and open files in it.
8. Log off LON-CL1.
9. Log on to **LON-CL1** as **Adatum\Aidan** with the password **Pa\$\$w0rd**.
10. Open Windows Explorer and try to access **\\LON-SVR1\Docs**. You should be able to open all files in this folder.
11. Log off LON-CL1.
12. Start and then log on to **LON-CL2** as **Adatum\Aidan** with the password **Pa\$\$w0rd**.
13. Open Windows Explorer and try to access **\\LON-SVR1\Docs**. You should be unable to see Doc1 and Doc2, because the LON-CL2 is not permitted to view secret documents.

Results: After completing this exercise, you will have validated Dynamic Access Control functionality.

Exercise 6: Implementing New Resource Policies

Scenario

As a final step in implementing Dynamic Access Control, you will test the effect of implementing a new resource policy.

The main tasks for this exercise are as follows:

1. Configure staging for a central access policy.
2. Configure staging permissions.
3. Verify staging.
4. Use effective permissions to test Dynamic Access Control.

► Task 1: Configure staging for a central access policy

1. On LON-DC1, open **Group Policy Management**.

2. Open the Group Policy Management Editor for **DAC Policy**.
3. Browse to **Computer Configuration\Policies\Windows Settings\Security Settings Advanced Audit Policy Configuration\Audit Policies**, and then select **Object Access**.
4. Double-click **Audit Central Access Policy Staging**, select all three check boxes, and then click **OK**.
5. Double-click **Audit File System**, select all three check boxes, and then click **OK**.
6. Close the Group Policy Management Editor and the Group Policy Management console.

► **Task 2: Configure staging permissions**

1. On LON-DC1, open Active Directory Administrative Center, and then open the properties for the **Department Match** central access rule.
2. In the **Proposed permissions** section, configure the condition for **Authenticated Users** as follows: **User-Company Department-Equals-Value-Marketing**.

► **Task 3: Verify staging**

1. Log on to **LON-CL1** as **Adatum\Adam** with the password **Pa\$\$w0rd**.
2. In Windows Explorer, try to access **\\LON-SVR1\Research**, and the files within it.
3. Switch to LON-SVR1.
4. Open Event Viewer, open **Security Log**, and then look for events with Event ID 4818.

► **Task 4: Use effective permissions to test Dynamic Access Control**

1. On LON-SVR1, open the properties for the **C:\Research** folder.
2. Open the **Advanced** options for **Security**, and then click **Effective Access**.
3. Click **select a user**.
4. In the Select User, Computer, Service Account, or Group window, type **April**, click **Check Names**, and then click **OK**.
5. Click **View effective access**.
6. Review the results. The user should not have access to this folder.
7. Click **Include a user claim**.
8. On the drop-down list, select **Company Department**.
9. In the **Value** text box, type **Research**.
10. Click **View Effective access**. The user should now have access.
11. Close all open windows.

Results: After completing this exercise, you will have implemented new resource policies.

► **To prepare for the next module**

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20412A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20412A-LON-SVR1**, **20412A-LON-CL1**, and **20412A-LON-CL2**.

Module Review and Takeaways

Question: What is a claim?

Question: What is the purpose of a central access policy?

Question: What is Access Denied Assistance?

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Claims are not populated with the appropriate values.	
A conditional expression does not allow access.	

Best Practices

- Use central access policies instead of configuring conditional expressions on resources.
- Enable Access Denied Assistance settings.
- Always test changes that you have made to central access rules and central access policies before implementing them.

Use file classifications to assign properties to files.

Tools

Tool	Use	Location
Active Directory Administrative Center	For administering and creating claims, resource properties, rules and policies	Administrative tools
Group Policy Management Console	Managing group policy	Administrative tools
Group Policy Management Editor	Editing Group Policy Objects	Group Policy Management Console

Module 4

Implementing Network Load Balancing

Contents:

Module Overview	4-1
Lesson 1: Overview of NLB	4-2
Lesson 2: Configuring an NLB Cluster	4-5
Lesson 3: Planning an NLB Implementation	4-10
Lab: Implementing Network Load Balancing	4-16
Module Review and Takeaways	4-21

Module Overview

Network Load Balancing (NLB) is a Windows Server network component. NLB uses a distributed algorithm to balance IP traffic load across multiple hosts. It helps to improve the scalability and availability of business-critical, IP-based services. NLB also provides high availability, because it detects host failures and automatically redistributes traffic to surviving hosts. To effectively deploy NLB, you must understand its functionality and the scenarios where its deployment is appropriate. The main change to NLB in Windows Server® 2012 is the inclusion of a comprehensive set of Windows PowerShell® cmdlets. These cmdlets enhance your ability to automate the management of Windows Server 2012 NLB clusters. The Network Load Balancing console, which is also available in Windows Server 2008 and Windows Server 2008 R2, is also present in Windows Server 2012.

This module introduces you to NLB, and shows you how to deploy this technology, the situations for which NLB is appropriate, how to configure and manage NLB clusters, and how to perform maintenance tasks on NLB clusters.

Objectives

After completing this module, you will be able to:

- Describe NLB.
- Explain how to configure an NLB cluster.
- Explain how to plan an NLB implementation.

Lesson 1

Overview of NLB

Before you deploy NLB, you need to have a firm understanding of the types of server workloads for which this high availability technology is appropriate. If you do not understand the functionality of NLB, it is possible that you will deploy it in a manner that does not accomplish your overall objectives. For example, you need to understand why NLB is appropriate for web applications, but not for Microsoft® SQL Server® databases.

This lesson provides an overview of NLB, and the features new to NLB in Windows Server 2012. It also describes how NLB works normally, and during server failure and server recovery.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe NLB technology.
- Describe how NLB works.
- Explain how NLB accommodates server failures and recovery.
- Describe new NLB features in Windows Server 2012.

What Is NLB?

NLB is a scalable, high availability feature that you can install on all editions of Windows Server 2012. A *scalable technology* is one where you can add additional components (in this case additional cluster nodes) to meet increasing demand. A *node* in a Windows Server 2012 NLB cluster is a computer, either physical or virtual, that is running the Windows Server 2012 operating system.

Windows Server 2012 NLB clusters can have between two and 32 nodes. When you create an NLB cluster, it creates a virtual network address and virtual network adapter. The virtual network adapter has an IP address and a media access control (MAC) address. Network traffic to this address is evenly distributed across the nodes in the cluster. In a basic NLB configuration, each node in an NLB cluster will service requests at a rate that is approximately equal to that of all other nodes in the cluster. When an NLB cluster receives a request, it will forward that request to the node that is currently least utilized. You can configure NLB to preference some nodes over others.

NLB is failure-aware. This means that if one of the nodes in the NLB cluster goes offline, requests will no longer be forwarded to that node, but other nodes in the cluster will continue to accept requests. When the failed node returns to service, incoming requests will be redirected until traffic is balanced across all nodes in the cluster.

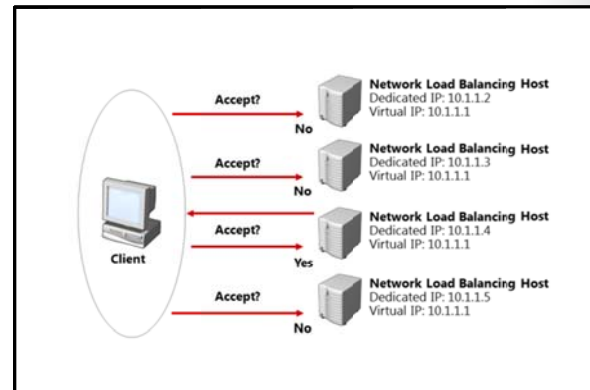
Network Load Balancing:

- Scalable high availability technology
- Balances traffic based on node utilization
 - New traffic will be directed to the node that is being utilized the least
 - You can configure NLB to preference some nodes over others.
- Used with stateless applications such as:
 - Web tiers of multi-tier applications
- Not used with stateful applications such as:
 - Traditional File servers
 - Database servers

How NLB Works

When you configure an application to use NLB, clients address the application using the NLB cluster address rather than the address of nodes that participate in the NLB cluster. The *NLB cluster address* is a virtual address that is shared between the hosts in the NLB cluster.

NLB directs traffic in the following manner: All hosts in the NLB cluster receive the incoming traffic, but only one node in the cluster, which is determined through the NLB process, will accept that traffic. All other nodes in the NLB cluster will drop the traffic.



Which node in the NLB cluster accepts the traffic depends on the configuration of port rules and affinity settings. Through these settings, you can determine if traffic that uses a particular port and protocol will be accepted by a particular node, or whether any node in the cluster will be able to accept and respond.

NLB also sends traffic to nodes based on current node utilization. New traffic is directed to nodes that are being least utilized. For example, if you have a four node cluster where three of the nodes are responding to requests from 10 clients and one node is responding to requests from 5 clients, the node that has fewer clients will receive more incoming traffic until utilization is more evenly balanced across the nodes.

How NLB Works with Server Failures and Recovery

NLB is able to detect the failure of cluster nodes. When a cluster node is in a failed state, it is removed from the cluster, and the cluster does not direct new traffic to the node. Failure is detected by using heartbeats. NLB cluster heartbeats are transmitted every second between nodes in a cluster. A node is automatically removed from a NLB cluster if it misses five consecutive heartbeats. When a node is added or removed from a cluster, a process known as convergence occurs. Convergence allows the cluster to determine its current configuration. Convergence can only occur if each node is configured with the same port rules.

NLB cluster heartbeats are transmitted every second between nodes in a cluster
 Convergence occurs when:

- A node misses five consecutive heartbeats, at which time it is automatically removed from a NLB cluster.
- A node that was member of a cluster returns to functionality.
- An administrator adds or removes a node manually.

Nodes can be configured to rejoin a cluster automatically by setting the Initial host state setting on the node's properties using the Network Load Balancing Manager. By default, a host that is a member of a cluster will attempt to rejoin that cluster automatically. For example, if you reboot a server that is a member of an NLB cluster after applying a software update, the server will rejoin the cluster automatically after the reboot process completes.

Administrators can manually add or remove nodes from NLB clusters. When an administrator removes a node, they can choose to perform a Stop or a Drainstop action. The Stop action terminates all existing connections to the cluster node and stops the NLB service. The Drainstop action blocks all new connections without terminating existing sessions. Once all current sessions end, the NLB service is stopped.

NLB can only detect server failure; it cannot detect application failure. This means that if a web application fails but the server remains operational, the NLB cluster will continue to forward traffic to the cluster node that hosts the failed application. One method of managing this problem is to implement a monitoring solution such as Microsoft System Center 2012 - Operations Manager. With Operations Manager, you can monitor functionality of applications. You can also configure Operations Manager to generate an alert in the event that an application on a cluster node fails. An alert in turn can configure a remediation action, such as restarting services, restarting the server, or withdrawing the node from the NLB cluster so that it does not receive further incoming traffic.

NLB Features in Windows Server 2012

The most substantial change to NLB features in Windows Server 2012 is the inclusion of Windows PowerShell support. The `NetworkLoadBalancingClusters` module contains 35 NLB-related cmdlets. This module becomes available on a server when the NLB Remote Server Administration Tools (RSAT) are installed. The Windows PowerShell cmdlets have the following nouns:

- **NlbClusterNode.** Lets you manage a cluster node. Includes the **Add, Get, Remove, Resume, Set, Start, Stop,** and **Suspend** verbs.
- **NlbClusterNodeDip.** Lets you configure the cluster node's dedicated management IP. Includes the **Add, Get, Remove,** and **Set** verbs.
- **NlbClusterPortRule.** Lets you manage port rules. Includes the **Add, Disable, Enable, Get, Remove,** and **Set** verbs.
- **NlbClusterVip.** Lets you manage the NLB cluster's virtual IP. Includes the **Add, Get, Remove,** and **Set** verbs.
- **NlbCluster.** Lets you manage the NLB cluster. Includes the **Get, New, Remove, Resume, Set, Start, Stop,** and **Suspend** verbs.
- **NlbClusterDriverInfo.** Provides information about the NLB cluster driver. Includes the **Get** verb.
- **NlbClusterNodeNetworkInterface.** Lets you retrieve information about a cluster node's network interface driver. Includes the **Get** verb.
- **NlbClusterIpv6Address.** Lets you configure the cluster's IPv6 address. Includes the **New** verb.
- **NlbClusterPortRuleNodeHandlingPriority.** Lets you set priority on a per-port rule basis. Supports the **Set** verb.
- **NlbClusterPortRuleNodeWeight.** Lets you set node weight on a per-port rule basis. Supports the **Set** verb.

Use 35 new NLB Windows PowerShell cmdlets to manage all aspects of NLB configuration

- Use `NlbCluster` noun to manage the cluster
- Use `NlbClusterNode` noun to manage individual nodes



Note: To see the list of Windows PowerShell cmdlets for NLB, you can use the `get-command -module NetworkLoadBalancingClusters` command.

Lesson 2

Configuring an NLB Cluster

To deploy NLB successfully, you must first have a firm understanding of its deployment requirements. You must also have planned the manner in which you are going to use port rules and affinity settings to ensure that traffic to the application that is being hosted on the NLB cluster is handled appropriately.

This lesson provides you with information about the infrastructure requirements that you must consider prior to deploying NLB. It also provides you with important information on how to configure NLB clusters and nodes to best suit your objectives.

Lesson Objectives

After completing this lesson you will be able to:

- Describe NLB deployment requirements.
- Describe how to implement NLB.
- Explain configuration options for NLB.
- Explain how to configure NLB affinity and port rules.
- Describe network considerations for NLB.

Deployment Requirements for NLB

NLB requires that all hosts in the NLB cluster reside on the same TCP/IP subnet. Although TCP/IP subnets can be configured to span multiple geographic locations, NLB clusters are unlikely to achieve convergence successfully if the latency between nodes exceeds 250 milliseconds (ms). When you are designing geographically dispersed NLB clusters, you should instead choose to deploy an NLB cluster at each site, and then use Domain Name System (DNS) round robin to distribute traffic between sites.

All network adapters within an NLB cluster must be configured as either unicast or multicast. You cannot configure an NLB cluster where there is a mixture of unicast and multicast adapters. When using unicast mode, the network adapter must support changing its MAC address.

You can only use TCP/IP protocol with network adapters that participate in NLB clusters. NLB supports IPv4 and IPv6. The IP addresses of servers that participate in an NLB cluster must be static and must not be dynamically allocated. When you install NLB, Dynamic Host Configuration Protocol (DHCP) is disabled on each interface that you configure to participate in the cluster.

All editions of Windows Server 2012 support NLB. Microsoft supports NLB clusters with nodes that are running different editions of Windows Server 2012. However, as a best practice, NLB cluster nodes should be computers with similar hardware specifications, and that are running the same edition of the Windows Server 2012 operating system.

- All hosts must be on the same subnet
- All adapters must be used with NLB configured with static IP address
- All adapters must be configured as either unicast or multicast
- Only TCP/IP protocol can be used on adapters

Demonstration: Deploying NLB

This demonstration shows how to create a Windows Server 2012 NLB cluster.

Demonstration Steps

Create a Windows Server 2012 NLB Cluster

1. Log on to **LON-SVR1** using the **Adatum\Administrator** account.
2. From the Tools menu, open the Windows PowerShell Integrated Scripting Environment (ISE).
3. Enter the following commands, and then press Enter:

```
Invoke-Command -Computername LON-SVR1,LON-SVR2 -command {Install-WindowsFeature
NLB,RSAT-NLB}
New-NlbCluster -InterfaceName "Local Area Connection" -OperationMode Multicast -
ClusterPrimaryIP 172.16.0.42 -ClusterName LON-NLB
Add-NlbClusterNode -InterfaceName "Local Area Connection" -NewNodeName "LON-SVR2" -
NewNodeInterface "Local Area Connection"
```

4. Open Network Load Balancing Manager from the Tools menu and view the cluster.

Configuration Options for NLB

Configuring NLB clusters involves specifying how hosts in the cluster will respond to incoming network traffic. How NLB directs traffic depends on the port and protocol that it is using, and whether the client has an existing network session with a host in the cluster. You can configure these settings by using port rules and affinity settings.

Port Rules

With port rules, you can configure how requests to specific IP addresses and ports are directed by the NLB cluster. You can load balance traffic on Transmission Control Protocol (TCP) port 80 across all nodes in an NLB cluster, while directing all requests to TCP port 25 to a specific host.

To specify how you want to distribute requests across nodes in the cluster, you configure a filtering mode when creating a port rule. You can do this in the **Add/Edit Port Rule** dialog box, which you can use to configure one of the following filtering modes:


- **Multiple hosts.** When you configure this mode, all NLB nodes respond according to the weight assigned to each node. Node weight is calculated automatically, based on the performance characteristics of the host. If a node fails, other nodes in the cluster continue to respond to incoming requests. Multiple host filtering increases availability and scalability, as you can increase capacity by adding nodes, and the cluster continues to function in the event of node failure.
- **Single host.** When you configure this mode, the NLB cluster directs traffic to the node that is assigned the highest priority. In the event that the node assigned the highest priority is unavailable, the host assigned the next highest priority handles the incoming traffic. Single host rules increase availability, but do not increase scalability.

Port Rules determine how traffic is routed to cluster nodes depending on TCP or UDP port

- Multiple hosts
- Single host
- Disable port range

Affinity settings determine how reconnection occurs.


- None
- Single
- Class C

 **Note:** Highest priority is the lowest number, with a priority of 1 being higher priority than a priority of 10.

- **Disable this port range.** When you configure this option, all packets for this port range are dropped, without being forwarded to any cluster nodes. If you do not disable a port range and there is no existing port rule, the traffic is forwarded to the host with the lowest priority number.

You can use the following Windows PowerShell cmdlets to manipulate port rules:

- **Add-NlbClusterPortRule.** Use this cmdlet to add a new port rule.
- **Disable-NlbClusterPortRule.** Use this cmdlet to disable an existing port rule.
- **Enable-NlbClusterPortRule.** Use this cmdlet to enable a disabled port rule.
- **Set-NlbClusterPortRule.** Use this cmdlet to modify the properties of an existing port rule.
- **Remove-NlbClusterPortRule.** Use this cmdlet to remove an existing port rule.

 **Note:** Each node in a cluster must have identical port rules. The exception to this is the load weight (in **multiple-hosts** filter mode) and handling priority (in **single-host** filter mode). Otherwise, if the port rules are not identical, the cluster will not converge.

Affinity

Affinity determines how the NLB cluster distributes requests from a specific client. Affinity settings only come into effect when you are using the **multiple hosts** filtering mode. You can select from the following affinity modes:

- **None.** In this mode, any cluster node responds to any client request, even if the client is reconnecting after an interruption. For example, the first webpage on a web application might be retrieved from the third node, the second web page from the first node, and the third web page from the second node. This affinity mode is suitable for stateless applications.
- **Single.** When you use this affinity mode, a single cluster node handles all requests from a single client. For example, if the third node in a cluster handles a client's first request, then all subsequent requests are also handled by that node. This affinity mode is useful for stateful applications.
- **Class C.** When you set this mode, a single node will respond to all requests from a class C network (one that uses the 255.255.255.0 subnet mask). This mode is useful for stateful applications where the client is accessing the NLB cluster through load-balanced proxy servers. These proxy servers will have different IP addresses, but will be within the same class C (24 bit) subnet block.

Host Parameters

You configure the host parameters for a host by clicking the host in the **Network Load Balancing Manager** console, and then from the **Host** menu, clicking **Properties**. You can configure the following host settings for each NLB node:

- **Priority.** Each NLB node is assigned a unique priority value. If no existing port rule matches the traffic that is addressed to the cluster, traffic will be assigned to the NLB node that is assigned the lowest priority value.
- **Dedicated IP address.** You can use this parameter to specify an address the host uses for remote management tasks. When you configure a dedicated IP address, NLB configures port rules so that they do not affect traffic to that address.

- Subnet Mask. When you are selecting a subnet mask, ensure that there are enough host bits to support the number of servers in the NLB cluster, and any routers that connect the NLB cluster to the rest of the organizational network. For example, if you plan to have a cluster that has 32 nodes and supports two routes to the NLB cluster, you will need to set a subnet mask that supports 34 host bits or more—such as 255.255.255.192.
- Initial host state. You can use this parameter to specify the actions the host will take after a reboot. The default Started state will have the host rejoin the NLB cluster automatically. The Suspended state pauses the host, allowing you to perform operations that require multiple reboots without triggering cluster convergence. The Stopped state stops the node.

Demonstration: Configuring NLB Affinity and Port Rules

This demonstration shows how to configure affinity for NLB cluster nodes, and how to configure NLB port rules.

Demonstration Steps

Configure Affinity for NLB Cluster Nodes

1. On LON-SVR2, on the taskbar, click the **Windows PowerShell** icon.
2. In Windows PowerShell, enter each of the following commands, pressing Enter after each command:

```
Cmd.exe
Mkdir c:\porttest
Xcopy /s c:\inetpub\wwwroot c:\porttest
Exit
New-Website -Name PortTest -PhysicalPath "C:\porttest" -Port 5678
New-NetFirewallRule -DisplayName PortTest -Protocol TCP -LocalPort 5678
```

Configure NLB Port Rules

1. On LON-SVR1, open the Network Load Balancing Manager.
2. Remove the **All port** rule.
3. In Network Load Balancing Manager, edit the properties of the LON-NLB cluster.
4. Add a port rule with the following properties:
 - Port range: **80 to 80**
 - Protocols: **Both**
 - Filtering mode: **Multiple Host**
 - Affinity: **None**
5. Create a port rule with the following properties:
 - Port range: **5678 to 5678**
 - Protocols: **Both**
 - Filtering mode: **Single Host**
6. Edit the host properties of LON-SVR1.
7. Configure the port rule for port 5678 and set handling priority to **10**.

Network Considerations for NLB

You must consider several factors when you are designing a network to support an NLB cluster. The primary decision is whether you want to configure the NLB cluster to use Unicast or Multicast cluster operation mode.

Unicast Mode

When you configure a NLB cluster to use unicast mode, all cluster hosts use the same unicast MAC address. Outgoing traffic uses a modified MAC address that is determined by the cluster host's priority setting. This prevents the switch that handles outbound traffic from having problems with all cluster hosts using the same MAC address.

When you use unicast mode with a single network adapter on each node, only computers that use the same subnet can communicate with the node using the node's assigned IP address. If you have to perform any node management tasks, such as connecting using Remote Desktop to apply software updates, you will need to perform these tasks from a computer that is on the same TCP/IP subnet as the node.

When you use unicast mode with two or more network adapters, one adapter will be used for dedicated cluster communication, and the other adapter or adapters can be used for management tasks. When you use unicast mode with multiple network adapters, you can perform cluster management tasks such as connecting using Remote PowerShell to add or remove roles and features.

Unicast mode can also minimize problems that occur when cluster nodes also host other non-NLB related roles or services. For example, using unicast mode means that a server that participates in a web server cluster on port 80 may also host another service such as DNS or DHCP. Although this is possible, Microsoft recommends that all cluster nodes have the same configuration.

Multicast Mode

When you configure an NLB cluster to use multicast mode, each cluster host keeps its original MAC address, but also is assigned an additional multicast MAC address. Each node in the cluster is assigned the same additional MAC multicast address. Multicast mode requires network switches and routers that support multicast MAC addresses.

Internet Group Management Protocol Multicast

Internet Group Management Protocol (IGMP) multicast mode is a special form of multicast mode that prevents the network switch from being flooded with traffic. When you deploy IGMP multicast mode, traffic is forwarded only through switch ports that participate in the NLB cluster. IGMP multicast mode requires switch hardware that supports this functionality.

Network Considerations

You can improve NLB cluster performance when using unicast mode by using separate virtual local area networks (VLANs) for cluster traffic and management traffic. Using VLANs segments traffic, thereby preventing management traffic from affecting cluster traffic. When you host NLB nodes on virtual machines using Windows Server 2012, you can also use network virtualization to segment management traffic from cluster traffic.

Unicast Mode

- Suitable for clusters that have multiple network adapters

Multicast Mode

- Suitable for NLB clusters that have single network adapters
- Network devices must support multicast MAC addresses

IGMP Multicast

- Improves switch performance
- Requires a network switch that supports this functionality

Lesson 3

Planning an NLB Implementation

When you are planning an NLB implementation, you must ensure that the applications that you deploy are appropriate for NLB. Not all applications are suitable for deployment on NLB clusters, and it is important for you to be able to identify which ones can benefit from this technology. You also need to know what steps you can take to secure NLB, and be familiar with the options that you have to scale NLB, should the application hosted on the NLB cluster require greater capacity.

Lesson Objectives

After completing this lesson you will be able to:

- Explain how to design application and storage support for NLB.
- Describe the special considerations for deploying NLB clusters on virtual machines.
- Describe the options that you can implement to secure NLB.
- Describe the options for scaling NLB.

Describe the method you can use to upgrade an NLB cluster to Windows Server 2012.


Designing Applications and Storage Support for NLB

Because clients can be redirected to any node in an NLB cluster, each node in the cluster must be able to provide a consistent experience. Therefore, when you are designing applications and storage support for NLB applications, you must ensure that you configure each node in the same way, and that each node has access to the same data.

When a highly available application has multiple tiers—such as a web application that includes an SQL Server database tier—the web application tier is hosted on an NLB cluster. SQL Server, as a stateful application, is not made highly available

using NLB. Instead, you use technologies such as failover clustering, mirroring, or AlwaysOn Availability Groups, to make the SQL Server database tier highly available.

All hosts in an NLB cluster should run the same applications and be configured in the same way. When you are using web applications, you can use Internet Information Services (IIS) 8.0's shared configuration functionality to ensure that all nodes in the NLB cluster are configured in the same manner.

 **Reference Links:** You can find out more about IIS 8.0's shared configuration with Windows Server 2012 at <http://learn.iis.net/page.aspx/264/shared-configuration/>

You can also use technologies such as file shares that are hosted on Cluster Shared Volumes (CSV) to host application configuration information. File shares hosted on CSVs allow multiple hosts to have access to application data and configuration information. File shares that are hosted on CSVs are a feature of Windows Server 2012.

- Each node in an NLB cluster needs to have the same configuration
- Each node needs access to the same consistent application data.
- Use IIS Shared Configuration to ensure that web application configuration is consistent across NLB nodes.
- Use Cluster Shared Volumes to host shared application and configuration data for NLB applications.

Considerations for Deploying an NLB Cluster on Virtual Machines

As organizations transition from physical to virtual deployments, administrators must consider several factors when determining the placement of NLB cluster nodes on Hyper-V hosts. This includes the network configuration of virtual machines, the configuration of the Hyper-V hosts, and the benefits of using Hyper-V's high availability features in conjunction with NLB.

Virtual Machine Placement

You should place NLB cluster nodes on separate hard disks on the Hyper-V® host. That way, should a disk or disk array fail, even if one node becomes unavailable, other NLB cluster nodes that are hosted on the same Hyper-V host will remain online. As a best practice, you should configure the Hyper-V host with redundant hardware, including redundant disks, network adapters, and power supplies. This will minimize the chance that hardware failure on the Hyper-V host will lead to all nodes in an NLB cluster becoming unavailable. When you are using multiple network adapters, configure network teaming to ensure that virtual machines are able to maintain access to the network even in the event that individual network adapter hardware suffers a failure.

Where possible, deploy NLB virtual machine nodes on separate hyper-V hosts. When you are planning this type of configuration, ensure that the virtual machines that participate in the NLB cluster are located on the same TCP/IP subnet. This protects the NLB cluster from other types of server failure, such as the failure of a motherboard or any other single point of failure.

Virtual Machine Network Configuration

Because adding additional virtual network adapters is a straightforward process, you can configure the NLB cluster to use unicast mode, and then deploy each virtual machine with multiple network adapters. You should create separate virtual switches for cluster traffic and node management traffic, because segmenting traffic can improve performance. You can also use network virtualization to partition cluster traffic from node management traffic. You can use VLAN tags as a method of partitioning cluster traffic from node management traffic.

When you are using unicast mode, ensure that you enable MAC address spoofing for the virtual network adapter on the Hyper-V host. You can do this by editing the virtual network adapter's settings on the **Virtual Machine Settings** dialog box, which is available through the Hyper-V Manager. Enabling MAC address spoofing allows unicast mode to configure MAC address assignment on the virtual network adapter.

NLB Cluster vs. Virtual Machine High Availability

Virtual machine high availability is the process of placing virtual machines on failover clusters. When a failover cluster node fails, the virtual machine fails over, so that it is hosted on another node. Though failover clustering and NLB are both high availability technologies, they serve different purposes. Failover clustering supports stateful applications such as SQL Server, whereas NLB is suited to stateless applications such as websites. Highly available virtual machines do not allow an application to scale, because you cannot add nodes to increase capacity. However, it is possible to deploy NLB cluster nodes as highly available virtual machines. In this scenario, the NLB cluster nodes fail over to a new Hyper-V host in the event that the original Hyper-V host fails.

- Configure virtual machines with multiple network adapters
- Configure one network adapter on each node member to use a shared private network switch
- Configure the NLB cluster to use unicast mode
- Use the shared private network switch for cluster communication
- When NLB nodes span multiple sites, use network virtualization to separate cluster network.

The degree of availability and redundancy required will fluctuate, depending on the application. A business-critical application that costs an organization millions of dollars when it is down requires an availability that differs from that of an application that causes minimal inconvenience if it is offline.

Considerations for Securing NLB

NLB clusters are almost always used to host web applications that are important to the organization. Because of this importance, you should take steps to secure NLB, both by restricting the traffic that can address the cluster, and by ensuring that appropriate permissions are applied.

Configure Port Rules

When securing NLB clusters, you must first ensure that you create port rules to block traffic to all ports other than those used by applications hosted on the NLB cluster. When you do this, all incoming traffic that is not specifically addressed to applications that are running on the NLB cluster will be dropped, before being forwarded to cluster nodes. If you do not do this first step, all incoming traffic that is not managed by a port rule will be forwarded to the cluster node with the lowest cluster priority value.

- Use NLB cluster port rules to discard traffic not related to cluster applications.
- Use firewall rules on cluster nodes to drop traffic not related to cluster applications or node management.
- Configure applications to respond only to traffic that is addressed to the cluster.
- Use SANs to create certificates that support the application name and node names.
- Implement principle of least privilege to ensure that only authorized users have appropriate permissions on nodes.

Configure Firewall Rules

You should also ensure that Windows Firewall with Advanced Security is configured on each NLB cluster node. When you enable NLB on a cluster node, the following firewall rules are created and enabled automatically. This allows NLB to function and communicate with other nodes in the cluster:

- Network Load Balancing (DCOM-In)
- Network Load Balancing (ICMP4-ERQ-In)
- Network Load Balancing (ICMP6-ERQ-In)
- Network Load Balancing (RPCSS)
- Network Load Balancing (WinMgmt-In)
- Network Load Balancing (ICMP4-DU-In)
- Network Load Balancing (ICMP4-ER-In)
- Network Load Balancing (ICMP6-DU-In)
- Network Load Balancing (ICMP6-EU-In)

When created, these firewall rules do not include scope settings. In high-security environments, you would configure an appropriate local IP address or IP address range, and a remote IP address for each of these rules. The remote IP address or address range should include the addresses that are used by other hosts in the cluster.

When you are configuring additional firewall rules, remember the following:

- When you are using multiple network adapters in unicast mode, configure different firewall rules for each network interface. For the interface used for management tasks, you should configure the firewall rules to allow inbound management traffic only. For example, enabling the use of remote Windows PowerShell, Windows Remote Manager (Windows RM), and Remote Desktop for

management tasks. You should configure the firewall rules on the network interface the cluster node uses, to provide an application to the cluster, and to allow access to that application. For example, allow incoming traffic on TCP ports 80 and 443 on an application that uses the HTTP and HTTPS protocols.

- When you are using multiple network adapters in multicast mode, configure firewall rules that allow access to applications that are hosted on the cluster, but block access to other ports.

Configure Applications to Respond Only to Traffic Addressed to the Cluster

You should configure applications on each node to respond only to traffic that is addressed to the cluster, and to ignore application traffic that is addressed to the individual node. For example, if you deploy a web application that is designed to respond to traffic addressed to `www.adatum.com`, there will be a website on each node that will accept traffic on port 80. Depending on the NLB cluster configuration, it is possible that traffic that is addressed to the node on port 80 will generate a direct response. For example, getting the Adatum web application by entering the address `http://nlb-node-3.adatum.com` in a browser instead of entering the address `http://www.adatum.com`. You can secure applications from this type of direct traffic by configuring them to respond only to traffic that uses the NLB cluster address. For web applications, you can do this by configuring the website to use a host header. Each application that runs on an NLB cluster will have its own unique method of allowing you to configure the application to respond only to traffic directed at the cluster, rather than at the individual cluster node.

Securing Traffic with SSL

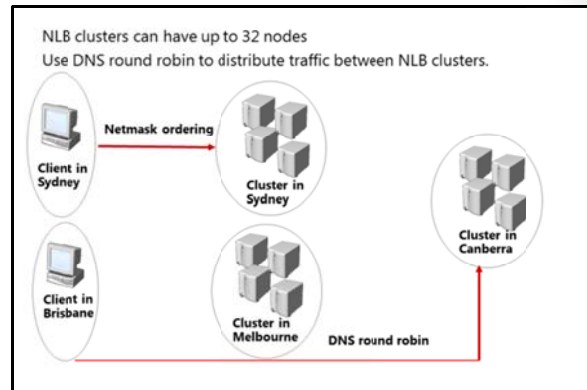
NLB websites must all use the same website name. When you are securing websites that you make highly available using NLB, you need to ensure that each website has an SSL certificate that matches the website name. You can use host headers on each node. In most cases, you will install the same website certificate on each node in the NLB cluster, because this is simpler than procuring separate certificates for each cluster node. In some cases, you will need to procure certificates that support subject alternative names (SANs). Certificates that support SANs allow a server to be identified by multiple names, such as the name used by the clustered application and the name of the cluster node. For example, a certificate with a SAN might support the names `www.adatum.com`, `node1.adatum.com`, `node2.adatum.com`, `node3.adatum.com`, and `node4.adatum.internal`.

Principle of Least Privilege

Ensure that users are only delegated permissions for tasks that they need to perform on the NLB node. Members of the local Administrators group on any single node are able to add and remove cluster nodes, even if they are not members of the local Administrators group on those nodes. Applications that run on NLB clusters should be configured in such a way that they do not require application administrators to have local Administrator privileges on the servers that host the application. Only users whose job role requires them to be able to make remote management connections to NLB cluster nodes should be able to make those connections.

Considerations for Scaling NLB

Scaling is the process of increasing the capacity of an NLB cluster. For example, if you have a four-node NLB cluster and each node in the cluster is being heavily utilized to the point where the cluster cannot manage more traffic, you can add additional nodes. Adding nodes will spread the same load across more computers, reducing the load on each current cluster node. Capacity increases because a larger number of similarly configured computers can manage a higher workload than a smaller number of similarly configured computers.



An NLB cluster supports up to 32 nodes. This means that you can scale-out a single NLB cluster so that 32 separate nodes participate in that cluster. When you consider scaling an application so that it is hosted on a 32-node NLB cluster, remember that each node in the cluster must be on the same TCP/IP subnet.

An alternative to building single NLB clusters is to build multiple NLB clusters, and then to use DNS round robin to share traffic between them. DNS round robin is a technology that allows a DNS server to provide requesting clients with different IP addresses to the same hostname in sequential order. For example, if there are three addresses associated with a hostname, the first requesting host receives the first address, the second receives the second address, the third receives the third, and so forth. When you use DNS round robin with NLB, you associate the IP addresses of each cluster with the hostname that is used by the application.

Distributing traffic between NLB clusters using DNS round robin also allows you to deploy NLB clusters across multiple sites. DNS round robin supports netmask ordering. This technology ensures that clients on a subnet are provided with an IP address of a host on the same network, if one is available. For example, you might deploy three four-node NLB clusters in the cities of Sydney, Melbourne, and Canberra, and use DNS round robin to distribute traffic between them. With netmask ordering, a client that is accessing the application in Sydney will be directed to the NLB cluster hosted in Sydney. A client that is not on the same subnet as the NLB cluster nodes, such as a client in the city of Brisbane, would be directed by DNS round robin to the Sydney, Melbourne, or Canberra NLB cluster.

Considerations for Upgrading NLB Clusters

Upgrading NLB clusters involves moving cluster nodes from one host operating system—for example Windows Server 2003 or Windows Server 2008—to Windows Server 2012. Upgrading the cluster might not involve performing an operating system upgrade on each node, because in some cases the original host operating system might not support a direct upgrade to Windows Server 2012. In cases where the original host operating system does not support a direct upgrade to Windows Server 2012, you can perform a migration.

A key consideration when you are upgrading NLB clusters is to remember that NLB supports having clusters that are running a mixture of operating systems. This means that you can have a cluster running a mixture of Windows Server 2003, Windows Server 2008,

Possible to run NLB clusters that have different operating systems

- Windows Server 2012 NLB clusters can interoperate with:
 - Windows Server 2003 & Windows Server 2003 R2
 - Windows Server 2008 & Windows Server 2008 R2

Piecemeal upgrade:

- Add Windows Server 2012 cluster nodes
- Remove nodes running earlier operating systems

Upgrade clusters:

1. Remove node from NLB cluster
2. Upgrade to Windows Server 2012
3. Rejoin node to NLB cluster

and Windows Server 2012. Keep in mind that while mixed operating system NLB clusters are supported, they are not recommended. You should configure the NLB cluster so that all hosts are running the same operating system as soon as possible.



Note: In some situations, it will not be possible to upgrade the operating system of a cluster node.

When you are performing an upgrade, you can use one of the following strategies:

- **Piecemeal Upgrade.** During this type of upgrade, you add new Windows Server 2012 nodes to an existing cluster, and then remove the nodes that are running earlier versions of the Windows Server operating system. This type of upgrade is appropriate when the original hardware and operating system does not support a direct upgrade to Windows Server 2012.
- **Rolling upgrade.** During this type of upgrade, you upgrade one node in the cluster at a time. You do this by taking the node offline, performing the upgrade, and then rejoining the node back to the cluster.



Reference Links: You can learn more about upgrading NLB clusters at the following link:
[http://technet.microsoft.com/en-us/library/cc731691\(W.S.10\).aspx](http://technet.microsoft.com/en-us/library/cc731691(W.S.10).aspx)

Lab: Implementing Network Load Balancing

Scenario

A. Datum Corporation is an engineering and manufacturing company. The organization is based in London, England, and is quickly expanding into Australia. As the company expands, the need for scalable web applications has increased. With this in mind, you are developing a pilot program to test the deployment of Windows NLB on hosts running the Windows Server 2012 operating system.

As you intend to automate the process of deploying Windows NLB clusters, you will use Windows PowerShell to perform many of the cluster setup and configuration tasks. You will also configure port rules and affinity, which will allow you to deploy multiple load-balanced web applications on the same Windows NLB clusters.

Objectives

- Create a Windows NLB cluster.
- Configure and manage an NLB cluster.
- Validate high availability for the NLB cluster.

Lab Setup

Estimated Time: 45 minutes

20412A-LON-DC1

20412A-LON-SVR1

20412A-LON-SVR2

Username: **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 **20412A-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on using the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
5. Repeat steps 2-4 for **20412A-LON-SVR1** and **20412A-LON-SVR2**.

Exercise 1: Implementing an NLB Cluster

Scenario

You eventually want to automate the process of deploying Windows Server 2012 NLB clusters. With this in mind, you will be using Windows PowerShell to perform the majority of the NLB cluster deployment tasks.

The main tasks for this exercise are as follows:

1. Verify website functionality for standalone servers.
2. Install the Windows Network Load Balancing feature.
3. Create a new Windows Server 2012 NLB cluster.

4. Add a second host to the cluster.
5. Validate the NLB cluster.

► **Task 1: Verify website functionality for standalone servers**

1. On LON-SVR1, navigate to the folder **c:\inetpub\wwwroot**.
2. Open **iis-8.png** in Microsoft® Paint, and use the Paint Brush tool and the color red to mark the IIS Logo in a distinctive manner.
3. Close Windows Explorer.
4. Switch to LON-DC1 and then click to the Start screen.
5. Open Internet Explorer.
6. Navigate to **http://LON-SVR1**. Verify that the web page is marked in a distinctive manner with the color red.
7. Navigate to **http://LON-SVR2**. Verify that the website is not marked in a distinctive manner.
8. Close Internet Explorer.

► **Task 2: Install the Windows Network Load Balancing feature**

1. On LON-SVR1, open Windows PowerShell ISE.
2. Type the following command, and then press Enter:

```
Invoke-Command -Computername LON-SVR1,LON-SVR2 -command {Install-WindowsFeature
NLB,RSAT-NLB}
```

► **Task 3: Create a new Windows Server 2012 NLB cluster**

1. On LON-SVR1, in Windows PowerShell ISE, type the following command, and then press Enter:

```
New-NlbCluster -InterfaceName "Local Area Connection" -OperationMode Multicast -
ClusterPrimaryIP 172.16.0.42 -ClusterName LON-NLB
```

2. In Windows PowerShell ISE, type the following command, and then press Enter:

```
Invoke-Command -Computername LON-DC1 -command {Add-DNSServerResourceRecordA
zonename adatum.com -name LON-NLB -Ipv4Address 172.16.0.42}
```

► **Task 4: Add a second host to the cluster**

- On LON-SVR1, in the Windows PowerShell ISE window, type the following command, and then press Enter:

```
Add-NlbClusterNode -InterfaceName "Local Area Connection" -NewNodeName "LON-SVR2" -
NewNodeInterface "Local Area Connection"
```

► **Task 5: Validate the NLB cluster**

1. On LON-SVR1, open the Network Load Balancing Manager, and verify that nodes LON-SVR1 and LON-SVR2 display with the status of **Converged**.
2. View the properties of the LON-NLB cluster, and verify the following:
 - The cluster is set to use the **Multicast** operations mode.
 - There is a single port rule named **All** that starts at port **0** and ends at port **65535** for both **TCP** and **UDP** protocols, and that it uses **Single** affinity.

Results: After this exercise, you should have successfully implemented an NLB cluster.

Exercise 2: Configuring and Managing the NLB Cluster

Scenario

As you will want to deploy multiple separate websites to the NLB cluster and differentiate these websites based on port address, you want to ensure that you are able to configure and validate port rules. You also want to experiment with affinity settings to ensure that requests are distributed evenly across hosts.

The main tasks for this exercise are as follows:

1. Configure port rules and affinity.
2. Validate port rules.
3. Manage host availability in the NLB Cluster.

► Task 1: Configure port rules and affinity

1. On LON-SVR2, open Windows PowerShell.
2. In Windows PowerShell, enter the following commands, pressing Enter after each command:

```
Cmd.exe
Mkdir c:\porttest
Xcopy /s c:\inetpub\wwwroot c:\porttest
Exit
New-Website -Name PortTest -PhysicalPath "C:\porttest" -Port 5678
New-NetFirewallRule -DisplayName PortTest -Protocol TCP -LocalPort 5678
```

3. Open Windows Explorer and then browse to and open **c:\porttest\iis-8.png** in Microsoft Paint.
4. Use the **Blue** paintbrush to mark the IIS Logo in a distinctive manner.
5. Switch to LON-DC1.
6. Open Internet Explorer and navigate to **http://LON-SVR2:5678**.
7. Verify that the IIS Start page with the image marked with blue displays.
8. Switch to LON-SVR1.
9. On LON-SVR1, open Network Load Balancing Manager, and view the cluster properties of LON-NLB.
10. Remove the **All port** rule.
11. Add a port rule with the following properties:
 - Port range: **80 to 80**
 - Protocols: **Both**
 - Filtering mode: **Multiple Host**
 - Affinity: **None**
12. Create a new port rule with the following properties:
 - Port range: **5678 to 5678**
 - Protocols: **Both**
 - Filtering mode: **Single Host**
13. Click **OK** to close the Cluster Properties dialog box.

14. Edit the host properties of LON-SVR1.
15. Configure the **Handling Priority** value of the port rule for port **5678** as **10**.

► **Task 2: Validate port rules**

1. Switch to LON-DC1.
2. Using Internet Explorer, navigate to **http://lon-nlb**, refresh the Web page 20 times, and verify that web pages with and without the distinctive red marking display.
3. On LON-DC1, navigate to address **http://LON-NLB:5678**, refresh the web page 20 times, and verify that only the web page with the distinctive blue marking displays.

► **Task 3: Manage host availability in the NLB Cluster**

1. Switch to LON-SVR1.
2. Use the Network Load Balancing Manager on LON-SVR1 to suspend LON-SVR1.
3. Verify that node LON-SVR1 displays as **Suspended**, and that node LON-SVR2 displays as **Converged**.
4. Resume and then start LON-SVR1.
5. Verify that both node LON-SVR1 and LON-SVR2 now display as **Converged**.

Results: After this exercise, you should have successfully configured and managed an NLB cluster.

Exercise 3: Validating High Availability for the NLB Cluster

Scenario

As part of preparing for the deployment of NLB in your organization's environment, you want to ensure that it is possible to perform maintenance tasks such as reboot operations, without affecting the availability of the websites that are hosted on the cluster. With this in mind, you will verify availability by rebooting one of the hosts while attempting to access the clustered website. You will also explore the Drainstop functionality.

The main tasks for this exercise are as follows:

1. Validate website availability when the host is unavailable.
2. Configure and validate Drainstop.

► **Task 1: Validate website availability when the host is unavailable**

1. Restart LON-SVR1.
2. Switch to LON-DC1.
3. On LON-DC1, open Internet Explorer, and navigate to **http://LON-NLB**.
4. Refresh the website 20 times. Verify that the website is available, but that it does not display the distinctive red mark on the IIS logo until LON-SVR1 has restarted.

► **Task 2: Configure and validate Drainstop**

1. On LON-SVR1, open Network Load Balancing Manager and initiate a Drainstop on LON-SVR2.
2. On LON-DC1, navigate to **http://lon-nlb** and verify that only the welcome page with the red IIS logo displays.

Results: After this exercise, you should have successfully validated high availability for the NLB cluster.

► **To prepare for the next module**

When you finish the lab, revert the virtual machines to their initial state.

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

Lab Review

Question: How many additional nodes can you add to the LON-NLB cluster?

Question: What steps would you take to ensure that LON-SVR1 always manages requests for web traffic on port 5678, given the port rules that exist at the end of this exercise?

Question: What is the difference between a Stop and a Drainstop action?

Module Review and Takeaways

Question: You have created a four-node Windows Server 2012 NLB cluster. The cluster hosts a website that is hosted on IIS. What happens to the cluster if you shut down the World Wide Web publishing service on one of the nodes?

Question: You want to host the `www.contoso.com`, `www.adatum.com`, and `www.fabrikam.com` websites on a four-node NLB cluster. The cluster IP address will be a public IP address, and each fully qualified domain name (FQDN) is mapped in DNS to the cluster's public IP address. What steps should you take on each node to ensure that traffic is directed to the appropriate site?

Question: You have an eight-node Windows NLB cluster that hosts a web application. You want to ensure that traffic from a client that uses the cluster remains with the same node throughout their session, but that traffic from separate clients is distributed equitably across all nodes. Which option do you configure to accomplish this goal?

Real-world Issues and Scenarios

To become a true high availability solution, use a monitoring solution with NLB that will detect application failure. This is because NLB clusters will continue to direct traffic to nodes with failed applications as long as NLB—which is independent of the application—continues to send heartbeat traffic.

MCT USE ONLY. STUDENT USE PROHIBITED

Module 5

Implementing Failover Clustering

Contents:

Module Overview	5-1
Lesson 1: Overview of Failover Clustering	5-2
Lesson 2: Implementing a Failover Cluster	5-14
Lesson 3: Configuring Highly Available Applications and Services on a Failover Cluster	5-20
Lesson 4: Maintaining a Failover Cluster	5-25
Lesson 5: Implementing a Multi-Site Failover Cluster	5-30
Lab: Implementing Failover Clustering	5-36
Module Review and Takeaways	5-41

Module Overview

Providing high availability is important for any organization that wants to provide continuous services to its users. High availability is a term that denotes the capability of a system or device to be usable when it is required. You can express high availability as a percentage, which is calculated by dividing the actual service time by the required service time. High availability does not mean that the system will be free of any downtime. However, a network that has an uptime of 99.999 percent often is considered highly available. Failover clustering is one of the main technologies in Windows Server® 2012 that can provide high availability for various applications and services. In this module, you will learn about failover clustering, its components, and implementation techniques.

Objectives

After completing this module, you will be able to:

- Describe failover clustering.
- Implement a failover cluster.
- Configure highly available applications and services.
- Maintain a failover cluster.
- Implement multi-site failover clustering.

Lesson 1

Overview of Failover Clustering

Failover clustering is a high availability process, wherein an instance of a service or application that is running over one machine can fail-over onto a different machine in the failover cluster if the first machine fails. Failover clusters in Windows Server 2012 provide a high availability solution for many server roles and applications. By implementing failover clusters, you can maintain application or service availability if one or more computers in the failover cluster fail.

Before you implement failover clustering, you should be familiar with general high availability concepts. You must also understand clustering terminology, and how failover clusters work. Finally, you must be familiar with new clustering features in Windows Server 2012.

Lesson Objectives

After completing this lesson, you will be able to:

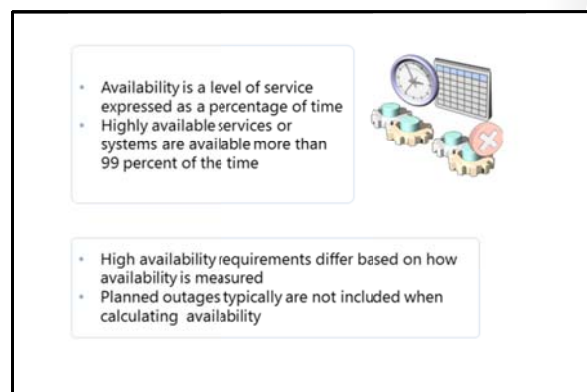
- Describe high availability.
- Describe failover clustering improvements in Windows Server 2012.
- Describe failover cluster components.
- Describe Cluster Shared Volumes (CSV).
- Define failover and failback.
- Describe a quorum.
- Describe quorum modes in Windows Server 2012 failover clusters.
- Describe failover cluster networks.
- Describe failover cluster storage.

What Is High Availability?

Availability refers to a level of service that applications, services, or systems provide. Availability is expressed as the percentage of time that a service or system is available. Highly available systems have minimal downtime—whether planned or unplanned—and are available more than 99 percent of the time, depending on an organization's needs and budget. For example, a system that is unavailable for 8.75 hours per year would have a 99.9 percent availability rating, and would be considered highly available.

To improve availability, you must implement fault-tolerance mechanisms that mask or minimize how failures of the service's components and dependencies affect the system. You can achieve fault tolerance by implementing redundancy to single points of failure.

Miscommunication about service-level expectations between the customer and the IT organization can result in poor business decisions, such as unsuitable investment levels and customer dissatisfaction. Be sure to express availability requirements clearly, so that there are no misunderstandings about the implications.



The availability measurement period can also have a significant effect on the definition of availability. For example, a requirement for 99.9 percent availability over a one-year period allows for 8.75 hours of downtime, whereas a requirement for 99.9 percent availability over a rolling four-week window allows for only 40 minutes of downtime per period.

For high availability, you also should identify and negotiate planned outages, service and support hours, maintenance activities, service pack updates, and software updates. These are scheduled outages, and typically not included as downtime; you typically calculate availability based on unplanned outages only. However, you have to negotiate exactly which planned outages you will consider as downtime.

Failover Clustering in Windows Server 2012

While most of the failover clustering features and administration techniques from Windows Server 2008 R2 are present in Windows Server 2012, some new features and technologies in Windows Server 2012 increase scalability and cluster storage availability, and provide better and easier management and faster failover.

The important new features in Windows Server 2012 failover clustering include:

- **Increased scalability.** In Windows Server 2012, a failover cluster can have 64 physical nodes and can run 4,000 virtual machines on each cluster. This is a significant improvement over Windows Server 2008 R2, which supports only 16 physical nodes and 1,000 virtual machines per cluster. Each cluster you create is now available from the Server Manager console. Server Manager in Windows Server 2012 can discover and manage all clusters that are created in an Active Directory® Domain Services (AD DS) domain. If you deploy the cluster in a multi-site scenario, the administrator can now control which nodes in a cluster have votes for establishing quorum. Failover clustering scalability is also improved for virtual machines that are running on clusters. This will be discussed in more detail in "Module 6: Implementing Hyper-V Availability."
- **Improved CSVs.** This technology was introduced in Windows Server 2008 R2, and it became very popular for providing virtual machine storage. In Windows Server 2012, CSVs appear as CSV file systems, and they support server message block (SMB) version 3.0 storage for Hyper-V® and other applications. In addition, CSV can use the Server Message Block (SMB) Multichannel and SMB Direct features to enable traffic to stream across multiple networks in a cluster. It is also possible to implement file server on CSVs, in scale-out mode. For additional security, you can use Windows BitLocker® drive encryption for CSV disks, and you can make CSV storage visible only to a subset of nodes in a cluster. For reliability, you can scan and repair CSV volumes with zero offline time.
- **Cluster-Aware Updating.** In earlier versions of Windows Server, updating cluster nodes to minimize or avoid downtime required significant preparation and planning. In addition, updating cluster nodes was a mostly manual procedure, which caused additional administrative effort. Windows Server 2012 introduces Cluster-Aware Updating, a new technology for this purpose. Cluster-Aware Updating automatically updates cluster nodes with the Windows Update hotfix, while keeping the cluster online and minimizing downtime. This technology will be explained in more detail in "Lesson 4: Maintaining a Failover Cluster."

Windows Server 2012 failover clustering improvements include:

- Increased scalability
- Improved CSVs
- Cluster-aware updating
- Active Directory integration improvements
- Management improvements

The **Cluster.exe** command-line tool, Cluster Automation Server (MSCluster) COM interface and **Add-ClusterPrintServerRole** cmdlet are removed in Windows Server 2012.

- Active Directory integration improvements. In Windows Server 2008, failover clustering is integrated in AD DS. Windows Server 2012 improves on this integration. Administrators can now create cluster computer objects in targeted organizational units (OUs), or by default in the same OUs as the cluster nodes. This aligns failover cluster dependencies on AD DS with the delegated domain administration model that is used in many IT organizations. In addition, you can now deploy failover clusters with access only to read-only domain controllers.
- Management improvements. Although failover clustering in Windows Server 2012 still uses almost the same management console and the same administrative techniques, Windows Server 2012 brings some important management improvements. In the Validation wizard, the validation speed for large failover clusters is improved and new tests for CSVs, the Hyper-V role, and virtual machines are added. In addition, new Windows PowerShell® cmdlets are available for managing clusters, monitoring clustered virtual machine applications, and creating highly available Internet Small Computer System Interface (iSCSI) targets.

Removed and Deprecated Features

In Windows Server 2012 clustering, some of the features from older failover clustering versions are removed or deprecated. If you are upgrading from an older version, you should be aware of these changes:

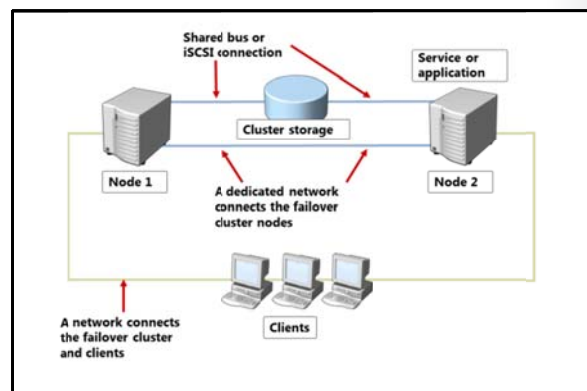
- The Cluster.exe command-line tool is deprecated. However, you can still optionally install it with the failover clustering tools. Failover clustering Windows PowerShell cmdlets provide a functionality that is generally the same as **cluster.exe** commands.
- The Cluster Automation Server (MSClus) Component Object Model (COM) interface is deprecated, but you can optionally install it with the failover clustering tools.
- The Support for 32-bit cluster resource dynamic-link libraries (DLLs) is deprecated, but you can optionally install 32-bit DLLs. Cluster resource DLLs should be updated to 64-bit.
- The Print Server role is removed from the High Availability Wizard, and it cannot be configured in the Failover Cluster Manager.
- The **Add-ClusterPrintServerRole** cmdlet is deprecated, and it is not supported in Windows Server 2012.

Failover Cluster Components

As a failover cluster, a group of independent computers work together to increase the availability of applications and services. Physical cables and software connect the clustered servers. Servers that participate in the cluster are also known as *nodes*. If one of the cluster nodes fails, another node begins to provide services. This process is known as *failover*. With failover, users experience minimal to no service disruptions.

A failover clustering solution consists of several components, which include:

- *Nodes*. These are computers that are members of a failover cluster. These computers run Cluster service, and resources and applications associated to cluster.



- *Network.* This is a network across which cluster nodes can communicate with one another, and with clients. There are three types of networks that can be used in a cluster: public, private, and public-and-private. These networks are discussed in more detail in the "Failover Cluster Networks" topic.
- *Resource.* This is an entity that is hosted by a node. It is managed by the Cluster service, and can be started, stopped, and moved to another node.
- *Cluster storage.* This is a storage system that is usually shared between cluster nodes. In some scenarios, such as clusters of servers running Microsoft® Exchange Server, shared storage is not required.
- *Clients.* These are computers (or users) that are using the Cluster service.
- *Service or application.* This is a software entity that is presented to clients and used by clients.

In a failover cluster, each node in the cluster:

- Has full connectivity and communication with the other nodes in the cluster.
- Is aware when another node joins or leaves the cluster. In addition, each node is aware when a node or resource is failing, and has the ability to take those services over.
- Is connected to a network through which client computers can access the cluster.
- Is usually connected through a shared bus or iSCSI connection to shared storage.
- Is aware of the services or applications that are running locally, and the resources that are running on all other cluster nodes.

Cluster storage usually refers to logical devices—typically hard disk drives or logical unit numbers (LUNs)—to which all the cluster nodes attach through a shared bus. This bus is separate from the bus that contains the system and boot disks. The shared disks store resources such as applications and file shares that the cluster will manage.

A failover cluster typically defines at least two data communications networks: one network enables the cluster to communicate with clients, and the second isolated network enables the cluster node members to communicate directly with one another. If a directly-connected shared storage is not being used, then a third network segment (for iSCSI or Fibre Channel) can exist between the cluster nodes and a data storage network.

Most clustered applications and their associated resources are assigned to one cluster node at a time. The node that provides access to those cluster resources is the *active node*. If the nodes detect the failure of the active node for a clustered application, or if the active node is taken offline for maintenance, the clustered application is started on another cluster node. To minimize the impact of the failure, client requests are immediately and transparently redirected to the new cluster node.

What Are CSVs?

In classic failover cluster deployment, only a single node at a time controls a LUN on the shared storage. This means that another node cannot see shared storage, until it becomes an active node. CSV is a new technology introduced in Windows Server 2008 R2, which enables multiple nodes to share a single LUN concurrently. Each node obtains exclusive access to individual files on the LUN instead of to the entire LUN. In other words, CSV provides a solution so that multiple nodes in the cluster can access the same NTFS file system simultaneously.

CSV benefits:

- Fewer LUNs required
- Better use of disk space
- Resources are in a single logical location
- No special hardware required
- Increased resiliency

To implement CSV:

- Create and format volumes on shared storage
- Add the disks to failover cluster storage
- Add the storage to the CSV

In the first version in Windows Server 2008 R2, CSV was designed only for hosting virtual machines that are running on a Hyper-V server in a failover cluster. This enabled administrators to have a single LUN that hosts multiple virtual machines in a failover cluster. Multiple cluster nodes have access to the LUN, but each virtual machine runs only on one node at a time. If the node on which a virtual machine is running fails, CSV enables the virtual machine to restart on a different node in the failover cluster. Additionally, this provides simplified disk management for hosting virtual machines, as compared to each virtual machine requiring a separate LUN.

In Windows Server 2012, CSV has additional improvements. It is now possible to use CSV for other roles, and not just Hyper-V. For example, you can now configure the file server role in failover clustering in the Scale-Out File Server scenario. Scale-Out File Server is designed to provide scale-out file shares that are continuously available for file-based server application storage. Scale-out file shares provide the ability to share the same folder from multiple nodes in the same cluster. In this context, CSV in Windows Server 2012 introduces support for a read cache, which can improve performance in certain scenarios. In addition, a CSV file system (CSVFS) can perform **Chkdsk** without impacting applications with open handles on the file system.

Other important improvements in CSV in Windows Server 2012 are:

- CSV proxy file system (CSVFS): In the Disk Management console, CSV volumes now appear as CSVFS. However, this is not a new file system. The underlying technology is still the NTFS file system, and CSVFS volumes are still formatted with NTFS. However, because volumes appear as CSVFS, applications can discover that they are running on CSVs, which helps improve compatibility. Additionally, because of the single file namespace, all files have the same name and path on any node in a cluster.
- Multisubnet support for CSVs: CSVs have been enhanced to integrate with SMB Multichannel to help achieve faster throughput for CSV volumes.
- Support for BitLocker volume encryption: Windows Server 2012 supports BitLocker volume encryption for both traditional clustered disks and CSVs. Each node performs decryption by using the computer account for the cluster itself.
- Support for SMB 3.0 storage: SMB 3.0 storage is supported for Hyper-V and applications such as Microsoft SQL Server®. This means that, for example, you can host Hyper-V virtual machine files on a shared folder.
- Integration with SMB Multichannel and SMB Direct: This integration allows CSV traffic to stream across multiple networks in the cluster and to leverage network adapters that support Remote Direct Memory Access (RDMA).

- Integration with the Storage Spaces feature in Windows Server 2012: This integration can provide virtualized storage on clusters of inexpensive disks.
- Reduced downtime: CSV in Windows Server 2012 lets you scan and repair volumes with zero offline time.

Implementing CSV

You can configure CSV only when you create a failover cluster. After you create the failover cluster, you can enable CSV for the cluster, and then add storage to the CSV.

However, before you can add storage to the CSV, you must make the LUN available as shared storage to the cluster. When you create a failover cluster, all of the shared disks that you configured in Server Manager are added to the cluster, and you can then add them to a CSV. If you add more LUNs to the shared storage, you must first create volumes on the LUN, add the storage to the cluster, and then add the storage to the CSV.

As a best practice, you should configure CSV before you make any virtual machines highly available. However, you can convert from regular disk access to CSV after deployment. The following considerations apply:

- When you convert from regular disk access to CSV, this removes the LUN's drive letter or mount point. This means that you must recreate all virtual machines that are stored on the shared storage. If you must retain the same virtual machine settings, consider exporting the virtual machines, switching to CSV, and then importing the virtual machines in Hyper-V.
- You cannot add shared storage to CSV if it is in use. If you have a running virtual machine that is using a cluster disk, you must shut down the virtual machine, and then add the disk to CSV.

What Are Failover and Failback?

Failover transfers from one node to another the responsibility for providing access to resources in a cluster. Failover can occur when an administrator intentionally moves resources to another node for maintenance, or due to unplanned downtime of a node due to hardware failure or other reasons. In addition, service failure on an active node can initiate failover to another node.

A failover attempt consists of the following steps:

1. The Cluster service takes all the resources in the instance offline, in an order that is determined by the instance's dependency hierarchy. This means that dependent resources are taken offline first, followed by the resources on which they depend. For example, if an application depends on a physical disk resource, the Cluster service takes the application offline first, which enables the application to write changes to the disk before the disk is taken offline.
2. After all the resources are offline, the Cluster service attempts to transfer the instance to the node that is listed next on the instance's list of preferred owners.
3. If the Cluster service successfully moves the instance to another node, it attempts to bring all the resources online. This time, it starts at the lowermost part of the dependency hierarchy. Failover is complete when all the resources are online on the new node.

During failover, the clustered instance and all associated resources are moved from one node to another

Failover occurs when:

- The node that currently hosts the instance becomes inactive for any reason
- One of the resources within the instance fails
- An administrator forces a failover

Cluster service can fail back after the offline node becomes active again

Once the offline node becomes active again, the Cluster service can fail back instances that were originally hosted on the offline node. When the Cluster service fails back an instance, it uses the same procedures that it performs during failover. The Cluster service takes all the resources in the instance offline, moves the instance, and then brings all the resources in the instance back online.

What Is a Quorum?

A *quorum* is the number of elements that must be online for a cluster to continue running. Each cluster node is an element, and in effect, each element can cast one vote to determine whether the cluster continues to run. If there is an even number of nodes, then an additional element—which is known as a *witness*—is assigned to the cluster. The witness element can be either a disk or a file share. Each voting element contains a copy of the cluster configuration, and the Cluster service works to keep all copies synchronized at all times.

In failover clusters, a quorum defines the consensus that enough cluster members are available to provide services


Quorum:

- Is based on votes
- Allows nodes, file shares, or a shared disk to have a vote, depending on the quorum mode
- Allows the failover cluster to remain online when sufficient votes are available

The cluster will stop providing failover protection if most of the nodes fail, or if there is a problem with communication between the cluster nodes. Without a quorum mechanism, each set of nodes could continue to operate as a failover cluster. This results in a partition within the cluster. Quorum prevents two or more nodes from concurrently operating a failover cluster resource. If a clear majority is not achieved between the node members, then the vote of the witness becomes crucial to maintain the validity of the cluster.

Concurrent operation could occur when network problems prevent one set of nodes from communicating with another set of nodes. That is, a situation might occur where more than one node tries to control access to a resource. If that resource is, for example, a database application, damage such as corruption of the database could result. Imagine the consequence if two or more instances of the same database are made available on the network, or if data was accessed and written to a target from more than one source at a time. If the application itself is not damaged, the data could easily become corrupted.

Because a given cluster has a specific set of nodes and a specific quorum configuration, the cluster can calculate the number of votes that are required for it to continue providing failover protection. If the number of votes drops below the majority, the cluster stops running, which means it will not provide failover protection if there is a node failure. Nodes will still listen for the presence of other nodes, in case another node appears again on the network. However, the nodes will not function as a cluster until a majority consensus or quorum is achieved.

 **Note:** A fully functioning cluster depends not just on quorum, but also on the capacity of each node to support the services and applications that fail over to that node. For example, a cluster that has five nodes could still have quorum after two nodes fail, but each remaining cluster node would continue serving clients only if it has enough capacity (such as disk space, processing power, random access memory (RAM), or network bandwidth) to support the services and applications that failed over to it. An important part of the design process is planning each node's failover capacity. A failover node must be able to run its own load and the load of additional resources that might fail over to it.

The Process of Achieving Quorum

Because a given cluster has a specific set of nodes and a specific quorum configuration, the cluster software on each node stores information about how many votes constitute a quorum for that cluster. If the number of votes drops below the majority, the cluster stops providing services. Nodes will continue to listen for incoming connections from other nodes on port 3343, in case they appear again on the network, but the nodes will not begin to function as a cluster until quorum is achieved.

There are several phases a cluster must complete to achieve a quorum. As a given node comes online, it determines whether there are other cluster members, with which it can communicate. This process may be in progress on multiple nodes simultaneously. After establishing communication with other members, the members compare their membership views of the cluster until they agree on one view (based on timestamps and other information). A determination is made whether this collection of members has a quorum, or has enough members the total of which creates sufficient votes so that a split scenario cannot exist. A split scenario means that another set of nodes that are in this cluster are running on a part of the network that is inaccessible to these nodes. Therefore, more than one node could be actively trying to provide access to the same clustered resource. If there are not enough votes to achieve quorum, the voters (the currently recognized members of the cluster) wait for more members to appear. After at least the minimum vote total is attained, the Cluster service begins to bring cluster resources and applications into service. With quorum attained, the cluster becomes fully functional.

Quorum Modes in Windows Server 2012 Failover Clustering

The quorum modes in Windows Server 2012 failover clustering are the same modes that are present in Windows Server 2008. As before, a majority of votes determines whether a cluster achieves quorum. Nodes can vote, and where appropriate, either a disk in cluster storage (known as a *disk witness*) or a file share (known as a *file share witness*) can vote. There is also a quorum mode called No Majority: Disk Only, which functions like the disk-based quorum in Windows Server 2003. Other than the No Majority: Disk Only mode, there is no single point of failure with the quorum modes, because only the number of votes is important and not whether a particular element is available to vote.

Quorum mode	What has the vote?	When is quorum maintained?
Node Majority	Only nodes in the cluster have a vote	Quorum is maintained when more than half of the nodes are online
Node and Disk Majority	The nodes in the cluster and a disk witness have a vote	Quorum is maintained when more than half of the votes are online
Node and File Share Majority	The nodes in the cluster and a file share witness have a vote	Quorum is maintained when more than half of the votes are online
No Majority: Disk Only	Only the quorum-shared disk has a vote	Quorum is maintained when the shared disk is online

The Windows Server 2012 failover clustering quorum mode is flexible. You can choose the mode best suited to your cluster. Be aware that most of the time it is best to use the quorum mode that the cluster software selects. If you run the Quorum Configuration Wizard, the quorum mode that the wizard lists as recommended is the quorum mode that the cluster software will choose. You should change the quorum configuration only if you have determined that the change is appropriate for your cluster.


Windows Server 2012 failover clustering has four quorum modes:

- **Node Majority.** Each node that is available and in communication with other nodes can vote. The cluster functions only with a majority (more than half) of the votes. This model is preferred when the cluster consists of an odd number of server nodes, and no witness is needed to maintain or achieve quorum.

- **Node and Disk Majority.** Each node plus the disk witness, which is a designated disk in the cluster storage, can vote when they are available and in communication. The cluster functions only with a majority (more than half) of the votes. This model is based on an even number of server nodes being able to communicate with one another in the cluster, in addition to the disk witness.
- **Node and File Share Majority.** Each node plus the file share witness, which is a designated file share created by the administrator, can vote when they are available and in communication. The cluster functions only with a majority (more than half) of the votes. This model is based on an even number of server nodes being able to communicate with one another in the cluster, in addition to the file share witness.
- **No Majority: Disk Only.** The cluster has quorum if one node is available and in communication with a specific disk in the cluster storage. Only the nodes that are also in communication with that disk can join the cluster.


Except for the No Majority: Disk Only mode, all quorum modes in Windows Server 2012 failover clusters are based on a simple majority vote model. As long as a majority of the votes are available, the cluster continues to function. For example, if there are five votes in the cluster, the cluster continues to function as long as there are at least three available votes. The source of the votes is not relevant—the vote could be a node, a disk witness, or a file share witness. The cluster will stop functioning if a majority of votes is not available.

In the No Majority: Disk Only mode, the quorum-shared disk can veto all other possible votes. In this mode, the cluster will continue to function as long as the quorum-shared disk and at least one node are available. This type of quorum also prevents more than one node from assuming the primary role.

 **Note:** If the quorum-shared disk is not available, the cluster will stop functioning, even if all nodes are still available. In the No Majority: Disk Only mode, the quorum-shared disk is a single point of failure, so this mode is not recommended.

When you configure a failover cluster in Windows Server 2012, the Installation Wizard automatically selects one of two default configurations. By default, failover clustering selects:

- Node Majority configuration if there is an odd number of nodes in the cluster.
- Node and Disk Majority configuration if there is an even number of nodes in the cluster.

 **Note:** You should modify this setting only if you determine that a change is appropriate for your cluster, and only once you understand the implications of making the change.

In addition to planning your quorum mode, you should also consider the capacity of the nodes in your cluster, and their ability to support the services and applications that may fail over to that node. For example, a cluster that has four nodes and a disk witness will still have quorum after two nodes fail. However, if you have several applications or services deployed on the cluster, each remaining cluster node may not have the capacity to provide services.

Failover Cluster Networks

Networks and network adapters are important parts of each cluster implementation. You cannot configure a cluster without configuring the networks that the cluster will use. A network can perform one of the following roles in a cluster:

- *Private network.* A private network carries internal cluster communication. By using this network, cluster nodes exchange heartbeats and check for another node or nodes. The failover cluster authenticates all internal communication. However, administrators who are especially concerned about security may want to restrict internal communication to physically secure networks.
- *Public network.* A public network provides client systems with access to cluster application services. IP address resources are created on networks that provide clients with access to the Cluster service.
- *Public-and-private network.* A public-and-private network (also known as a *mixed network*) carries internal cluster communication and connects clients to cluster application services.

Network	Description
Public network	Clients use this network to connect to the clustered service
Private network	Nodes use this network to communicate with each other
Public-and-private network	Required to communicate with external storage systems

- One network can support both client and node communications
- Multiple network cards are recommended to provide enhanced performance and redundancy

When you configure networks in failover clusters, you should also determine which network to connect to the shared storage. If you use iSCSI for the shared storage connection, the network will use an IP-based Ethernet communications network. However, you should not use this network for node or client communication. Sharing the iSCSI network in this manner may result in contention and latency issues for both users and for the resource that is being provided by the cluster.

Although not a best practice, you can use the private and public networks for both client and node communications. Preferably, you should dedicate an isolated network for the private node communication. The reasoning for this is similar to using a separate Ethernet network for iSCSI, which is to avoid resource bottleneck and contention issues. The public network is configured to enable client connections to the failover cluster. Although the public network can provide backup for the private network, a better design practice is to define alternative networks for the primary private and public networks, or at least team the network interfaces that are used for these networks.

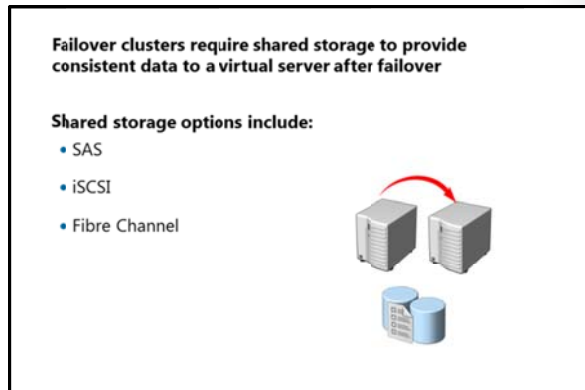
The networking features in Windows Server 2012–based clusters include the following:


- The nodes transmit and receive heartbeats by using User Datagram Protocol (UDP) unicast, instead of UDP broadcast (which was used in legacy clusters). The messages are sent on port 3343.
- You can include clustered servers on different IP subnets, which reduces the complexity of setting up multi-site clusters.
- The *Failover Cluster Virtual Adapter* is a hidden device that is added to each node when you install the failover clustering feature. The adapter is assigned a media access control (MAC) address based on the MAC address that is associated with the first enumerated physical network adapter in the node.
- Failover clusters fully support IPv6 for both node-to-node and node-to-client communication.
- You can use Dynamic Host Configuration Protocol (DHCP) to assign IP addresses, or you can assign static IP addresses to all nodes in the cluster. However, if some nodes have static IP addresses and you configure others to use DHCP, the Validate a Configuration Wizard will respond with an error. The cluster IP address resources are obtained based on the configuration of the network interface that is supporting that cluster network.

Failover Cluster Storage

Most failover clustering scenarios require shared storage to provide consistent data to a highly available service or application after failover. There are three shared-storage options for a failover cluster:

- *Shared serial attached SCSI (SAS)*. Shared SAS is the lowest-cost option. However, shared SAS is not very flexible for deployment because the two cluster nodes must be physically close together. In addition, the shared storage devices that are supporting shared SAS have a limited number of connections for cluster nodes.
- *iSCSI*. iSCSI is a type of storage area network (SAN) that transmits small computer system interface (SCSI) commands over IP networks. Performance is acceptable for most scenarios when the physical medium for data transmission is between 1 gigabit per second (Gbps) and 10 Gbps Ethernet. This type of SAN is fairly inexpensive to implement, because no specialized networking hardware is required. In Windows Server 2012, you can implement iSCSI target software on any server, and present local storage over iSCSI interface to clients.
- *Fibre Channel*. Fibre Channel SANs typically have better performance than iSCSI SANs, but are more expensive. Specialized knowledge and hardware are required to implement a Fibre Channel SAN.



 **Note:** The Microsoft iSCSI Software Target is now an integrated feature in Windows Server 2012. This feature can provide storage from a server over a TCP/IP network, including shared storage for applications that are hosted in a failover cluster. In addition, in Windows Server 2012, you can configure a highly available iSCSI Target Server as a clustered role by using Failover Cluster Manager or Windows PowerShell.

Storage Requirements

Before choosing the storage solution, you should also be aware of the following storage requirements:

- To use the native disk support that is included in failover clustering, use basic disks and not dynamic disks.
- You should format the partitions with NTFS. For the disk witness, the partition must be NTFS, because FAT is not supported.
- For the partition style of the disk, you can use either master boot record (MBR) or globally unique identifier (GUID) partition table (GPT).
- Because improvements in failover clustering require that the storage respond correctly to specific SCSI commands, the storage must follow the SCSI Primary Commands-3 (SPC-3) standard. In particular, the storage must support Persistent Reservations, as specified in the SPC-3 standard.
- The miniport driver used for the storage must work with the Storport storage driver. Storport offers a higher performance architecture and better Fiber Channel compatibility in Windows operating systems.

- You must isolate storage devices (one cluster per device). You should not allow servers that belong to different clusters to access the same storage devices. You can achieve this by using LUN masking or zoning. This prevents LUNs used on one cluster from being seen on another cluster. Consider using multipath input/output (I/O) software. Cluster nodes commonly use multiple host bus adapters to access storage, and this lets you achieve additional high availability. To be able to use multiple host bus adapters, you must use multipath software. For Windows Server 2012, your multipath solution must be based on Microsoft Multipath I/O (MPIO). Your hardware vendor usually supplies an MPIO device-specific module (DSM) for your hardware, although Windows Server 2012 includes one or more DSMs as part of the operating system.

Lesson 2

Implementing a Failover Cluster

Failover clusters in Windows Server 2012 have specific recommended hardware and software configurations that enable Microsoft to support the cluster. Failover clusters are intended to provide a higher level of service than stand-alone servers. Therefore, cluster hardware requirements are frequently stricter than requirements for stand-alone servers.

This lesson describes how to prepare for cluster implementation. In this lesson, you will also discuss the hardware, network, storage, infrastructure, and software requirements for Windows Server 2012 failover clusters. Finally, this lesson also outlines the steps for using the Validate a Configuration Wizard to ensure correct cluster configuration, and how to migrate failover clusters.

Lesson Objectives

After completing this lesson, you will be able to:

- Explain how to prepare for implementing failover clustering.
- Describe hardware requirements for failover clustering.
- Describe network requirements for failover clustering.
- Describe infrastructure requirements for failover clustering.
- Describe software requirements for failover clustering.
- Explain how to validate and configure a failover cluster.
- Explain how to migrate failover clusters.

Preparing for Failover Cluster Implementation

Before you implement failover clustering, you must identify services and applications that you want to make highly available. Failover clustering cannot be applied to all applications, and sometimes, applications have their own redundancy mechanisms. In addition, you should be aware that failover clustering does not provide improved scalability by adding nodes. You can only obtain scalability by scaling up and using more powerful hardware for the individual nodes. Therefore, you should only use failover clustering when your goal is high availability, and not scalability. In Windows Server 2012, there is one exception to this: if you implement File Services on CSVs, you can also achieve a level of scalability.

Failover clustering is best suited for stateful applications that are restricted to a single set of data. One example of such an application is a database. Data is stored in a single location and can only be used by one database instance. You can also use failover clustering for Hyper-V virtual machines. The best results for failover clustering occur when the client can reconnect to the application automatically after failover. If the client does not reconnect automatically, then the user must restart the client application.

Failover clustering uses only IP-based protocols and is, therefore, suited only to IP-based applications. Failover clustering now supports both IPv4 and IPv6.

Use failover clustering when:

- High availability is required
- Scalability is not required
- Application is stateful
- Client automatically reconnects to the application
- Application uses IP-based protocols



Consider the following guidelines when planning node capacity in a failover cluster:

- Spread out the highly available applications from a failed node. When all nodes in a failover cluster are active, the highly available services or applications from a failed node should be spread out among the remaining nodes to prevent a single node from being overloaded.
- Ensure that each node has sufficient idle capacity to service the highly available services or applications that are allocated to it when another node fails. This idle capacity should be a sufficient buffer to avoid nodes running at near capacity after a failure event. Failure to plan resource utilization adequately can result in a decrease in performance following node failure.
- Use hardware with similar capacity for all nodes in a cluster. This simplifies the planning process for failover, because the failover load will be evenly distributed among the surviving nodes.
- Use standby servers to simplify capacity planning. When a passive node is included in the cluster, then all highly available services or applications from a failed node can fail over to the passive node. This avoids the need for complex capacity planning. If this configuration is selected, it is important that the standby server has sufficient capacity to run the load from more than one node failure.

You should also examine all cluster configuration components to identify single points of failure. You can remedy many single points of failure with simple solutions, such as adding storage controllers to separate and stripe disks, or teaming network adapters, and using multipathing software. These solutions reduce the probability that a failure of a single device will cause a failure in the cluster. Typically, server class computer hardware has options for multiple power supplies for power redundancy, and for creating redundant array of independent disk (RAID) sets for disk data redundancy.

Hardware Requirements for Failover Cluster Implementation

It is very important to make good decisions when you select hardware for cluster nodes. Failover clusters have to satisfy the following criteria to meet availability and support requirements:

- All hardware that you select for a failover cluster should meet the Certified for Windows Server 2012 logo requirements. Hardware that has this logo has been independently tested to meet the highest technical bar for reliability, availability, stability, security, and platform compatibility. This logo also means that official support options exist in case malfunctions arise.
- You should install the same or similar hardware on each failover cluster node. For example, if you choose a specific model of network adapter, you should install this adapter on each of the cluster nodes.
- If you are using SAS or Fiber Channel storage connections, the mass storage device controllers that are dedicated to the cluster storage should be identical in all clustered servers. They should also use the same firmware version.
- If you are using iSCSI storage connections, each clustered server must have one or more network adapters or host bus adapters dedicated to the cluster storage. The network that you use for iSCSI storage connections should not be used for network communication. In all clustered servers, the network adapters that you use to connect to the iSCSI storage target should be identical, and we recommend that you use 1 Gbps Ethernet or more.

The hardware requirements for a failover implementation include the following:

- Server hardware components must be marked with the Certified for Windows Server 2012 logo
- Server nodes should all have the same configuration and contain the same or similar components
- All tests in the Validate a Configuration Wizard must be passed

- After you configure the servers with the hardware, all tests provided in the Validate a Configuration Wizard must be passed before the cluster is considered a configuration that will be supported by Microsoft.

Network Requirements for Failover Cluster Implementation

One of the network requirements for failover cluster implementation is that failover cluster network components must have the Certified for Windows Server 2012 logo, and must also pass the tests in the Validate a Configuration Wizard.

Additionally:

- The network adapters in each node should be identical and have the same IP protocol version, speed, duplex, and flow control capabilities that are available.
- The networks and network equipment to which you connect the nodes should be redundant so that even a single failure allows for the nodes to continue communicating with one another. You can use network adapter teaming to provide single network redundancy. We recommend multiple networks to provide multiple paths between nodes for inter-node communication. Otherwise, a warning will generate during the validation process.
- The network adapters in a cluster network must have the same IP address assignment method, which means either that they all use static IP addresses, or that they all use DHCP.

The network requirements for a failover implementation include the following:

- The network hardware components must be marked with the Certified for Windows Server 2012 logo
- The server should be connected to multiple networks for communication redundancy, or to a single network with redundant hardware, to remove single points of failure
- The network adapters should be identical and have the same IP protocol versions, speed, duplex, and flow control capabilities



Note: If you connect cluster nodes with a single network, the network passes the redundancy requirement in the Validate a Configuration Wizard. However, the report from the wizard will include a warning that the network should not have single points of failure.

Infrastructure Requirements for Failover Cluster Implementation

Failover clusters depend on infrastructure services. Each server node must be in the same Active Directory domain, and if you use Domain Name System (DNS), the nodes should use the same DNS servers for name resolution.

We recommend that you install the same Windows Server 2012 features and roles on each node. Inconsistent configuration on cluster nodes can cause instability and performance issues. In addition, you should not install the AD DS role on any of the cluster nodes, because AD DS has its own fault-tolerance mechanism. If you install the AD DS role on one of the nodes, you must install it on all nodes. However, this is not recommended.

The infrastructure requirements for a failover implementation include the following:

- The nodes in the cluster must use DNS for name resolution
- All servers in the cluster must be in the same Active Directory domain
- The user account that creates the cluster must have administrator rights and permissions on all servers, and the Create Computer Objects permission in the domain.

Failover cluster infrastructure recommendations include:

- The same roles should be installed on each cluster node
- The AD DS role should not be installed on any of the cluster nodes

You must have the following network infrastructure for a failover cluster:

- Network settings and IP addresses. When you use identical network adapters for a network, also use identical communication settings such as speed, duplex mode, flow control, and media type on those adapters. Also, compare the settings between the network adapter and the switch to which it connects, and ensure that no settings are in conflict. Otherwise, network congestion or frame loss might occur, which could adversely affect how the cluster nodes communicate among themselves, with clients or with storage systems.
- Unique subnets. If you have private networks that are not routed to the rest of the network infrastructure, ensure that each of these private networks uses a unique subnet. This is necessary even if you give each network adapter a unique IP address. In addition, these private network addresses should not be registered in DNS. For example, if you have a cluster node in a central office that uses one physical network, and another node in a branch office that uses a separate physical network, do not specify 10.0.0.0/24 for both networks, even if you give each adapter a unique IP address. This avoids routing loops and other network communications problems if, for example, the segments are accidentally configured into the same collision domain because of incorrect virtual local area network (VLAN) assignments.
- DNS. The servers in the cluster typically use DNS for name resolution. DNS dynamic update protocol is a supported configuration.
- Domain role. All servers in the cluster must be in the same Active Directory domain. As a best practice, all clustered servers should have the same domain role (either member server or domain controller). The recommended role is member server because AD DS inherently includes its own failover protection mechanism.
- Account for administering the cluster. To be able to administer the cluster, you must have an account with appropriate permissions. You must have local administrator rights on all nodes that participate in the cluster. In addition, when you create the cluster, you must have the right to create new objects in domains. You can do this with the Domain Admin account, or you can delegate these rights to another domain account.

In Windows Server 2012, there is no cluster service account. Instead, the Cluster service automatically runs in a special context that provides the specific permissions and credentials that are necessary for the service (similar to the local system context, but with reduced credentials). When a failover cluster is created and a corresponding computer object is created in AD DS, that object is configured to prevent accidental deletion. In addition, the cluster Network Name resource has an additional health check logic, which periodically verifies the health and properties of the computer object that represents the Network Name resource.


Software Requirements for Failover Cluster Implementation

Failover clusters require that each cluster node must run the same edition of Windows Server 2012. The edition can be either Windows Server 2012 Standard or Windows Server 2012 Datacenter. The nodes should also have the same software updates and service packs. Depending on the role that will be clustered, a Windows Server 2012 Server Core installation may also meet the software requirements. In Windows Server 2012, Server Core is the default installation option, and therefore you should consider it as a cluster node. However, Failover Clustering can also be installed on a full GUI version.

The software requirements for a failover implementation include the following:

- All nodes must run the same edition of Windows Server 2012, which can be any of the following:
 - Windows Server 2012 Standard, Full or Server Core installation
 - Windows Server 2012 Datacenter, Full or Server Core installation
- All nodes must run the same processor architecture (32-bit, x64-based, or Itanium architecture-based)
- All nodes should have the same service pack and updates

It is also important that the same version of service packs or any operating system updates exist on all nodes that are parts of a cluster.

 **Note:** Windows Server 2012 provides Cluster-Aware Updating technology that can help you maintain updates on cluster nodes. This feature will be discussed in more detail in “Lesson 4: Maintaining a Failover Cluster”.

Each node must run the same processor architecture. This means that each node must have the same processor family, which might be, for example, the Intel Xeon processor family with Extended Memory 64Technology, the AMD Opteron AMD64 family, or the Intel Itanium-based processor family.

Demonstration: Validating and Configuring a Failover Cluster

The Validate a Configuration Wizard runs tests that confirm if the hardware and software settings are compatible with failover clustering. Using the wizard, you can run the complete set of configuration tests or a subset of the tests. You should run the tests on servers and storage devices before you configure the failover cluster, and again after you make any major changes to the cluster. You can access the test results in the %windir%\cluster\Reports directory.

Demonstration Steps

Validate and Configure a Cluster

1. Start the Failover Cluster Manager on the LON-SVR3.
2. Start the Validate Configuration Wizard.
3. Review the report.
4. Create a new cluster. Add **LON-SVR3** and **LON-SVR4** as cluster nodes.
5. Name the cluster **Cluster1**.
6. Use **172.16.0.125** as the **IP address**.

Migrating Failover Clusters

In some scenarios, such as replacing cluster nodes, or upgrading to a newer version of a Windows operating system, you will need to migrate clustered roles or services from one cluster to another. In Windows Server 2012, it is possible to migrate clustered roles and cluster configuration from clusters running Windows Server 2012, Windows Server 2008 R2, or Windows Server 2008. You can migrate these roles and configurations in one of two ways:

- Migrate from an existing cluster to a new cluster that is running Windows Server 2012:
In this scenario, you have two new cluster nodes running Windows Server 2012, and you then perform migration from an existing cluster with nodes running Windows Server 2008 or later.
- Perform an in-place migration on a two-node cluster: This is a more complex scenario, where you want to migrate a cluster to a new version of the Windows operating system. In this scenario, you do not have additional computers for new cluster nodes. For example, you may want to upgrade a cluster that is currently running on Windows Server 2008 R2 to a cluster running Windows Server 2012. To achieve this, you must first remove resources from one node, and evict that node from a cluster. Next, you perform a clean installation of Windows Server 2012 on that server. After Windows Server 2012 is installed, you create a one-node failover cluster, migrate the clustered services and applications from the old cluster node to that failover cluster, and then remove the old node from cluster. The last step is to install Windows Server 2012 on another cluster node, together with failover cluster feature, add the server to the failover cluster, and run validation tests to confirm that the overall configuration works correctly.

The Cluster Migration Wizard is a tool that lets you perform the migration of clustered roles. Because the Cluster Migration Wizard does not copy data from one storage location to another, you must copy or move data or folders (including shared folder settings) during a migration. In addition, the Cluster Migration Wizard does not migrate mount-point information (information about hard disk drives that do not use drive letters, and are mounted in a folder on another hard disk drive). However, it can migrate physical disk resource settings to and from disks that use mount points.

You can migrate clustered roles from one cluster to another. You can perform migration by:

- Migrating clustered roles to a new cluster with new servers
- Performing in-place migration with only two nodes

Cluster Migration Wizard is used in the migration process. It migrates roles, but not data or folders.

Lesson 3

Configuring Highly Available Applications and Services on a Failover Cluster

After you have configured your clustering infrastructure, you should configure specific roles or services to be highly available. Not all roles can be clustered. Therefore, you should first identify the resource that you want to put in a cluster, and then verify whether that resource is supported. In this lesson, you will learn about configuring roles and applications in clusters, and you will learn about configuring cluster settings.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe and identify cluster resources and services.
- Describe the process for clustering server roles.
- Cluster a file server role.
- Explain how to configure failover cluster properties.
- Explain how to manage cluster nodes.
- Explain how to configure application failover settings.

Identifying Cluster Resources and Services

Clustered services are services or applications that are made highly available by installing them on a failover cluster. Clustered services are active on one node, but can be moved to another node. A clustered service that contains an IP address resource and a network name resource (and other resources) is published to a client on the network under a unique server name. Because this group of resources displays as a single logical server to clients, it is called a *clustered instance*.

Users access applications or services on an instance in the same manner as they would if the applications or services were on a non-clustered server. Usually, applications or users do not know that they are connecting to a cluster, or the node to which they are connected.

Resources are physical or logical entities—such as a file share, disk, or IP address—that the failover cluster manages. Resources are the most basic and smallest configurable units that may provide a service to clients, or may be important parts of the cluster. At any time, a resource can run only on a single node in a cluster, and is online on a node when it provides its service to that specific node.

Server Cluster Resources

A *cluster resource* is any physical or logical component that has the following characteristics:

- It can be brought online and taken offline.
- It can be managed in a server cluster.
- It can be hosted (owned) by only one node at a time.

Clustered services:

- Are services or applications that are made highly available by installing them on a failover cluster
- Are active on one node, but can be moved to another node

Resources:

- Are the components that make up a clustered service
- Can only run on one node at a time
- Are moved to another node when one node fails
- Include components such as shared disks, names, and IP addresses

To manage resources, the Cluster service communicates to a resource DLL through a resource monitor. When the Cluster service makes a request for a resource, the resource monitor calls the appropriate entry point function in the resource DLL to check and control the resource state.

Dependent Resources

A *dependent resource* is one that requires another resource to operate. For example, because a network name must be associated with an IP address, a network name is considered a dependent resource. Because of this requirement, a network name resource depends on an IP address resource. Dependent resources are taken offline before the resources upon which they depend are taken offline. Similarly, they are brought online after the resources on which they depend are brought online. A resource can specify one or more resources on which it is dependent. Resource dependencies also determine bindings. For example, clients will be bound to the particular IP address on which a network name resource depends.

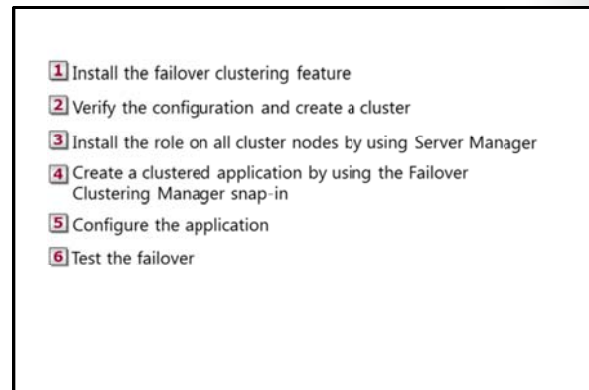
When you create resource dependencies, consider the fact that although some dependencies are strictly required, others are not required but are recommended. For example, a file share that is not a Distributed File System (DFS) root has no required dependencies. However, if the disk resource that holds the file share fails, the file share will be inaccessible to users. Therefore, it is logical to make the file share dependent on the disk resource.

A resource can also specify a list of nodes on which it can run. Possible nodes and dependencies are important considerations when administrators organize resources into groups.

Process for Clustering Server Roles

Failover clustering supports the clustering of several Windows Server roles, such as File Services, DHCP, and Hyper-V. To implement clustering for a server role, or for external applications such as Microsoft SQL Server or Exchange Server, perform the following procedure:

1. Install the failover clustering feature. Use Server Manager, `dism.exe`, or Windows PowerShell to install the failover clustering feature on all computers that will be cluster members. In Windows Server 2012, you can install roles and features on multiple servers simultaneously from single Server Manager console.
2. Verify configuration and create a cluster with the appropriate nodes. Use the Failover Cluster Manager snap-in to first validate a configuration, and then create a cluster with selected nodes.
3. Install the role on all cluster nodes. Use Server Manager, `dism.exe`, or Windows PowerShell to install the server role that you want to use in the cluster.
4. Create a clustered application by using the Failover Cluster Manager snap-in.
5. Configure the application. Configure options on the application that is being used in the cluster.
6. Test failover. Use the Failover Cluster Management snap-in to test failover by intentionally moving the service from one node to another.



After you create the cluster, you can monitor its status by using the Failover Cluster Management console, and manage available options.

Demonstration: Clustering a File Server Role

Demonstration Steps

Cluster a File Server Role

1. On LON-SVR3, add **Cluster Disk 2** as cluster storage for Cluster1.
2. Configure **File Server** as a clustered role. Configure a **File Server for general use**.
3. For the **Client Access Point** name, type **AdatumFS** with the address of **172.16.0.55**.
4. Use **Cluster Disk 2** for the storage for AdatumFS.

Configuring Failover Cluster Properties

Once you create a cluster, the newly created cluster has many properties that you can configure.

When you open Cluster Properties, you can configure cluster name or change the name, you can add various types of resources such as IP address and network name to the cluster, and you can also configure cluster permissions. By configuring permissions, you determine who can have full control over that specific cluster and who can just read the cluster configuration.

In addition, you can perform some standard management tasks on each cluster periodically, or on demand. These tasks range from adding and removing cluster nodes, to modifying the quorum settings. Some of the most frequently used configuration tasks include:

- **Managing cluster nodes:** For each node in a cluster, you can stop cluster service temporarily, pause the service, initiate remote desktop to the node, or evict the node from the cluster.
- **Managing cluster networks:** You can add or remove cluster networks, and configure networks that will be dedicated to inter-cluster communication.
- **Managing permissions:** By managing permissions, you can delegate rights to administer a cluster.
- **Configuring cluster quorum settings:** By configuring quorum settings, you determine the way in which quorum is achieved, as well as who can have vote in a cluster.
- **Migrating services and applications to a cluster:** You can implement existing services to the cluster and make them highly available.
- **Configuring new services and applications to work in a cluster:** You can implement new services to the cluster.
- **Removing a cluster:** You can remove a cluster if you decide to stop using clustering, or if you want to move the cluster to another set of nodes.

You can perform most of these administrative tasks by using the Failover Cluster Management console, or by using Windows PowerShell. However, **Cluster.exe**, which was used for some of these tasks in previous Windows Server operating system versions, is no longer supported in Windows Server 2012, and is not part of the default installation.

In Failover Cluster properties, you can configure cluster names, add resources, and configure cluster permissions

Common cluster management tasks include:

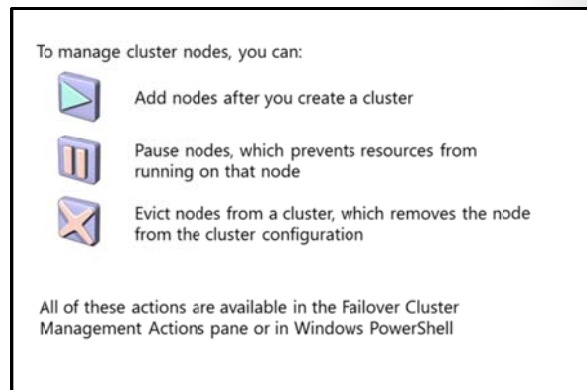
- Managing nodes
- Managing networks
- Managing permissions
- Configuring cluster quorum settings
- Migrating services and applications to a cluster
- Configuring new services and applications
- Removing the cluster

Managing Cluster Nodes

Cluster nodes are mandatory for each cluster. After you create a cluster and put it into production, you might have to manage cluster nodes occasionally.

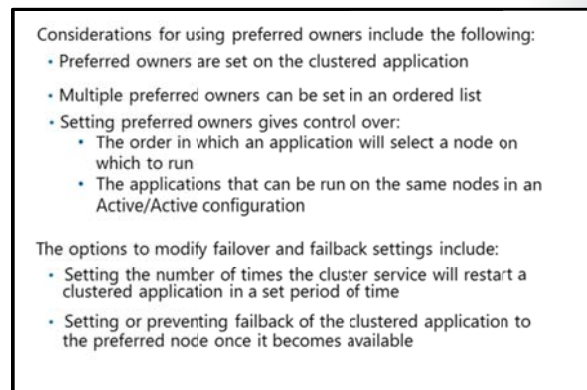
You can manage cluster nodes by using the Failover Cluster Management console or Windows PowerShell. There are three aspects to managing cluster nodes:

- You can add a node to an established failover cluster by selecting **Add Node** in the Failover Cluster Management Actions pane in the Failover Cluster Manager console. The Add Node Wizard prompts you for information about the additional node.
- You can pause a node to prevent resources from being failed over or moved to the node. You typically pause a node when it is undergoing maintenance or troubleshooting. When you are pausing a node, you can choose to drain roles from that node.
- You can evict a node, which is an irreversible process for a cluster node. After you evict the node, it must be re-added to the cluster. You evict nodes when a node is damaged beyond repair, or is no longer needed in the cluster. If you evict a damaged node, you can repair or replace it, and then add it back to the cluster by using the Add Node Wizard.



Configuring Application Failover Settings

You can adjust failover settings, including preferred owners and failback settings, to control how the cluster responds when the application or service fails. You can configure these settings on the property sheet for the clustered service or application (either on the **General** tab or on the **Failover** tab).



The following table provides examples that show how these settings work.

Setting	Result
Example 1: General tab, Preferred owner: Node1 Failover tab, Failback setting: Allow failback (Immediately)	If the service or application fails over from Node1 to Node2, when Node1 is once again available, the service or application will fail back to Node1.
Example 2: Failover tab, Maximum failures in the specified period: 2 Failover tab, Period (hours): 6	<p>In a six-hour period, if the application or service fails no more than two times, it will be restarted or failed over every time. If the application or service fails a third time in the six-hour period, it will be left in the failed state.</p> <p>The default value for the maximum number of failures is $n-1$, where n is the number of nodes. You can change the value, but we recommend a fairly low value so that if multiple node failures occur, the application or service will not be moved between nodes indefinitely.</p>

Lesson 4

Maintaining a Failover Cluster

Once your cluster infrastructure is running, it is important that you establish monitoring to prevent possible failures. In addition, it is important that you have backup and restore procedures for cluster configuration. Windows Server 2012 has new technology that allows you to update cluster nodes without downtime. In this lesson, you will learn about monitoring failover clusters, backing up and restoring cluster configurations, and updating cluster nodes.

Lesson Objectives

After completing this lesson, you will be able to:

- Explain how to monitor failover clusters.
- Explain how to back up and restore a failover cluster configuration.
- Explain how to maintain and troubleshoot failover clusters.
- Describe Cluster-Aware Updating.
- Configure Cluster-Aware Updating.

Monitoring Failover Clusters

Many tools are available in Windows Server 2012 to help you monitor failover clusters. You can use standard Windows Server tools such as the Event Viewer and the Performance and Reliability Monitor snap-in to review cluster event logs and performance metrics. You can also use Tracerpt.exe to export data for analysis. Additionally, you can use the Multipurpose Internet Mail Extension Hypertext Markup Language (MHTML)-formatted cluster configuration reports and the Validate a Configuration Wizard to troubleshoot problems

with the cluster configuration and hardware changes. Since the cluster.exe command-line tool is deprecated in Windows Server 2012, you can use Windows PowerShell instead to perform similar tasks.

Some of the tools you can use to monitor clusters include:

- Event Viewer
- Tracerpt.exe
- Performance and Reliability Monitor snap-in
- MHTML-formatted cluster configuration reports
- Validate a Configuration Wizard

Event Viewer

When problems arise in the cluster, use the Event Viewer to view events with a Critical, Error, or Warning severity level. Additionally, informational-level events are logged to the failover clustering Operations log, which can be found in the Event Viewer in the Applications and Services Logs\Microsoft\Windows folder. Informational-level events are usually common cluster operations, such as cluster nodes leaving and joining the cluster, or resources going offline or coming online.

In previous Windows Server versions, event logs were replicated to each node in the cluster. This simplified cluster troubleshooting, because you could review all event logs on a single cluster node. Windows Server 2012 does not replicate the event logs between nodes. However, the Failover Cluster Management snap-in has a **Cluster Events** option that enables you to view and filter events across all cluster nodes. This feature is helpful in correlating events across cluster nodes. The Failover Cluster Management snap-in also provides a **Recent Cluster Events** option that will query all the Error and Warning events from all the cluster nodes in the last 24 hours. You can access additional logs, such as the

Analytic and Debug logs, in the Event Viewer. To display these logs, modify the view on the top menu by selecting the **Show Analytic and Debug Logs** option.

Windows Event Tracing

Windows event tracing is a kernel component that is available early after startup, and late into shutdown. It is designed to allow for fast tracing and delivery of events, to trace files and to consumers. Because it is designed to be fast, Windows event tracing enables only basic in-process filtering of events based on event attributes.

The event trace log contains a comprehensive accounting of the failover cluster actions. Depending on how you want to view the data, use Windows PowerShell or Tracerpt.exe to access the information in the event trace log.

Tracerpt.exe will parse the event trace logs only on the node on which it is run. All the individual logs are collected in a central location. To transform the XML file into a text file or an HTML file that can be opened in Windows Internet Explorer, you can parse the XML-based file by using the Microsoft XSL parsing command prompt msxsl.exe tool, and an XSL style sheet.

Performance and Reliability Monitor Snap-In

The Performance and Reliability Monitor snap-in lets you:

- Trend application performance on each node. To determine how an application is performing, you can view and trend specific information on system resources that are being used on each node.
- Trend application failures and stability on each node. You can pinpoint when application failures occur, and match the application failures with other events on the node.
- Modify trace log settings. You can start, stop, and adjust trace logs, including their size and location.

Backing Up and Restoring Failover Cluster Configuration

Cluster configuration can be a time-consuming process with many details. Therefore, backing up your cluster configuration is important. You can perform cluster configuration backup and restore using Windows Server Backup or a non-Microsoft backup tool.

When you back up the cluster configuration, be aware of the following:

- You must test your backup and recovery process before putting a cluster into production.
- You must first add the Windows Server Backup feature, if you decide to use it. You can do this by using Server Manager, by using the `dism.exe` utility, or by using Windows PowerShell.

Windows Server Backup is the built-in backup and recovery software for Windows Server 2012. To complete a successful backup, consider the following:

- For a backup to succeed in a failover cluster, the cluster must be running and must have quorum. In other words, enough nodes must be running and communicating (perhaps with a witness disk or witness file share—depending on the quorum configuration,) that the cluster has achieved quorum.
- You must back up all clustered applications. If you cluster a SQL Server database, you must have a backup plan for the databases and configuration outside the cluster configuration.

When backing up failover clusters, keep in mind that:

- Windows Server Backup is a Windows Server 2012 feature
- You install Windows Server Backup as a feature
- Backup and restore operations involve the Volume Shadow Copy Service
- Non-Microsoft tools are available to perform backups and restores
- You must perform system state backups

A non-authoritative restore restores a single node in the cluster

An authoritative restore restores the entire cluster configuration to a point in time

- If application data must be backed up, the disks on which you store the data must be made available to the backup software. You can achieve this by running the backup software from the cluster node that owns the disk resource, or by running a backup against the clustered resource over the network. If you are using CSVs, you can run backup from any node that is attached to the CSV volume.
- The cluster service tracks which cluster configuration is the most recent, and it replicates that configuration to all cluster nodes. If the cluster has a witness disk, the Cluster service also replicates the configuration to the witness disk.

Restoring a Cluster

There are two types of restore:

- **Non-authoritative restore.** Use a non-authoritative restore when a single node in the cluster is damaged or rebuilt, and the rest of the cluster is operating correctly. Perform a non-authoritative restore by restoring the system recovery (system state) information to the damaged node. When you restart that node, it will join the cluster and receive the latest cluster configuration automatically.
- **Authoritative restore.** Use an authoritative restore when the cluster configuration must be rolled back to a previous point in time. For example, use an authoritative restore if an administrator accidentally removed clustered resources or modified other cluster settings. Perform the authoritative restore by stopping the cluster resource on each node, and then performing a system recovery (system state) on a single node by using Windows Server Backup interface. After the restored node restarts the cluster service, the remaining cluster nodes can also start the cluster service.

Maintaining and Troubleshooting Failover Clusters

Cluster validation functionality implemented in Windows Server 2012 failover clustering, prevents misconfigurations and non-working clusters. However, in some cases you may still have to perform maintenance or cluster troubleshooting.

Some common maintenance tasks can help you prevent problems in cluster configuration:

- Use the Validate a Configuration Wizard to highlight configuration issues that might cause cluster problems.
- Review cluster events and trace logs to identify application or hardware issues that might cause an unstable cluster.
- Review hardware events and logs to help pinpoint specific hardware components that might cause an unstable cluster.
- Review SAN components, switches, adapters, and storage controllers to help identify any potential problems.

When troubleshooting failover clusters:


- Identify the perceived problem by collecting and documenting the symptoms of the problem.
- Identify the scope of the problem so that you can understand what is being affected by the problem, and the impact of that effect on the application and the clients.

The failover cluster maintenance techniques include:

- ✓ Reviewing events in logs (cluster, hardware, storage)
- ✓ Using the Validate a Configuration Wizard
- ✓ Defining a process for troubleshooting failover clusters
- ✓ Reviewing storage configuration
- ✓ Checking for group and resource failures

- Collect information so that you can accurately understand and pinpoint the possible problem. After you identify a list of possible problems, you can prioritize them by probability, or by the impact of a repair. If you cannot pinpoint the problem, you should attempt to re-create the problem.
- Create a schedule for repairing the problem. For example, if the problem only affects a small subset of users, you can delay the repair to an off-peak time so that you can schedule downtime.
- Complete and test each repair one at a time so that you can identify the fix.

To troubleshoot SAN issues, start by checking physical connections and by checking each of the hardware component logs. Additionally, run the Validate a Configuration Wizard to verify that the current cluster configuration is still supportable.

 **Note:** When you run the Validate a Configuration Wizard, ensure that the storage tests that you select can be run on an online failover cluster. Several of the storage tests cause loss of service on the clustered disk when the tests are run.

Troubleshooting Group and Resource Failures

To troubleshoot group and resource failures:

- Use the Dependency Viewer in the Failover Cluster Management snap-in to identify dependent resources.
- Check the Event Viewer and trace logs for errors from the dependent resources.
- Determine whether the problem only happens on a specific node or nodes, by trying to re-create the problem on different nodes.

What Is Cluster-Aware Updating?

Applying Windows operating system updates to nodes in a cluster requires extra attention. If you want to provide zero downtime for a clustered role, you must update cluster nodes manually one after another, and you must move resources manually from the node that you are updating to another node. This procedure can be very time-consuming. In Windows Server 2012, Microsoft has implemented a new feature for automatic update of cluster nodes.

Cluster-Aware Updating (CAU) is a Windows Server 2012 feature that lets administrators update cluster nodes automatically, with little or no loss in availability during the update process. During an update procedure, CAU transparently takes each cluster node offline, installs the updates and any dependent updates, performs a restart if necessary, brings the node back online, and then moves to update the next node in a cluster.

For many clustered roles, this automatic update process triggers a planned failover, and it can cause a transient service interruption for connected clients. However, for continuously available workloads in Windows Server 2012—such as Hyper-V with live migration or file server with SMB Transparent Failover—CAU can orchestrate cluster updates with no effect on service availability.

Cluster-Aware Updating is an automated feature specific to Windows Server 2012, which updates nodes in a cluster with minimal or zero downtime

Cluster-Aware Updating can work in two modes:

- Remote-updating mode
- Self-updating mode

Cluster Updating Modes

CAU can orchestrate the complete cluster updating operation in two modes:

- *Remote-updating mode.* In this mode, a computer that is running Windows Server 2012 or Windows 8 is called and configured as an orchestrator. To configure a computer as a CAU orchestrator, you must install failover clustering administrative tools on it. The orchestrator computer is not a member of the cluster that is updated during the procedure. From the orchestrator computer, the administrator triggers on-demand updating by using a default or custom Updating Run profile. Remote-updating mode is useful for monitoring real-time progress during the Updating Run, and for clusters that are running on Server Core installations of Windows Server 2012.
- *Self-updating mode.* In this mode, the CAU clustered role is configured as a workload on the failover cluster that is to be updated, and an associated update schedule is defined. In this scenario, CAU does not have a dedicated orchestrator computer. The cluster updates itself at scheduled times by using a default or custom Updating Run profile. During the Updating Run, the CAU orchestrator process starts on the node that currently owns the CAU clustered role, and the process sequentially performs updates on each cluster node. In the self-updating mode, CAU can update the failover cluster by using a fully automated, end-to-end updating process. An administrator can also trigger updates on demand in this mode, or use the remote-updating approach, if desired. In the self-updating mode, an administrator can access summary information about an Updating Run in progress by connecting to the cluster and running the **Get-CauRun** Windows PowerShell cmdlet.

To use CAU, you must install the failover clustering feature in Windows Server 2012, and you must create a failover cluster. The components that support CAU functionality are installed automatically on each cluster node.

You must also install the CAU tools, which are included in the failover clustering Tools (which are also part of the Remote Server Administration Tools (RSAT)). The CAU tools consist of the CAU user interface (UI) and the CAU Windows PowerShell cmdlets. The failover clustering Tools are installed by default on each cluster node when you install the failover clustering feature. You can also install these tools on a local or a remote computer that is running Windows Server 2012 or Windows 8, and that has network connectivity to the failover cluster.

Demonstration: Configuring CAU

Demonstration Steps

Configure CAU

1. Make sure that the failover cluster is configured and running on LON-SVR3 and LON-SVR4.
2. Add **failover clustering Feature** to LON-DC1.
3. Run **Cluster-Aware Updating** on LON-DC1, and configure it to connect to **Cluster1**.
4. Preview updates that are available for nodes LON-SVR3 and LON-SVR4.
5. Review available options for the Updating Run profile.
6. Apply available updates to **Cluster1** from LON-DC1.
7. After updates are applied, configure **Add CAU Clustered Role with Self-Updating Enabled** on LON-SVR3.

Lesson 5

Implementing a Multi-Site Failover Cluster

In some scenarios, you may have to deploy cluster nodes on different sites. Usually, you do this when you build disaster recovery solutions. In this lesson, you will learn about multi-site failover clusters, and the prerequisites for implementing them. You will also learn about synchronous and asynchronous replication, and the process of choosing a quorum mode for multi-site clusters.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe a multi-site failover cluster.
- Describe prerequisites for implementing a multi-site cluster.
- Describe synchronous and asynchronous replication.
- Explain how to choose a quorum mode for multi-site clusters.
- Describe the process for deploying multi-site clusters.
- Describe the challenges for implementing multi-site clusters.

What Is a Multi-Site Failover Cluster?

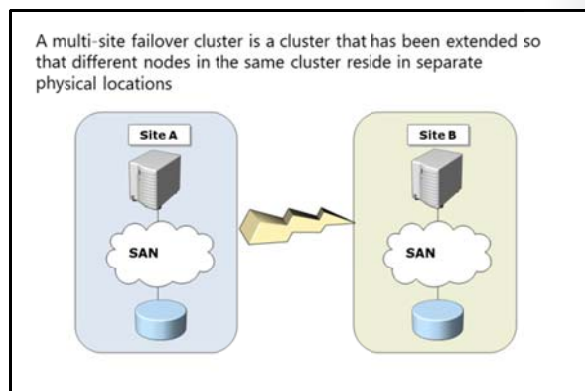
A multi-site failover cluster is a cluster that has been extended so that different nodes in the same cluster reside in separate physical locations. A multi-site failover cluster thereby provides highly available services in more than one location. Multi-site failover clusters can solve several specific problems, but they also present specific challenges.

In a multi-site failover cluster, each site usually has a separate storage system with replication between the sites. Multi-site cluster storage replication enables each site to be independent, and provides fast access to the local disk. With separate storage systems, you cannot share a single disk between sites.

A multi-site failover cluster has three main advantages in a failover site, as compared to a remote server:

- When a site fails, a multi-site cluster automatically fails over the clustered service or application to another site.
- Because the cluster configuration is replicated automatically to each cluster node in a multi-site cluster, there is less administrative overhead than a cold standby server, which requires you to replicate changes manually.
- The automated processes in a multi-site cluster reduce the possibility of human error, which is present in manual processes.

Because of increased cost and complexity of a multi-site failover cluster, it might not be an ideal solution for every application or business. When you are considering whether to deploy a multi-site cluster, you should evaluate the importance of the applications to the business, the type of applications, and any alternative solutions. Some applications can easily provide multi-site redundancy with log shipping or



other processes, and can still achieve sufficient availability with only a modest increase in cost and complexity. Examples for this are SQL Server log shipping, Exchange Server continuous replication, and DFS replication.

The complexity of a multi-site cluster requires more architectural and hardware planning. It also requires you to develop business processes to test the cluster functionality routinely.

Prerequisites for Implementing a Multi-Site Failover Cluster

Prerequisites for implementation of multi-site cluster are different from those for single-site cluster implementation. It is important to understand what you must prepare before you start implementation of multi-site cluster.

Prior to implementing multi-site failover cluster, you must ensure the following:

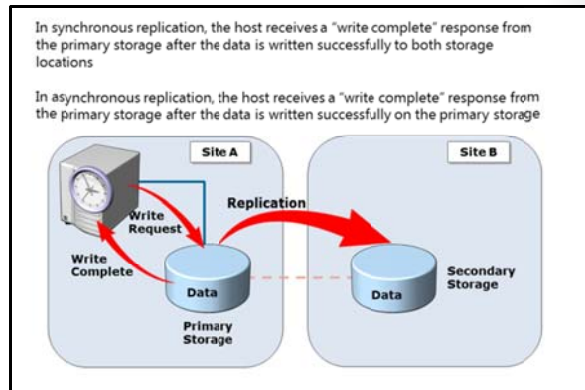
- You must have enough nodes and votes on each site, so that the cluster can be online even if one site is down. This setup requires additional hardware, and can come with significant financial costs.
- All nodes must have the same operating system and service pack version.
- You must provide at least one low-latency and reliable network connection between sites. This is important for cluster heartbeats. By default, regardless of subnet configuration, heartbeat frequency (also known as subnet delay) is once every second (1,000 milliseconds). The range for heartbeat frequency is once every 250-2000 milliseconds on a common subnet, and 250-4,000 milliseconds across subnets. By default, when a node misses a series of 5 heartbeats, another node will initiate failover. The range for this value (also known as subnet threshold) is from 3 through 10.
- You must provide a storage replication mechanism. Failover clustering does not provide any storage replication mechanism, so you must provide another solution. This also requires that you have multiple storage solutions, one for each cluster you create.
- You must ensure that all other necessary services for cluster, such as AD DS and DNS are also available on a second site.
- You must ensure that client connections can be redirected to a new cluster node when failover happens.

To implement a multi-site failover cluster, you must provide the following:

- Additional hardware to ensure enough nodes on each site
- Same operating systems and service packs on each node
- Reliable low-latency intercluster network
- Storage replication mechanism
- Infrastructure services on each site

Synchronous and Asynchronous Replication

It is not possible for a geographically dispersed failover cluster to use shared storage between physical locations. Wide area network (WAN) links are too slow and have too much latency to support shared storage. This means that you must have separate instances of data. To have exact copies of data on both sides, geographically dispersed failover clusters must synchronize data between locations by using specialized hardware. Multi-site data replication can be either synchronous or asynchronous:



- When you use synchronous replication, the host receives a *write complete* response from the primary storage after the data is written successfully on both storage systems. If the data is not written successfully to both storage systems, the application must attempt to write to the disk again. With synchronous replication, both storage systems are identical.
- When you use asynchronous replication, the node receives a write complete response from the storage after the data is written successfully on the primary storage. The data is written to the secondary storage on a different schedule, depending on the hardware or software vendor's implementation. Asynchronous replication can be storage-based, host-based, or even application-based. However, not all forms of asynchronous replication are sufficient for a multi-site cluster. For example, DFS Replication provides file-level asynchronous replication. However, it does not support multi-site failover clustering replication. This is because DFS Replication replicates smaller documents that are not held open continuously, and was not designed for high-speed, open-file replication.

When to Use Synchronous or Asynchronous Replication

Use synchronous replication when data loss is not acceptable. Synchronous replication solutions require low-disk write latency, because the application waits for both storage solutions to acknowledge the data writes. The requirement for low latency disk writes also limits the distance between the storage systems, because increased distance can cause higher latency. If the disk latency is high, the performance and even the stability of the application can be affected.

Asynchronous replication overcomes latency and distance limitations by acknowledging local disk writes only, and by reproducing the disk write on the remote storage system in a separate transaction. Because asynchronous replication writes to the remote storage system after it writes to the local storage system, the possibility of data loss during a failure increases.

Selecting a Quorum Mode for Multi-Site Clusters

Each failover cluster must have quorum mode defined, so that a majority vote can be easily determined at any time. For a geographically dispersed cluster, you cannot use quorum configurations that require a shared disk, because geographically dispersed clusters do not use shared disks. Both the Node and Disk Majority and No Majority: Disk Only quorum modes require a shared witness disk to provide a vote for determining quorum. You should only use these two quorum modes if the hardware vendor specifically recommends and supports them.

When designing automatic failover for geographically dispersed clusters:

- Use Node Majority or Node Majority with File Share quorum
- Use three locations to allow automatic failover of a single virtual server:
 - All three locations must be linked directly to each other
 - One location is only a file-share witness

To use the Node and Disk Majority and No Majority: Disk Only modes in a multi-site cluster, the shared disk requires that:

- You preserve the semantics of the SCSI commands across the sites, even if a complete communication failure occurs between sites.
- You replicate the witness disk in real-time synchronous mode across all sites.

Because multi-site clusters can have WAN failures in addition to node and local network failures, Node Majority and Node and File Share Majority are better solutions for multi-site clusters. If there is a WAN failure that causes the primary and secondary sites to lose communication, a majority must still be available to continue operations.


If there are an odd number of nodes, then use the Node Majority quorum. If there is an even number of nodes, which is typical in a geographically-dispersed cluster, you can use the Node Majority with File Share Majority quorum.

If you are using Node Majority and the sites lose communication, you need a mechanism to determine which nodes remain in the cluster, and which nodes leave the cluster. The second site requires another vote to obtain quorum after a failure. To obtain another vote for quorum, you must join another node to the cluster, or create a file share witness.

The Node and File Share Majority mode can help maintain quorum without adding another node to the cluster. To provide for a single-site failure and enable automatic failover, the file share witness might have to exist at a third site. In a multi-site cluster, a single server can host the file share witness. However, you must create a separate file share for each cluster.

You must use three locations to enable automatic failover of a highly available service or application. Locate one node in the primary location that runs the highly available service or application. Locate a second node in a disaster-recovery site, and locate the third node for the file share witness in a different location.

There must be direct network connectivity between all three locations. In this manner, if one site becomes unavailable, the two remaining sites can still communicate and have enough nodes for a quorum.

 **Note:** In Windows Server 2008 R2, administrators could configure the quorum to include nodes. However, if the quorum configuration included nodes, all nodes were treated equally according to their votes. In Windows Server 2012, you can adjust cluster quorum settings so that when the cluster determines whether it has quorum, some nodes have a vote and some do not. This adjustment can be useful when you implement solutions across multiple sites.

Process for Configuring a Multi-Site Failover Cluster

Configuration of multi-site cluster is somewhat different from configuring a single-site cluster. Multi-site clusters are more complex to configure and maintain, and require more administrative effort to support. High-level steps to configure a multi-site cluster are as follows:

1. Ensure that you have enough cluster nodes on each site. In addition, ensure that cluster nodes have similar hardware configurations, and have the same version of operating system and service pack.
2. Ensure that networking between sites is operational, and that network latency is acceptable for configuring the cluster. (You can validate this by using Validate Configuration Wizard in Failover Cluster Manager.)
3. Ensure that you have deployed reliable storage replication mechanism between sites. Also, choose the type of replication for use.
4. Ensure that key infrastructure services such as AD DS, DNS, and DHCP, are present on each site.
5. Run the Validate a Configuration Wizard on all of the cluster nodes to determine if your configuration is acceptable for creating a cluster.
6. Determine the role that you will configure in a cluster.
7. Determine the cluster quorum mode that you will use.
8. Create a clustered role.
9. Configure failover/failback settings.
10. Validate failover and failback.

You should be aware that multi-site clusters require more administrative effort during failover and failback. While single-site cluster failover/failback is mostly automatic, with multi-site clusters this is not the case.

Challenges for Implementing a Multi-Site Cluster

Implementing multi-site clusters is more complex than implementing single-site clusters, and can also present several challenges to the administrator. Storage and network issues are the most challenging aspects of implementing multi-site clusters.

In a multi-site cluster, there is no shared storage that the cluster node uses. This means that every node on each site must have its own storage instance. On the other hand, failover clustering does not include any built-in functionality to replicate data between sites. There are three

High level steps for implementing a multi-site failover cluster:

1. Ensure that enough nodes are available
2. Ensure that network connections between sites is reliable
3. Provide a storage replication mechanism
4. Provide key infrastructure services on both sites
5. Validate cluster configuration
6. Configure the clustered role and quorum
7. Configure and validate failover/failback

Storage Challenge	Description
Requires a separate or non-Microsoft data replication solution	<ul style="list-style-type: none"> • Hardware (block level) storage-based replication • Software (file system level) host-based replication • Application-based replication (such as Exchange 2007 Cluster Continuous Replication)
Can be either synchronous or asynchronous replication	<ul style="list-style-type: none"> • Synchronous: No acknowledgement of data changes made in Site A until the data is successfully written to Site B • Asynchronous: Data changes made in Site A will eventually be written to the storage in Site B
<ul style="list-style-type: none"> • Intra-node communications are time-sensitive. You might need to configure these thresholds to meet the higher WAN latency • DNS replication might impact client reconnect times when failover is based on hostname • Active Directory replication latency might affect application data availability • Some applications might require all of the nodes to be in the same Active Directory site 	

options for replicating data: block level hardware-based replication, software-based file replication installed on the host, or application-based replication.

Multi-site data replication can be either synchronous or asynchronous. Synchronous replication does not acknowledge data changes that are made in, for example, Site A until the data successfully writes to Site B. With asynchronous replication, data changes that are made in Site A are eventually written to Site B.

When you deploy a multi-site cluster and run the Validate a Configuration Wizard, the disk tests will not find any shared storage, and will therefore not run. However, you can still create a failover cluster. If you follow the hardware manufacturer's recommendations for Windows Server failover clustering hardware, Microsoft will support the solution.

Windows Server 2012 enables cluster nodes to exist on different IP subnets, which enables a clustered application or service to change its IP address based on that IP subnet. DNS updates the clustered application's DNS record so that clients can locate the IP address change. Because clients rely on DNS to find a service or application after a failover, you might have to adjust the DNS records' **Time to Live** setting, and the speed at which DNS data replicates. Additionally, when cluster nodes are in multiple sites, network latency might require you to modify the inter-node communication (heartbeat) delay and time-out thresholds.

Lab: Implementing Failover Clustering

Scenario

As A. Datum Corporation's business grows, it is becoming increasingly important that many of the applications and services on the network be available at all times. A. Datum has many services and applications that need to be available to their internal and external users who are working in different time zones around the world. Because many of these applications cannot be made highly available by using Network Load Balancing, you will need to use a different technology.

As one of the senior network administrators at A. Datum Corporation, you are responsible for implementing failover clustering on the Windows Server 2012 servers, to provide high availability for network services and applications. You will be responsible for planning the Failover Cluster configuration, and deploying applications and services on the Failover Cluster.

Objectives

- Configure a failover cluster with CSV storage.
- Deploy and configure a highly available file server on the failover cluster.
- Validate the high availability of the failover cluster and storage.
- Configure Cluster-Aware Updating on the failover cluster.

Lab Setup

Estimated Time: 60 minutes

- 20412-LON-DC1
- 20412-LON-SVR1
- 20412-LON-SVR3
- 20412-LON-SVR4
- MSL-TMG1

Username: **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 **20412A-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on using the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
5. Repeat steps 2-4 for **20412A-LON-SVR1**, **20412A-LON-SVR3** and **20412A-LON-SVR4**.
6. For **MSL-TMG1**, just repeat step 2.

Exercise 1: Configuring a Failover Cluster

Scenario

A. Datum Corporation has some critical applications and services that they want to make highly available. Some of these services cannot use Network Load Balancing. Therefore, you have decided to implement failover clustering with the use of iSCSI storage, which is already in place. To start this process, you need to implement the core components for failover clustering, validate the cluster, and then create the failover cluster.

The main tasks for this exercise are as follows:

1. Connect cluster nodes to the iSCSI targets.
2. Install the failover clustering feature.
3. Validate the servers for failover clustering.
4. Create the failover cluster.
5. Configure Cluster Shared Volumes.

► Task 1: Connect cluster nodes to the iSCSI targets

1. On LON-SVR3, start **iSCSI Initiator**, and configure **Discover Portal** with IP address **172.16.0.21**.
2. Connect to the discovered target in the **Targets** list.
3. Repeat steps 1 and 2 on LON-SVR4.
4. Open **Disk Management** on LON-SVR3.
5. Bring online and initialize the three new disks.
6. Make a simple volume on each disk and format it with NTFS.
7. On LON-SVR4, open **Disk Management**, refresh the console and bring online the three new disks.

► Task 2: Install the failover clustering feature

1. On LON-SVR3, install the **failover clustering** feature by using Server Manager.
2. On LON-SVR4, install the **failover clustering** feature by using Server Manager.

► Task 3: Validate the servers for failover clustering

1. On LON-SVR3, open the Failover Cluster Manager console.
2. Start the **Validate a Configuration Wizard**.
3. Use LON-SVR3 and LON-SVR4 as nodes for test.
4. Review the report and then close the report.
5. On the **Summary** page, remove the check mark next to **Create the cluster now using the validated nodes**, click **Finish**.

► Task 4: Create the failover cluster

1. On LON-SVR3, in the Failover Cluster Manager, start the **Create Cluster Wizard**.
2. Use LON-SVR3 and LON-SVR4 as cluster nodes.
3. Specify **Cluster1** as the **Access Point name**.
4. Specify the **IP address** as **172.16.0.125**.

► Task 5: Configure Cluster Shared Volumes

1. In the Failover Cluster Manager console on LON-SVR3, navigate to **Storage->Disks**.
2. Locate a disk that is assigned to **Available Storage**. (If possible use Cluster Disk 2).
3. Add this to **Cluster Shared Volumes**.

Results: After this exercise, you will have installed and configured the failover clustering feature.

Exercise 2: Deploying and Configuring a Highly Available File Server**Scenario**

In A. Datum Corporation, File Services is one of the important services that must be highly available, because it hosts very important data that is being used all the time. After you have created a cluster infrastructure, you decide to configure a highly available file server and implement settings for failover and failback.

The main tasks for this exercise are as follows:

1. Add the File Server application to the failover cluster.
2. Add a shared folder to a highly available file server.
3. Configure failover and failback settings.
4. Validate cluster quorum settings.

► Task 1: Add the File Server application to the failover cluster

1. Add the **File Server** role service to LON-SVR3 and LON-SVR4.
2. On LON-SVR3, open the Failover Cluster Manager console.
3. Add **File Server** as a cluster role.
4. Choose to implement **Scale-Out File Server for application data**.
5. Specify **AdatumFS** as **Client Access Name**.

► Task 2: Add a shared folder to a highly available file server

1. On LON-SVR3, in the Failover Cluster Manager, start a New Share Wizard to add a new shared folder to the AdatumFS cluster role.
2. Specify the profile for the share as **SMB Share – Quick**.
3. Name the shared folder as **Data**.
4. Enable **continuous availability**.

► Task 3: Configure failover and failback settings

1. On LON-SVR3, in the Failover Cluster Manager, open the **Properties** for the **AdatumFS** cluster role.
2. Enable failback between **4** and **5** hours.
3. Select both **LON-SVR3** and **LON-SVR4** as the preferred owners.
4. Move **LON-SVR4** to be first in the **Preferred Owners** list.

► **Task 4: Validate cluster quorum settings**

- In the Failover Cluster Manager console, review settings for **Quorum Configuration**. It should be set to **Node and Disk Majority**.

Results: After this exercise, you will have deployed and configured a highly available file server.

Exercise 3: Validating the Deployment of the Highly Available File Server

Scenario

In the process of implementing a failover cluster, you want to ensure that cluster s performing correctly, by performing failover and failback tests.

The main tasks for this exercise are as follows:

1. Validate the highly available file server deployment.
2. Validate the failover and quorum configuration for the file server role.

► **Task 1: Validate the highly available file server deployment**

1. On LON-DC1, open Windows Explorer, and attempt to access the **\\AdatumFS** location. Make sure that you can access the **Data** folder.
2. Create a test text document inside this folder.
3. On LON-SVR3, in the Failover Cluster Manager, move **AdatumFS** to the second node.
4. On LON-DC1, in Windows Explorer, verify that you can still access **\\AdatumFS** location.

► **Task 2: Validate the failover and quorum configuration for the file server role**

1. On LON-SVR3, determine the current owner for the **AdatumFS** role.
2. Stop the Cluster service on the node that is the current owner of the **AdatumFS** role.
3. Verify that **AdatumFS** has moved to another node, and that the **\\AdatumFS** location is still available from the LON-DC1 computer.
4. Start the Cluster service on the node in which you stopped it in step 2.
5. From the Disks node, take the disk witness offline.
6. Verify that the **\\AdatumFS** location is still available from LON-DC1.
7. Bring the disk witness back online.

Results: After this exercise, you will have tested the failover and failback scenarios.

Exercise 4: Configuring Cluster-Aware Updating on the Failover Cluster

Scenario

Earlier, implementing updates to servers with critical service was causing unwanted downtime. To enable seamless and zero-downtime cluster updating, you want to implement the Cluster-Aware Updating feature and test updates for cluster nodes.

The main tasks for this exercise are as follows:

1. Configure Cluster-Aware Updating.

2. Update the failover cluster and configure self-updating.

► **Task 1: Configure Cluster-Aware Updating**

1. On LON-DC1, install the Failover Clustering feature.
2. From Server Manager, open **Cluster-Aware Updating**.
3. Connect to **Cluster1**.
4. Preview the updates available for nodes in **Cluster1**.

► **Task 2: Update the failover cluster and configure self-updating**

1. On LON-DC1, start the update process for **Cluster1**.
2. After the process is complete, log on to LON-SVR3 with the username as **Adatum\Administrator** and password as **Pa\$\$w0rd**.
3. On LON-SVR3, open **Cluster-Aware Updating** and configure **self-updating for Cluster1**, to be performed **weekly**, on **Sundays** at **4:00A.M.**

Results: After this exercise, you will have configured Cluster-Aware Updating on the Failover Cluster.

► **To prepare for the next module**

When you finish the lab, revert the virtual machines to their initial state.

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20412A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20412A-LON-SVR1**, **20412A-LON-SVR3**, **20412A-LON-SVR4** and **MSL-TMG1**.

Lab Review

Question: What information will you have to collect as you plan a failover cluster implementation and choose a quorum mode?

Question: After running the Validate a Configuration Wizard, how can you resolve the network communication single point of failure?

Question: In which situations might it be important to enable failback of a clustered application only during a specific time?

Module Review and Takeaways

Question: Why is using a No Majority: Disk-Only quorum configuration generally not a good idea?

Question: What is the purpose of CAU?

Question: What is the main difference between synchronous and asynchronous replication in a multi-site cluster scenario?

Question: What is the multi-site clusters enhanced feature in Windows Server 2012?

Real-world Issues and Scenarios

Question: Your organization is considering the use of a geographically dispersed cluster that includes an alternative data center. Your organization has only a single physical location together with an alternative data center. Can you provide an automatic failover in this configuration?

Tools

The tools for implementing failover clustering include:

- Failover Cluster Manager console
- Cluster-Aware Updating console
- Windows PowerShell
- Server Manager
- iSCSI initiator
- Disk Management

MCT USE ONLY. STUDENT USE PROHIBITED

Module 6

Implementing Failover Clustering with Hyper-V

Contents:

Module Overview	6-1
Lesson 1: Overview of Integrating Hyper-V with Failover Clustering	6-2
Lesson 2: Implementing Hyper-V Virtual Machines on Failover Clusters	6-7
Lesson 3: Implementing Hyper-V Virtual Machine Movement	6-15
Lesson 4: Managing Hyper-V Virtual Environments by Using VMM	6-21
Lab: Implementing Failover Clustering with Hyper-V	6-31
Module Review and Takeaways	6-36

Module Overview

One benefit of implementing server virtualization is that it allows you to provide high availability, both for applications or services that have built-in high availability functionality, and for applications or services that do not provide high availability in any other way. With the Windows Server® 2012 Hyper-V® technology, failover clustering, and Microsoft® System Center 2012 - Virtual Machine Manager (VMM), you can configure high availability by using several different options.

In this module, you will learn about how to implement failover clustering in a Hyper-V scenario to achieve high availability for a virtual environment. You will also learn about basic virtual machine features.

Objectives

After completing this module, you will be able to:

- Describe how Hyper-V integrates with failover clustering.
- Implement Hyper-V virtual machines on failover clusters.
- Implement Hyper-V virtual machine movement.
- Manage a Hyper-V virtual environment by using VMM.

Lesson 1

Overview of Integrating Hyper-V with Failover Clustering

Failover clustering is a Windows Server 2012 feature that enables you to make applications or services highly available. To make virtual machines highly available in a Hyper-V environment, you must implement failover clustering on Hyper-V host machines.

This lesson summarizes the high availability options for Hyper-V–based virtual machines, and then focuses on how failover clustering works, and how to design and implement failover clustering for Hyper-V.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe options for making virtual machines highly available.
- Explain how failover clustering works with Hyper-V nodes.
- Describe new failover clustering features for Hyper-V in Windows Server 2012.
- Describe best practices for implementing high availability in a virtual environment.

Options for Making Virtual Machines Highly Available

Most organizations have some applications that are business critical and must be highly available. To make an application highly available, you must deploy it in an environment that provides redundancy for all components that the application requires. For virtual machines to be highly available, you can choose between several options:

- Host clustering, in which you implement virtualization hosts as a clustered role
- Guest clustering, in which you implement clustering inside virtual machines
- Network Load Balancing (NLB) inside virtual machines

High availability options	Description
Host clustering	<ul style="list-style-type: none"> • Virtual machines are highly available • Does not require virtual machine operating system or application to be cluster-aware
Guest clustering	<ul style="list-style-type: none"> • Virtual machines are failover cluster nodes • Virtual machine applications must be cluster-aware • Requires iSCSI or virtual Fibre Channel interface for shared storage connections
NLB	<ul style="list-style-type: none"> • Virtual machines are NLB cluster nodes • Use for web-based applications

Host Clustering

Host clustering enables you to configure a failover cluster by using the Hyper-V host servers. When you configure host clustering for Hyper-V, you configure the virtual machine as a highly available resource. Failover protection is implemented at the host server level. This means that the guest operating system and applications that are running within the virtual machine do not have to be cluster-aware. However, the virtual machine is still highly available. Some examples of non-cluster-aware applications are print server, or proprietary network-based applications such as an accounting application. Should the host node that controls the virtual machine unexpectedly become unavailable, the secondary host node assumes control and restarts the virtual machine as quickly as possible.

You can also move the virtual machine from one node in the cluster to another in a controlled manner. For example, you could move the virtual machine from one node to another while patching the host operating system. The applications or services that are running in the virtual machine do not have to be compatible with failover clustering, and they do not need to be aware that the virtual machine is

clustered. Because the failover is at the virtual machine level, there are no dependencies on software that is installed inside the virtual machine.

Guest Clustering

You configure guest failover clustering very similarly to physical server failover clustering, except that the cluster nodes are multiple virtual machines. In this scenario, you create two or more virtual machines, and enable failover clustering within the guest operating system. The application or service is then enabled for high availability between the virtual machines by using failover clustering in each virtual machine. Because you implement failover clustering within each virtual machine node's guest operating system, you can locate the virtual machines on a single host. This can be a quick and cost-effective configuration in a test or staging environment.

For production environments, however, you can more robustly protect the application or service if you deploy the virtual machines on separate failover clustering-enabled Hyper-V host computers. With failover clustering implemented at both the host and virtual machine levels, you can restart the resource regardless of whether the node that fails is a virtual machine or a host. This configuration is also known as a *Guest Cluster Across Hosts*. It is considered an optimal high availability configuration for virtual machines that are running critical applications in a production environment.

You should consider several factors when you implement guest clustering:

- The application or service must be failover cluster-aware. This includes any of the Windows Server 2012 services that are cluster-aware, and any applications, such as clustered Microsoft SQL Server® and Microsoft Exchange Server.
- Hyper-V virtual machines can use Fibre Channel-based connections to shared storage (this is specific only to Microsoft Hyper-V Server 2012), or you can implement internet small computer system interface (iSCSI) connections from the virtual machines to the shared storage.

You should deploy multiple network adapters on the host computers and the virtual machines. Ideally, when using an iSCSI connection you should dedicate a network connection to the iSCSI connection, to the private network between the hosts, and to the network connection used by the client computers.

Network Load Balancing

NLB works with virtual machines in the same manner with which it works with physical hosts. It distributes IP traffic to multiple instances of a TCP/IP service, such as a web server that is running on a host within the NLB cluster. NLB transparently distributes client requests among the hosts, and it enables the clients to access the cluster by using a virtual host name or a virtual IP addresses. From the client computer's point of view, the cluster seems to be a single server that answers these client requests. As enterprise traffic increases, you can add another server to the cluster.

Therefore, NLB is an appropriate solution for resources that do not have to accommodate exclusive read or write requests. Examples of NLB-appropriate applications are web-based front ends to database applications, or Exchange Server Client Access servers.

When you configure an NLB cluster, you must install and configure the application on all virtual machines. After you configure the application, you install the NLB feature in Windows Server 2012 within each virtual machine's guest operating system (not on the Hyper-V hosts), and then configure an NLB cluster for the application. Earlier versions of Windows Server also support NLB, which means that the guest operating system is not limited to only Windows Server 2012. Similar to a Guest Cluster Across Host, the NLB resource typically benefits from overall increased input/output (I/O) performance when the virtual machine nodes are located on different Hyper-V hosts.

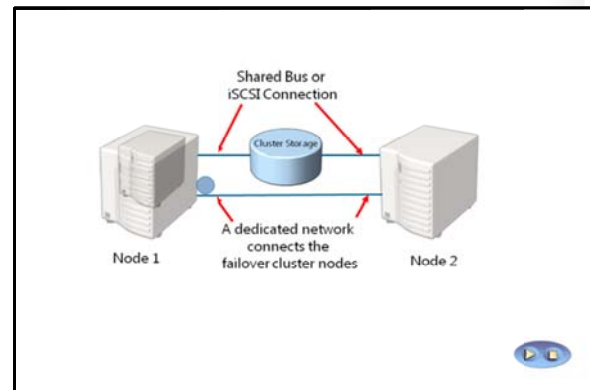


Note: As with earlier versions of Windows Server, you should not implement NLB and failover clustering within the same guest operating system because the two technologies conflict with one another.

How Does a Failover Cluster Work with Hyper-V Nodes?

When you implement failover clustering and configure virtual machines as highly available resources, the failover cluster treats the virtual machines like any other application or service. Namely, if a host fails, failover clustering will act to restore access to the virtual machine as quickly as possible on another host within the cluster. Only one node at a time runs the virtual machine. However, you can also move the virtual machine to any other node within the same cluster.

The failover process transfers the responsibility of providing access to resources within a cluster from one node to another. Failover can occur when an administrator moves resources to another node for maintenance or other reasons, or when unplanned downtime of one node occurs because of hardware failure or for other reasons.



The failover process consists of the following steps:

1. The node where the virtual machine is running owns the clustered instance of the virtual machine, controls access to the shared bus or iSCSI connection to the cluster storage, and has ownership of any disks, or logical unit numbers (LUNs), assigned to the virtual machine. All the nodes in the cluster use a private network to send regular signals, known as *heartbeat signals*, to one another. The heartbeat signals that a node is functioning and communicating on the network. The default heartbeat configuration specifies that each node send a heartbeat over TCP/UDP port 3343 each second (or 1,000 milliseconds).
2. Failover starts when the node hosting the virtual machine does not send regular heartbeat signals over the network to the other nodes. By default, this corresponds to five consecutively missed heartbeats (or 5,000 milliseconds). Failover may occur because of a node failure or network failure.
3. When heartbeat signals stop arriving from the failed node, one of the other nodes in the cluster begins taking over the resources that the virtual machines use. You define the node(s) that could take over by configuring the **Preferred and Possible Owners** properties. The preferred owner specifies the hierarchy of ownership if there is more than one possible failover node for a resource.

By default, all nodes are possible owners. Therefore, removing a node as a possible owner absolutely excludes it from taking over the resource in a failure situation. For example, suppose that you implement a failover cluster by using three nodes. However, only two nodes are configured as preferred owners. During a failover event, the resource could still be taken over by the third node if neither of the preferred owners are online. Although the third node is not configured as a preferred owner, as long as it is a possible owner, the failover cluster can use it if necessary to restore access to the resource.

Resources are brought online in order of dependency. For example, if the virtual machine references an iSCSI LUN, access to the appropriate host bus adapters (HBAs), network(s), and LUNs will be stored in that order. Failover is complete when all the resources are online on the new node. For clients interacting with the resource, there is a short service interruption, which most users will not notice.

4. You can also configure the cluster service to fail back to the offline node after it again becomes active. When the cluster service fails back, it uses the same procedures that it performs during failover. This means that the cluster service takes offline all the resources associated with that instance, moves the instance, and then brings all the resources in the instance back online.

What Is New in Failover Clustering for Hyper-V in Windows Server 2012

In Windows Server 2012, failover clustering is much improved with respect to Hyper-V clusters. Some of the most important improvements are:

- Failover clustering now supports up to 4,000 virtual machines, and the improved Failover Cluster Manager snap-in simplifies managing many virtual machines.
- Administrators can now perform multiselect actions to queue live migrations of multiple virtual machines, instead of one by one as in earlier versions.
- Administrators can also configure the virtual machine priority attribute to control the order in which virtual machines are started. Priority is also used to ensure that lower-priority virtual machines automatically release resources to higher priority virtual machines as needed.
- The Cluster Shared Volume (CSV) feature, which simplifies the configuration and operation of virtual machines, can help improve security and performance. It now supports scalable file-based server application storage, increased backup and restore and single consistent file namespace. In addition, you can now protect CSV volumes by using Windows® BitLocker® Drive Encryption, and by configuring the CSV volumes to make storage visible to only a subset of nodes.
- Virtual machine application monitoring. You can now monitor services that are running on clustered virtual machines. In clusters running Windows Server 2012, administrators can configure monitoring of services on clustered virtual machines that are also running Windows Server 2012. This functionality extends the high-level monitoring of virtual machines that is implemented in Windows Server 2008 R2 failover clusters.
- You can now store virtual machines on server message block (SMB) file shares in a file server cluster. This is a new method for providing highly available virtual machines. Instead of configuring a cluster between Hyper-V nodes, you can now have Hyper-V nodes out of cluster, but with virtual machine files on a highly available file share. To make this work, you should deploy a file server cluster in a scale-out file server mode. Scale-out file servers can also use CSV for storage.

Failover clustering for Hyper-V now includes:

- Support for up to 4,000 virtual machines per cluster
- Multi-select virtual machines for live migration
- Virtual machine priority attribute
- CSV improvements
- Virtual machine application monitoring
- Storing virtual machines on highly available SMB file shares

Best Practices for Implementing High Availability in a Virtual Environment

After you determine which applications you want to deploy on highly available failover clusters, you can plan and deploy the failover clustering environment. Consider the following recommendations when you implement the failover cluster:

- Use Windows Server 2012 as the Hyper-V host. Windows Server 2012 provides enhancements such as Hyper-V 3.0, improved CSVs, virtual machine migrations, and other features that improve flexibility and performance when you implement host failover clustering.
- Plan for failover scenarios. When you design the hardware requirements for the Hyper-V hosts, make sure that you include the hardware capacity that is required when hosts fail. For example, if you deploy a six-node cluster, you must determine the number of host failures that you want to accommodate. If you decide that the cluster must sustain the failure of two nodes, then the four remaining nodes must have the capacity to run all of the virtual machines in the cluster.
- Plan the network design for failover clustering. To optimize the failover cluster performance and failover, you should dedicate a fast network connection for internode communication. As with earlier versions, this network should be logically and physically separate from the network segment(s) used by clients to communicate with the cluster. You can also use this network connection to transfer virtual machine memory during a live migration. If you are using iSCSI for any virtual machines, also dedicate a network connection to the iSCSI network connection.
- Plan the shared storage for failover clustering. When you implement failover clustering for Hyper-V, the shared storage must be highly available. If the shared storage fails, the virtual machines will all fail, even if the physical nodes are functional. To ensure storage availability, plan for redundant connections to the shared storage and redundant array of independent disks (RAID) redundancy on the storage device.
- Use the recommended failover cluster quorum mode. If you deploy a cluster with an even number of nodes, and shared storage is available to the cluster, the Failover Cluster Manager automatically selects the Node and Disk Majority quorum mode. If you deploy a cluster with an odd number of nodes, the Failover Cluster Manager automatically selects the Node Majority quorum mode. You should not modify the default configuration unless you understand the implications of doing this.
- Deploy standardized Hyper-V hosts. To simplify the deployment and management of the failover cluster and Hyper-V nodes, develop a standard server hardware and software platform for all nodes.
- Develop standard management practices. When you deploy multiple virtual machines in a failover cluster, you increase the risk that a single mistake may shut down a large part of the server deployment. For example, if an administrator accidentally configures the failover cluster incorrectly, and the cluster fails, all virtual machines in the cluster will be offline. To avoid this, develop and thoroughly test standardized instructions for all administrative tasks.

Best practices for implementing high availability in a virtual environment include:

- Use Windows Server 2012 as the Hyper-V host
- Plan for failover scenarios
- Plan the network design for failover clustering
- Plan the shared storage for failover clustering
- Use the recommended failover cluster quorum mode
- Deploy standardized Hyper-V hosts
- Develop and test standard management practices

Lesson 2

Implementing Hyper-V Virtual Machines on Failover Clusters

Implementing highly available virtual machines is somewhat different from implementing other roles in a failover cluster. Failover clustering in Windows Server 2012 provides many features for Hyper-V clustering, in addition to tools for virtual machine high availability management. In this lesson, you will learn how to implement highly available virtual machines.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe components of a Hyper-V cluster.
- Describe prerequisites for implementing Hyper-V failover clusters.
- Implement failover clustering for Hyper-V virtual machines.
- Configure CSVs.
- Implement highly available virtual machines on SMB 3.0 file shares.
- Describe considerations for implementing Hyper-V virtual machines in a cluster.

Components of Hyper-V Clusters

Hyper-V as a role has some specific requirements for cluster components. To form a Hyper-V cluster, you must have at least two physical nodes. Whereas other clustered roles (such as Dynamic Host Configuration Protocol (DHCP) file server) allow for nodes to be virtual machines, Hyper-V nodes must be composed of physical hosts. You cannot run Hyper-V inside a virtual machine on a Hyper-V host.

In addition to having nodes, you must also have physical and virtual networks. Failover clustering requires a network for internal cluster

communication, and also a network for clients. You can also implement a storage network separately, depending on the type of storage that you are using. Again, specific to the Hyper-V role, you should also consider virtual networks for clustered virtual machines. It is important that you create the same virtual networks on all physical hosts that participate in one cluster. Failure to do this will cause a virtual machine to lose network connectivity when moved from one host to another.

Storage is an important component of virtual machine clustering. You can use any type of storage that is supported by Windows Server 2012 failover clustering. As a best practice, you should configure storage as a CSV. This is discussed further in a following topic within this module.

Virtual machines are components of a Hyper-V cluster. In the Failover Cluster Manager, you can create new highly available virtual machines, or you can make existing virtual machines highly available. In both cases, the virtual machine storage location must be on shared storage that all nodes can access. However, you might not want to make all virtual machines highly available. In Failover Cluster Manager, you can select which virtual machines that you want to be part of a cluster configuration.

Hyper-V cluster components include:

- Cluster nodes that are physical computers
- Cluster networks
- Virtual networks
- Storage for virtual machines
- Virtual machines

Prerequisites for Implementing Hyper-V Clusters


To deploy Hyper-V on a failover cluster, you must ensure that you meet the hardware, software, account, and network infrastructure requirements as detailed in the following sections.

Hardware Requirements for Failover Clustering with Hyper-V

You must have the following hardware for a two-node failover cluster:

- Server hardware. Hyper-V requires an x64-based processor, hardware-assisted virtualization, and hardware-enforced Data Execution Prevention (DEP). As a best practice, the servers should have similar hardware. If you are using Windows Server 2008, the processors on the servers must be the same version. If you are using Windows Server 2008 R2 or Windows Server 2012, the processors must use the same architecture.

Hardware requirements for cluster nodes and storage	<ul style="list-style-type: none"> • Server hardware • Network adapters • Storage adapters • Storage
Software requirements for cluster nodes	<ul style="list-style-type: none"> • x64-based version of Windows Server 2012, Enterprise or Datacenter edition • Same software update and service packs • Full installation or Server Core installation
Network infrastructure requirements	<ul style="list-style-type: none"> • Network settings and IP addresses • DNS • Domain role • Account for administering the cluster

 **Note:** Microsoft supports a failover cluster solution only if all the hardware features are marked as “Certified for Windows Server.” Additionally, the complete configuration (servers, network, and storage) must pass all tests in the Validate This Configuration wizard, which is included in the Failover Cluster Manager snap-in.

- Network adapters. As with the other features in the failover cluster solution, the network hardware must be marked as “Certified for Windows Server.” To provide network redundancy, you can connect cluster nodes to multiple, distinct networks, or you can connect the nodes to one network that uses teamed network adapters, redundant switches, redundant routers, or similar hardware to remove single points of failure. As a best practice, you should configure multiple network adapters on the host computer that you configure as a cluster node. You should connect one network adapter to the private network that the inter-host communications uses.
- Storage adapters. If you use a Serial Attached SCSI (SAS) or Fibre Channel, the mass-storage device controllers in all clustered servers should be identical and should use the same firmware version. If you are using iSCSI, each clustered server should have one or more network adapters that are dedicated to the cluster storage. The network adapters that you use to connect to the iSCSI storage target should be identical, and you should use a gigabit Ethernet or faster network adapter.
- Storage. You must use shared storage that is compatible with Windows Server 2012. If you deploy a failover cluster that uses a witness disk, the storage must contain at least two separate volumes (or LUNs). One volume functions as the witness disk, and additional volumes contain the virtual machine files that are shared between the cluster nodes. Storage considerations and recommendations include the following:
 - Use basic disks, not dynamic disks. Format the disks with the NTFS file system.
 - Use either master boot record (MBR) or GUID partition table (GPT).
 - If you are using a storage area network (SAN), the miniport driver that the storage uses must work with the Microsoft Storport storage driver.

- Consider using multipath I/O software: If your SAN uses a highly available network design with redundant components, you can deploy failover clusters with multiple host bus adapters by using multipath I/O software. This provides the highest level of redundancy and availability. For Windows Server 2008 R2 and Windows Server 2012, your multipath solution must be based on Multipath I/O (MPIO).

Software Requirements for Using Hyper-V and Failover Clustering

The following are the software requirements for using Hyper-V and failover clustering:

- All the servers in a failover cluster must run the x64-based version of Windows Server 2012 Standard Edition or Windows Server 2012 Datacenter Edition. The nodes in a single failover cluster cannot run different versions.
- All the servers should have the same software updates and service packs.
- All the servers must be installed as a Full installations or Server Core installations.

Network Infrastructure Requirements

The following network infrastructure is required for a failover cluster and an administrative account with the following domain permissions:

- Network settings and IP addresses. Use identical communication settings on all network adapters, including the speed, duplex mode, flow control, and media type settings. Ensure that all network hardware supports the same settings.
- If you use private networks that are not routed to your whole network infrastructure for communication between cluster nodes, ensure that each of these private networks uses a unique subnet.
- DNS. The servers in the cluster must use Domain Name System (DNS) for name resolution. You should use the DNS dynamic update protocol.
- Domain role. All servers in the cluster must be in the same Active Directory® Domain Services (AD DS) domain. As a best practice, all clustered servers should have the same domain role, either member server or domain controller. The recommended role is member server. You should avoid installing cluster nodes on domain controllers because AD DS has its own high availability mechanism.
- Account for administering the cluster. When you first create a cluster or add servers to a cluster, you must be logged on to the domain with an administrator's account on all the cluster's servers. Additionally, if the account is not a Domain Admins account, the account must have the **Create Computer Objects** permission in the domain.

Implementing Failover Clustering for Hyper-V Virtual Machines

To implement failover clustering for Hyper-V virtual machines, you must complete the following high-level steps:

1. Install and configure the required versions of Windows Server 2012. After you complete the installation, configure the network settings, join the computers to an Active Directory domain, and configure the connection to the shared storage.

To implement failover clustering for Hyper-V virtual machines:

1. Install and configure Windows Server 2012
2. Configure shared storage
3. Install the Hyper-V and failover clustering features
4. Validate the cluster configuration
5. Create the cluster
6. Create a virtual machine on one of the cluster nodes
7. Make the virtual machine highly available

2. Configure the shared storage. You must use Disk Manager to create disk partitions on the shared storage.
3. Install the Hyper-V and failover clustering features on the host servers. You can use Server Manager in the Microsoft Management Console (MMC) or Windows PowerShell® for this.
4. Validate the cluster configuration. The Validate This Cluster wizard checks all the prerequisite components that are required to create a cluster, and provides warnings or errors if any components do not meet the cluster requirements. Before you continue, resolve any issues that the Validate This Cluster wizard identifies.
5. Create the cluster. Once the components pass the Validate This Cluster wizard, you can create a cluster. When you configure the cluster, assign a cluster name and an IP address. A computer account for the cluster name is created in the Active Directory domain, and the IP address is registered in DNS.



Note: You can enable clustered shared storage for the cluster only after you configure the cluster. If you want to use CSV, you should configure CSV before you proceed to the next step.

6. Create a virtual machine on one of the cluster nodes. When you create the virtual machine, ensure that all files that are associated with the virtual machine—including both the virtual hard disk and virtual machine configuration files—are stored on the shared storage. You can create and manage virtual machines in either Hyper-V Manager or Failover Cluster Manager. When you create a virtual machine by using Failover Cluster Manager, the virtual machine is made highly available automatically.
7. Make the virtual machine highly available. To make the virtual machine highly available, in the Failover Cluster Manager, select the option to make a new service or application highly available. The Failover Cluster Manager then displays a list of services and applications that you can make highly available. When you select the option to make virtual machines highly available, you can select the virtual machine that you created on shared storage.



Note: When you make a virtual machine highly available, a list displays of all virtual machines that are hosted on all cluster nodes—including virtual machines that are not stored on the shared storage. If you make a virtual machine that is not located on shared storage highly available, you receive a warning, but Hyper-V will add the virtual machine to the services and applications list. However, when you try to migrate the virtual machine to a different host, the migration will fail.

8. Test virtual machine failover. After you make the virtual machine highly available, you can migrate the computer to another node in the cluster. If you are running Windows Server 2008 R2 or Windows Server 2012, you can select to perform a quick migration or a live migration.

Configuring CSVs

You do not have to configure and use CSV when you implement high availability for virtual machines in Hyper-V. In fact, you can configure a Hyper-V cluster without using CSV. However as a best practice, you use CSV due to the following advantages:

- Reduced LUNs for the disks. You can use CSV to reduce the number of LUNs that your virtual machines require. When you configure a CSV, you can store multiple virtual machines on a single LUN, and multiple host computers can access the same LUN concurrently.
- Better use of disk space. Instead of placing each virtual hard disk (.vhdx) file on a separate disk with empty space so that the .vhdx file can expand, you can oversubscribe disk space by storing multiple .vhdx files on the same LUN.
- Virtual machine files stored in a single logical location. You can track the paths of .vhdx files and other files that virtual machines use. Instead of using drive letters or GUIDs to identify disks, you can specify the path names. When you implement CSV, all added storage appears in the \ClusterStorage folder. The \ClusterStorage folder is created on the cluster node's system folder, and you cannot move it. This means that all Hyper-V hosts that are members of the cluster must use the same drive letter as their system drive, or virtual machine failovers will fail.
- No specific hardware requirements. CSV implementation does not have specific hardware requirements. You can implement CSV on any supported disk configuration, and on either the Fibre Channel or iSCSI SANs.
- Increased resiliency. CSV increases resiliency because the cluster can respond correctly even if connectivity between one node and the SAN is interrupted, or part of a network is down. The cluster reroutes the traffic to the CSV through an intact part of the SAN or network.

CSV benefits are:

- Fewer LUNs required
- Better use of disk space
- Virtual machine files stored in a single logical location
- No special hardware required
- Increased resiliency

To implement CSV, do the following:

1. Create and format volumes on shared storage
2. Add the disks to failover cluster storage
3. Add the storage to the CSV

Implementing CSV

After you create the failover cluster, you can enable CSV for the cluster, and then add storage to the CSV.

Before you can add storage to the CSV, the LUN must be available as shared storage for the cluster. When you create a failover cluster, all the shared disks that are configured in Server Manager are added to the cluster, and you can then add them to a CSV. If you add more LUNs to the shared storage, you must first create volumes on the LUN, add the storage to the cluster, and then add the storage to the CSV.

As a best practice, you should configure CSV before you make any virtual machines highly available. However, you can convert from regular disk access to CSV after deployment. When implementing CSV, the following considerations apply:

- The LUN's drive letter or mount point is removed when you convert from regular disk access to CSV. This means that you must re-create all virtual machines that are stored on the shared storage. If you must keep the same virtual machine settings, consider exporting the virtual machines, switching to CSV, and then importing the virtual machines in Hyper-V. Additionally, consider using the storage migration option that is available in the Hyper-V role in Windows Server 2012.
- You cannot convert shared storage to CSV. If you have any single running virtual machine that uses a cluster disk, you must shut down the virtual machine, and then add the disk to CSV.

Implementing Highly Available Virtual Machines on an SMB 3.0 File Share

In Windows Server 2012, you can use a new technique to make virtual machines highly available. Instead of using host or guest clustering, you can now store virtual machine files on a highly available SMB 3.0 file share. By using this approach, high availability is achieved not by clustering Hyper-V nodes, but by file servers that host virtual machine files on their file shares. With this new capability, Hyper-V can store all virtual machine files, including configuration, .vhdx files, and snapshots, on highly available SMB 3.0 file shares.

In Windows Server 2012 you can store virtual machine files on an SMB 3.0 file share

Requirements include:

- File servers running Windows Server 2012
- File server cluster running in scale-out mode

You can use Hyper-V Manager to create or move virtual machine files to the SMB 3.0 file share

This technology requires the following infrastructure:

- One or more computers that are running Windows Server 2012 with the Hyper-V role installed.
- One or more computers that are running Windows Server 2012 with the File and Storage Services role installed.
- Domain members in the Active Directory infrastructure. The servers running AD DS do not need to run Windows Server 2012.

Before you implement virtual machines on an SMB 3.0 file share, configure a file server cluster. To do this, you should have at least two cluster nodes, both with file services and failover clustering installed on them. In the Failover Clustering console, create a scale-out file server cluster. After you configure the cluster, deploy the new SMB file share for applications. This share stores virtual machine files. When the share is created, you can use the Hyper-V Manager to deploy new virtual machines on the SMB 3.0 file share, or you can migrate existing virtual machines to the SMB file share by using the storage migration method.

Demonstration: Implementing Virtual Machines on Clusters (optional)

In this demonstration, you will see how to implement virtual machines on a failover cluster.



Note: Before starting this demonstration, ensure that LON-HOST1 is the owner of the ClusterVMs disk. If it is not, then move the ClusterVMs resource to LON-HOST1 before doing this procedure.

Demonstration Steps

Move virtual machine storage to the iSCSI target

- On LON-HOST1, open Windows Explorer, browse to **E:\Program Files\Microsoft Learning\20412\20412A-LON-CORE\Virtual Hard Disks**, and then move **20412A-LON-CORE.vhd** to the **C:\ClusterStorage\Volume1** location.

Configure the machine as highly available

1. In Failover Cluster Manager, click **Roles**, and then start the New Virtual Machine Wizard.
2. In the New Virtual Machine Wizard, use the following settings:
 - Cluster node: **LON-HOST2**.

- Computer name: **TestClusterVM**
 - Store the file at **C:\ClusterStorage\Volume1**.
 - RAM for TestClusterVM: **1536 MB**
 - Connect machine to existing virtual hard disk drive 20412A-LON-CORE.vhd located at **C:\ClusterStorage\Volume1**.
3. From the **Roles** node, start the virtual machine.

Considerations for Implementing Hyper-V Clusters

By implementing host failover clustering, you can make virtual machines highly available. However, implementing host failover clustering also adds significant cost and complexity to a Hyper-V deployment. You must invest in additional server hardware to provide redundancy, and you should implement or have access to a shared storage infrastructure.

Consider the following recommendations to ensure that the failover clustering strategy meets the organization's requirements:

Recommendations for implementing Hyper-V clusters are:

1. Identify the applications that require high availability
2. Identify the application components that must be highly available
3. Identify the application characteristics
4. Identify the total capacity requirements
5. Create the Hyper-V design
 - Verify basic requirements
 - Configure a dedicated network adapter for the private virtual network
 - Use similar host hardware
 - Verify network configuration
 - Manage live migrations

- Identify the applications or services that require high availability. Unless you have the option of making all virtual machines highly available, you must develop priorities for which applications you will make highly available.
- Identify the components that must be highly available to make the applications highly available. In some cases, the application might run on a single server, and making that server highly available is all that you have to do. Other applications may require that several servers and components such as storage or the network, be made highly available. In addition, ensure that the domain controllers are highly available, and that you have at least one domain controller on separate hardware or virtualization infrastructure.
- Identify the application characteristics. You must understand several aspects about the application.
 - Is virtualizing the server that is running the application an option? Some applications are not supported in or recommended for a virtual environment.
 - What options are available for making the application highly available? You can make some applications highly available through options other than host clustering. If other options are available, evaluate the benefits and disadvantages of each option.
 - What are the performance requirements for each application? Collect performance information on the servers currently running the applications to gain an understanding of the hardware requirements that are required when you virtualize the server.
 - What capacity is required to make the Hyper-V virtual machines highly available? As soon as you identify all the applications that you must make highly available by using host clustering, you can start to design the actual Hyper-V deployment. By identifying the performance requirements and network and storage requirements for applications, you can define the hardware that you have to implement for all the applications in a highly available environment.

Live migration is one of the most important aspects of Hyper-V clustering. You use the Live Migration feature in Windows Server 2012 to perform live migrations of virtual machines. When implementing live migration, consider the following:

- Verify basic requirements. Live migration requires that all hosts be part of a Windows Server 2012 failover cluster, and that the host processors have the same architecture. All hosts in the cluster must have access to shared storage, which meets the requirements for CSV.
- Configure a dedicated network adapter for the private virtual network. When you implement failover clustering, you should configure a private network for the cluster heartbeat traffic. You use this network to transfer the virtual machine memory during a failover. To optimize this configuration, configure for this network a network adapter that has a capacity of one gigabit per second (Gbps) or higher.



Note: You must enable the Client for Microsoft Networks component, and the File and Printer Sharing for Microsoft Networks component, for the network adapter that you want to use for the private network.

- Use similar host hardware. As a best practice, all failover cluster nodes should use the same hardware for connecting to shared storage, and all cluster nodes must have processors that have the same architecture. Whereas you can enable failover for virtual machines on a host with different processor versions by configuring processor compatibility settings, the failover experience and performance is more consistent if all servers have similar hardware.
- Verify network configuration. All nodes in the failover cluster must connect through the same IP subnet so that the virtual machine can continue communicating through the same IP address after live migration. In addition, the IP addresses that are assigned to the private network on all nodes must be on the same logical subnet. This means that multisite clusters must use a stretched virtual local area network (VLAN), which is a subnet that spans a wide area network (WAN) connection.
- Manage live migrations. In Windows Server 2008 R2, each node in the failover cluster can perform only one live migration at a time. If you try to start a second live migration before the first migration finishes, the migration fails. In Windows Server 2012, you can now run multiple live migrations simultaneously.

Lesson 3

Implementing Hyper-V Virtual Machine Movement

Moving virtual machines from one location to another is a common procedure in Hyper-V environments. While moving virtual machines in previous Windows Server versions required downtime, Windows Server 2012 introduces new technologies to enable seamless virtual machine movement. In this lesson, you will learn about virtual machine movement and migration options.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe migration options for virtual machines.
- Describe storage migration.
- Describe live migration.
- Explain how Hyper-V replicas work.
- Configure a Hyper-V replica.

Virtual Machine Migration Options

There are several scenarios in which you would want to migrate virtual machines from one location to another. For example, you might want to move a virtual machine virtual hard disk from one physical drive to another on the same host. Alternatively, you may need to move a virtual machine from one node in a cluster to another, or simply move a computer from one host server to another host server without the hosts being members of a cluster. Compared with Windows Server 2008 R2, Windows Server 2012 provides significant enhancements for this process in addition to simplified procedures.

To move virtual machines, you can use:

- Virtual machine and storage migration
- Quick migration
- Live migration
- Hyper-V Replica

You can also move a virtual machine by exporting and then importing the virtual machine

In Windows Server 2012, you can perform migration of virtual machines by using the following methods:

- **Virtual machine and storage migration.** With this method, you move a powered-on virtual machine from one location to another (or from one host to another) by using a wizard in Hyper-V Manager. Virtual machine and storage migration do not require failover clustering or any other high availability technology. Additionally, you do not need shared storage when you move just the virtual machine.
- **Quick Migration.** This method is also available in Windows Server 2008. It requires you have failover clustering installed and configured. The quick migration process saves the state of the virtual machine before the failover, and then restarts the virtual machine after failover completes.
- **Live Migration.** This feature is an improvement over Quick Migration, and is also available in Windows Server 2008 R2. The Live Migration feature enables you to migrate a virtual machine from one host to another without downtime. Unlike the quick migration process, Live Migration does not save the state of virtual machine; instead, it synchronizes the state during failover.
- **Hyper-V Replica.** This new feature in Windows Server 2012 enables you to replicate rather than move a virtual machine to another host, and to synchronize all virtual machine changes from the primary host to the host that contains the replica.

- Exporting and importing virtual machine. This is an established method of moving virtual machines without using a cluster. You export a virtual machine on one host, and then physically move exported files to another host by performing an import operation. This is a time-consuming operation that requires you to turn off the virtual machines during export and import. In Windows Server 2012, this migration method is improved. You can import a virtual machine to a Hyper-V host without exporting it before import. The Hyper-V role in Windows Server 2012 is now capable of configuring all the necessary settings during the import operation.

How Does Virtual Machine and Storage Migration Work?

There are many cases in which an administrator might want to move virtual machine files to another location. For example, if the disk on which a virtual machine hard disk resides runs out of space, you must move the virtual machine to another drive or volume. Moving virtual machines to other hosts is a common procedure.

In Windows Server 2008 and Windows Server 2008 R2, moving a virtual machine resulted in downtime because the virtual machine had to be turned off. If you moved a virtual machine between two hosts, then you also had to perform export and import operations for that specific virtual machine. Export operations can be time-consuming, depending on the size of the virtual hard disks.

Storage migration enables you to move virtual machines and their storage to other locations without downtime

- During migration the virtual machine hard drive is copied from one location to another
- Changes are written to both the source drive and the destination drive
- You can move virtual machine storage to the same host, another host, or to a SMB share
- Storage and virtual machine configuration can be in different locations

In Windows Server 2012, virtual machine and storage migration enables you to move a virtual machine and its storage to another location on the same host, or to another host computer, without having to turn off the virtual machine.

Virtual machine and storage migration works as follows:

1. To copy a virtual hard disk, start live storage migration by using the Hyper-V console. Optionally, you can use Windows PowerShell cmdlets.
2. The migration process creates a new virtual hard disk in the destination location and starts the copy process.
3. During the copy process, the virtual machine is fully functional. However, all changes that occur during the copy process are written to both the source and destination locations. Read operations are performed only from the source location.
4. As soon as the disk copy process completes, Hyper-V switches virtual machines to run on the destination virtual hard disk. In addition, if you are moving the virtual machine to another host, the computer configuration is copied and the virtual machine is associated with the host. If a failure occurs on the destination side, there is always a fallback option to run back again on the source directory.
5. After the virtual machine completes migration, the process deletes the source virtual hard disks.

The time that is required to move a virtual machine depends on the source and destination locations, the speed of the hard disks, storage, or network, and the size of the virtual hard disks. The move process is faster if the source and destination locations are on storage, and the storage supports .odx files. Instead of using buffered read and buffered write operations, the .odx file starts the copy operation with an offload read command and retrieves a token representing the data from the storage device. It then uses

an offload write command with the token to request data movement from the source disk to the destination disk.

When you move a virtual machine's virtual hard disks to another location, the Virtual Machine Move wizard presents three available options:

- **Move all the virtual machine's data to a single location.** Specify a single destination location, such as disk file, configuration, snapshot, and smart paging.
- **Move the virtual machine's data to a different location.** Specify individual locations for each virtual machine item.
- **Move only the virtual machine's virtual hard disk.** Move only the virtual hard disk file.


How Does Live Migration Work?

Live migration enables you to move running virtual machines from one failover cluster node to another node in the same cluster. With live migration, users who are connected to the virtual machine should experience almost no server outage.



Note: Whereas you can also perform live migration of virtual machines by using the virtual machine and storage migration method described in the previous topic, you should be aware that live migration is based on failover clustering.

Unlike the storage migration scenario, you can only perform live migration if the virtual machines are highly available.



You can initiate live migration through one of the following methods:

- The Failover Cluster Manager console.
- The Virtual Machine Manager Administrator Console, if you use VMM to manage your physical hosts.
- A Windows Management Instrumentation (WMI) or Windows PowerShell script.



Note: Live migration enables you to significantly reduce the perceived outage of a virtual machine during a planned failover. During a planned failover, you start the failover manually. Live migration does not apply during an unplanned failover, such as when the node hosting the virtual machine fails.

Live Migration Process

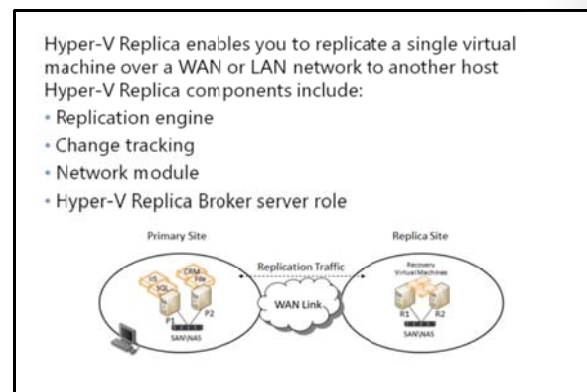
The live migration process consists of four steps that occur in the background:

1. **Migration setup.** When you start the failover of the virtual machine, the source node creates a Transmission Control Protocol (TCP) connection with the target physical host. This connection is used to transfer the virtual machine configuration data to the target physical host. Live migration creates a temporary virtual machine on the target physical host, and allocates memory to the destination virtual machine. The migration preparation also verifies that the virtual machine can be migrated.

2. Guest memory transfer. The guest memory is transferred iteratively to the target host while the virtual machine is still running on the source host. Hyper-V on the source physical host monitors the pages in the working set. As the system modifies memory pages, it tracks and marks them as being modified. During this phase of the migration, the migrating virtual machine continues to run. Hyper-V iterates the memory copy process several times, and each time a smaller number of modified pages are copied to the destination physical computer. A final memory copy process copies the remaining modified memory pages to the destination physical host. Copying stops as soon as the number of pages that have been modified in physical memory but not yet rewritten to disk—often called *dirty pages*—drops below a threshold, or after 10 iterations are complete.
3. State transfer. To actually migrate the virtual machine to the target host, Hyper-V stops the source partition, transfers the state of the virtual machine (including the remaining dirty memory pages) to the target host, and then restores the virtual machine on the target host. The virtual machine pauses during the final state transfer.
4. Clean up. The cleanup stage finishes the migration by dismantling the virtual machine on the source host, terminating the worker threads, and signaling the completion of the migration.

How Does Hyper-V Replica Work?

In some cases, you might want to have a spare copy of one virtual machine that you can run if the original virtual machine fails. However, when you implement high availability, you have only one instance of a virtual machine. High availability does not prevent corruption of software that is running inside the virtual machine. One way to address the issue of corruption is to copy the virtual machine. You can also back up the virtual machine and its storage. Although this solution achieves the desired result, it is resource-intensive and time-consuming.



To resolve this problem and to enable administrators to have an up-to-date copy of a single virtual machine, Microsoft has implemented Hyper-V Replica, a feature in Windows Server 2012. This feature enables virtual machines that are running at a primary site location or host to be replicated to a secondary site location or host across a WAN or LAN link. Hyper-V replica enables you to have two instances of a single virtual machine residing on different hosts: one as the primary (live) copy, and the other as a replica (offline) copy. These copies are synchronized, and you can perform failover at any time. In the event of a failure at a primary site, you can use Hyper-V Replica to execute a failover of the production workloads to replica servers at a secondary location within minutes, thus incurring minimal downtime.

The site configurations do not have to use the same server or storage hardware. Hyper-V Replica enables an administrator to restore virtualized workloads to a point in time depending on the Recovery History selections for the virtual machine.

Hyper-V Replica technology consists of several components:

- Replication engine: The replication engine manages the replication configuration details and manages initial replication, delta replication, failover, and test-failover operations. It also tracks virtual machine and storage mobility events, and takes appropriate actions as needed. That is, it pauses replication events until migration events complete, and then resumes where they left off.
- Change tracking: This component tracks changes that occur on the primary copy of the virtual machine. It is designed to track the changes regardless of where the virtual machine .vhdx files reside.

- Network module: The network module provides a secure and efficient way to transfer virtual machine replicas between primary hosts and replica hosts. It uses data compression, which is enabled by default. The transfer operation is secure because it relies on HTTPS and certification-based authentication.
- Hyper-V Replica Broker server role: This is a new server role that is implemented in Windows Server 2012, and you configure it during failover clustering. This server role enables you to have Hyper-V Replica functionality even when the virtual machine being replicated is highly available and can move from one cluster node to another. The Hyper-V Replica Broker server redirects all virtual machine specific events to the appropriate node in the replica cluster. The Broker queries the cluster database to determine which node should handle which events. This ensures that all events are redirected to the correct node in the cluster in the event that a quick migration, live migration, or storage migration process was executed.

Configuring Hyper-V Replica

Before you implement Hyper-V Replica, ensure that your infrastructure meets the following prerequisites:

- The server hardware supports the Hyper-V role on Windows Server 2012. Also, ensure that the server hardware has sufficient capacity to run all of the virtual machines to which you replicate it.
- Sufficient storage exists on both the primary and replica servers to host the files that replicated virtual machines use.
- Network connectivity exists between the locations that are hosting the primary and replica servers. Connectivity can be through a WAN or LAN link.
- Firewall rules are configured correctly to enable replication between the primary and replica sites. By default, traffic uses TCP port 80 or port 443.
- If you want to use certificate-based authentication, ensure that an X.509v3 certificate exists to support mutual authentication with certificates.

To configure Hyper-V replica, do the following:

1. Configure authentication options
2. Configure ports
3. Select replica servers
4. Select location for replica files
5. Enable replication on virtual machine

You do not have to install Hyper-V Replica separately. Hyper-V Replica is implemented as part of the Hyper-V server role. You can use it on Hyper-V servers that are standalone or servers that are part of a failover cluster (in which case, you should configure the Hyper-V Replica Broker server role). Unlike failover clustering, a Hyper-V role is not dependent on AD DS. You can use it with Hyper-V servers that are standalone, or that are members of different Active Directory domains (except in case when servers are part of a failover cluster).

To enable Hyper-V Replica, first configure the Hyper-V server settings. In the **Replication Configuration** group of options, enable the Hyper-V server as a replica server, and select the authentication and port options. You should also configure authorization options. You can choose to enable replication from any server that successfully authenticates, which is convenient in scenarios where all servers are part of same domain. Alternatively, you can type fully qualified domain names (FQDNs) of servers that you accept as replica servers. In addition, you must configure the location for replica files. You should configure these settings on each server that will serve as a replica server.

After you configure options at the server level, enable replication on a virtual machine. During this configuration, you must specify the replica server name and options for the connection. If the virtual

machine has more than one virtual hard disk, you can select which virtual hard disk drives that you want to replicate. You can also configure the recovery history and the initial replication method. Start the replication process after you configure these options.

Demonstration: Implementing Hyper-V Replica (optional)

In this demonstration, you will see learn how to implement Hyper-V Replica.

Demonstration Steps

Configure a replica

1. On LON-HOST1 and LON-HOST2, configure each server to be a Hyper-V Replica server.
 - o Authentication: **Kerberos (HTTP)**
 - o **Allow replication from any authenticated server**
 - o Create and use folder **E:\VMReplica** as a default location to store replica files.
2. Enable the firewall rule named **Hyper-V Replica HTTP Listener (TCP-In)** on both hosts.

Configure replication

1. On LON-HOST1, enable replication for the 20412A-LON-CORE virtual machine.
 - o Authentication: **Kerberos (HTTP)**
 - o Select to have only the latest recovery point available.
 - o Start replication immediately.
2. Wait for initial replication to finish and ensure that the 20412A-LON-CORE virtual machine appears in Hyper-V Manager on LON-HOST2.

Lesson 4

Managing Hyper-V Virtual Environments by Using VMM

VMM is member of the System Center 2012 family of products. It is a successor of System Center Virtual Machine Manager 2008 R2. VMM extends management functionality for Hyper-V hosts and virtual machines, and it provides deployment and provisioning for virtual machines and services. In this lesson, you will learn the basics of VMM.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe VMM.
- Describe the prerequisites for installing VMM.
- Describe private cloud infrastructure components.
- Describe how to manage hosts and host groups with VMM.
- Describe how to deploy virtual machines with VMM.
- Describe services and service templates.
- Describe physical-to-virtual (P2V) and virtual-to-virtual (V2V) migrations.
- Describe considerations for deploying a highly available VMM server.

What Is System Center 2012 - Virtual Machine Manager?

VMM is a management solution for a virtualized data center. VMM enables you to create and deploy virtual machines and services to private clouds by configuring and managing your virtualization host, networking, and storage resources. You can also use VMM to manage VMware ESX and Citrix XenServer hosts.

VMM is a component of System Center 2012 that discovers, captures, and aggregates knowledge of the virtualization infrastructure. VMM also manages policies, processes, and best practices by discovering, capturing, and aggregating knowledge of the virtualization infrastructure.

VMM succeeds VMM 2008 R2, and is a key component in enabling private cloud infrastructures, which helps transition enterprise IT from an infrastructure-focused deployment model into a service-oriented, user-centric environment.

The VMM architecture consists of several interrelated components. These components are:

- **Virtual Machine Manager server.** The Virtual Machine Manager server is the computer on which the VMM service runs. The Virtual Machine Manager server processes commands and controls communications with the Virtual Machine Manager database, the library server, and the virtual machine hosts. The Virtual Machine Manager server is the hub of a VMM deployment through which all other VMM components interact and communicate. The Virtual Machine Manager server also connects to a Microsoft SQL Server® database that stores all VMM configuration information.

VMM provides centralized administration and management of your virtual environment

VMM is used to:

- Manage Hyper-V hosts
- Manage VMware ESX and Citrix XenServer hosts
- Manage and deploy virtual machines
- Manage and deploy services
- Perform P2V and V2V conversions

- Virtual Machine Manager database. VMM uses a SQL Server database to store the information that you view in the VMM management console, such as managed virtual machines, virtual machine hosts, virtual machine libraries, jobs, and other virtual machine-related data.
- VMM management console. The management console is a program that you use to connect to a VMM management server, to view and manage physical and virtual resources, including virtual machine hosts, virtual machines, services, and library resources.
- Virtual Machine Manager library. A *library* is a catalog of resources, such as virtual hard disks, templates, and profiles, which are used to deploy virtual machines and services. A library server also hosts shared folders that store file-based resources. The VMM management server is always the default library server, but you can add additional library servers later.
- Command shell. Windows PowerShell is the command-line interface that you use to execute cmdlets that perform all available VMM functions. You can use these VMM-specific cmdlets to manage all the actions in a VMM environment.
- Self-Service Portal. The Self-Service Portal is a website that users who are assigned to a self-service user role can use to deploy and manage their own virtual machines.

Prerequisites for Installing VMM 2012

Before you deploy VMM and its components, ensure that your system meets the hardware and software requirements. While software requirements do not change based on the number of hosts that VMM manages, hardware prerequisites may vary. In addition, not all VMM components have the same hardware and software requirements.



Note: Windows Server 2008 R2 and Windows Server 2012 are the only supported operating systems for VMM 2012.

Software requirements for the Virtual Machine Manager server are:

- Windows Server 2008 R2
- SQL Server 2008 SP2 or SQL Server 2008 R2
- Microsoft .NET Framework 3.5 SP1 or later
- Windows AIK
- Windows PowerShell 2.0 if the Virtual Machine Manager management console will run on the same server
- WinRM 2.0

Hardware requirements:

- CPU: single core CPU 2 GHz
- RAM: 4 – 8 GB
- Disk space: 40 GB – 150 GB

The number of hosts determines the hardware requirements

Virtual Machine Manager Server

In addition to having Windows Server 2008 R2 or Windows Server 2012 installed, the following software must be installed on the server that will run the Virtual Machine Manager server:

- Microsoft .NET Framework 3.5 Service Pack 1 (SP1) or newer
- Windows Automated Installation Kit (AIK)
- Windows PowerShell 2.0, if the VMM management console will run on the same server
- Windows Remote Management 2.0. Note that this is installed by default in Windows Server 2008 R2, so you should just verify that the service is running.
- SQL Server 2008 Service Pack 2 (SP2) (Standard or Enterprise) or SQL Server 2008 R2 SP1 Standard, Enterprise, or Datacenter. This is necessary only when you install the VMM management server and SQL Server on same computer.

Hardware requirements vary depending on number of hosts, and have the following limits:

- CPU: Single core CPU 2 gigahertz (GHz), Dual core CPU 2.8 GHz

- Random access memory (RAM): 4–8 gigabytes (GB)
- Disk space: 40 GB – 150 GB, depending on whether you install a SQL Server database on the same server. In addition, if the library is on the same server, then disk space will also depend on library content.

Virtual Machine Manager Database

The Virtual Machine Manager database stores all VMM configuration information, which you can access and modify by using the VMM management console. The Virtual Machine Manager database requires SQL Server 2008 SP2 or newer. Because of this, the base hardware requirements for the Virtual Machine Manager database are equal to the minimum system requirements for installing SQL Server. Additionally, if you are managing more than 150 hosts, you should have at least 4 GB of RAM on the database server. Software requirements for the Virtual Machine Manager database are the same as for SQL Server.

Virtual Machine Manager Library

The Virtual Machine Manager library is the server that hosts resources for building virtual machines, services, and private clouds. In smaller environments, you usually install the Virtual Machine Manager library on the VMM management server. If this is the case, the hardware and software requirements are the same as for the VMM management server. In larger and more complex environments, we recommend that you maintain Virtual Machine Manager library on a separate server in a highly available configuration. If you want to deploy another Virtual Machine Manager library server, the server should meet the following requirements:

- Supported operating system: Windows Server 2008 or Windows Server 2008 R2
- Hardware management: Windows Remote Management 2.0
- CPU: at least 2.8 GHz
- RAM: at least 2 GB
- Hard disk space: varies based on the number and size of files that are stored

Private Cloud Infrastructure Components in VMM

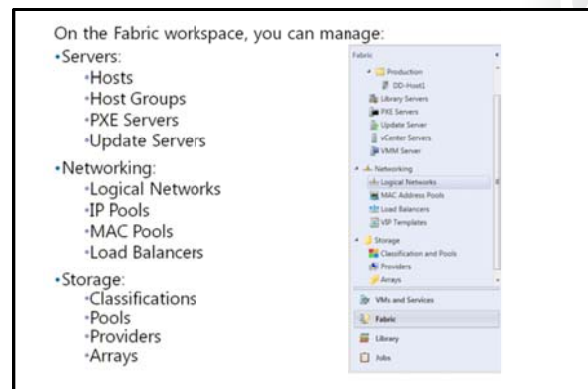
The key architectural concept in VMM is the private cloud infrastructure. Similar to public cloud solutions, such as in Windows Azure™, the private cloud infrastructure in VMM is an abstraction layer that shields the underlying technical complexities and allows you to manage defined resource pools that consist of servers, networking, and storage, in the enterprise infrastructure.

By using the VMM management console user interface (UI), you can create a private cloud from Hyper-V, VMware ESX, and Citrix XenServer hosts.

You can also benefit from cloud computing attributes, including self-servicing, resource pooling, and elasticity.

You can configure the following resources from the Fabric workspace in the VMM management console:

- Servers. In the Servers node, you can configure and manage several types of servers. Host groups contain virtualization hosts, which are the destinations you can use to deploy virtual machines. Library



servers are the repositories of building blocks—such as images, .iso files, and templates—for creating virtual machines.

- **Networking.** The Networking node is where you can define logical networks, assign pools of static IPs and media access control (MAC) addresses, and integrate load balancers. Logical networks are user-defined groupings of IP subnets and virtual local area networks (VLANs) that organize and simplify network assignments. Logical networks provide an abstraction of the underlying physical infrastructure, and they enable you to provision and isolate network traffic based on selected criteria such as connectivity properties and service level agreements (SLAs).
- **Storage.** You can discover, classify, and provision remote storage on supported storage arrays. VMM uses the Microsoft Storage Management Service—that is enabled by default during the installation of VMM—to communicate with external arrays.

Managing Hosts, Host Clusters, and Host Groups with VMM

In addition to virtual machine management, VMM can also help you manage and deploy Hyper-V hosts. In VMM, you can use technologies such as Windows Deployment Services (Windows DS) to deploy Hyper-V hosts on bare-metal machines and then manage them with VMM. When hosts are associated with VMM, you can configure several options, such as host reserves, quotas, permissions, and cloud memberships. VMM can also manage Hyper-V failover clusters.

VMM provides two new features that help optimize power and resource usage on hosts that are managed by VMM: dynamic optimization and power optimization. Dynamic optimization balances the virtual machine load within a host cluster, while power optimization enables VMM to evacuate balanced cluster hosts, and then turn them off to save power.

The recommended way to organize hosts in VMM is to create host groups. This simplifies management tasks. A host group enables you to apply settings to multiple hosts or host clusters with a single action. By default, there is a single host group in VMM named All Hosts. However, if necessary, you can create additional groups for your environment.

Host groups are hierarchical. When you create a new child host group, it inherits the settings from the parent host group. When a child host group moves to a new parent host group, the child host group maintains its original settings except for Performance and Resource Optimization (PRO) settings, which are managed separately. When the settings in a parent host group change, you can apply those changes to child host groups.

You use host groups in the following scenarios:

- Provide basic organization when you are managing many hosts and virtual machines. You can create custom views within the Hosts view and the Virtual Machines view to provide easy monitoring and access to a host. For example, you might create a host group for each branch office in your organization.
- Reserving resources for use by hosts. Host reserves are useful when placing virtual machines on a host. Host reserves determine the CPU, memory, disk space, disk I/O capacity, and network capacity that are continuously available to the host operating system.

- VMM can deploy and manage Hyper-V hosts, Hyper-V clusters, and host groups
- Host groups simplify management tasks by using a single action to apply settings to multiple hosts or host clusters
- Host group scenarios:
 - Provide basic organization when managing large numbers of hosts
 - Reserve resources for use by hosts
 - Designate hosts on which a user can create and operate their own virtual machines
 - Create private clouds

- Use the Host Group properties action for the root host group All Hosts, to set default host reserves for all hosts that VMM manages. If you want to use more of the resources on some hosts instead of on other hosts, you can set host reserves differently for each host group.
- Designating hosts on which users can create and operate their own virtual machines. When a VMM administrator adds self-service user roles, one part of role creation is to identify the hosts on which self-service users or groups in that role can create, operate, and manage their own virtual machines. As a best practice, you should designate a specific host group for this purpose.

Deploying Virtual Machines with VMM

One of the advantages of using VMM to manage a virtualized environment is the flexibility that VMM provides to create and deploy new virtual machines quickly.

Using VMM, you can manually create a new virtual machine with new configuration settings and a new hard disk. You can then deploy the new virtual machine from one of following sources:

- An existing .vhd or .vhdx file, either blank or preconfigured
- A virtual machine template
- A Virtual Machine Manager library

You can create new virtual machines either by converting an existing physical computer, or by cloning an existing virtual machine.

Creating a New Virtual Machine from an Existing VHD

You can create a new virtual machine based on either a blank virtual hard disk (VHD) or a preconfigured VHD that contains a guest operating system. VMM provides two blank VHD templates that you can use to create new disks:

- Blank Disk – Small
- Blank Disk – Large

You can also use a blank VHD when you want to use an operating system with a preboot execution environment (PXE). Alternatively, you can place an .iso image on a virtual DVD-ROM, and then install an operating system from scratch. This is an effective way to build a virtual machine's source image, which you can then use as a future template. To install the operating system on such a virtual machine, you can use an .iso image file from the library or from a local disk, then map a physical drive from the host computer, or start the guest operating system setup through a network service boot.

If you have a library of VHDs that you want to use in your VMM environment, you can create a virtual machine from an existing VHD. You can also select existing VHDs when you deploy any operating system from which VMM cannot create a template, such as an operating system that is not Windows-based.

When you create a new virtual machine using an existing VHD, you are essentially creating a new virtual machine configuration that is associated with the VHD file. VMM will create a copy of the source VHD so that you do not have to move or modify the original.

In VMM there are several ways to create and deploy new virtual machines:

- Create a new virtual machine from an empty hard disk
- Create a new virtual machine based on a predefined template
- Deploy a new virtual machine from the Virtual Machine Manager library

In this scenario, the source VHD must meet the following requirements:

- Leave the Administrator password blank on the VHD as part of the System Preparation Tool (Sysprep) process.
- Install the Virtual Machine Additions on the virtual machine.
- Use Sysprep to prepare the operating system for duplication.



Note: VMM 2012 will support the .vhdx virtual hard drive format when VMM 2012 SP1 is released.

Deploying from a Template

You can create a new virtual machine based on a template from the Virtual Machine Manager library. The template is a library resource, which links to a virtual hard disk drive that has a generalized operating system, hardware settings, and guest operating system settings. You use the guest operating system settings to configure operating system settings such as the computer name, local administrator password, and domain membership.

The deployment process does not modify the template, which you can reuse multiple times. If you are creating virtual machines in the Self-Service Portal, you must use a template.

The following requirements apply if you want to deploy a new virtual machine from a template:

- You must install a supported operating system on the VHD.
- You must leave the Administrator password blank on the VHD as part of the Sysprep process. However, you do not have to leave the Administrator password blank for the guest operating system profile.
- For customized templates, you must prepare the operating system on the VHD by removing the computer identity information. For Windows operating systems, you can prepare the VHD by using the Sysprep tool.

Deploying from the Virtual Machine Manager Library

If you deploy a virtual machine from the Virtual Machine Manager library, the virtual machine is removed from the library, and then placed on the selected host. When you use this method, you must provide the following details in the Deploy Virtual Machine wizard:

- The host for deployment. The template that you use provides a list of potential hosts and their ratings.
- The path of the virtual machine files on the host.
- The virtual networks used for the virtual machine. A list of existing virtual networks on the host will display.

What Are Services and Service Templates?

Services are a new concept in VMM. You must understand services fully before you deploy a private cloud infrastructure.

Traditional Services Scenario

Services usually refer to applications or sets of applications that provide services to end users. For example, you can deploy various types of web-based services, but you can also implement a service such as email. In a non-cloud computing scenario, deployment of any type of service usually requires users, developers, and administrators to work together through the phases of creating a service, deploying a service, testing the service, and maintaining the service.

A service frequently includes several computers that must work together to provide a service to end users. For example, a web-based service is usually an application that deploys on a web server, connects to a database server that might be hosted on another computer, and performs authentication on an Active Directory domain controller. Enabling this application requires three roles, and possibly three computers: a web server, a database server, and a domain controller. Deploying a test environment for a service such as this can be time consuming and resource consuming. Ideally, developers work with IT administrators to create an environment where they can deploy and test their web application.

Concept of a Service in a Private Cloud Scenario

With the concept of a private cloud, how you deal with services can change significantly. You can prepare the environment for a service, and then let developers deploy it by using a self-service application such as System Center 2012 - App Controller.

In VMM, a *service* is a set of one or more virtual machines that you deploy and manage together as a single entity. You configure these machines to run together to provide a service. In Windows Server 2008, users could deploy new virtual machines by using the Self-Service Portal. In VMM, end users can deploy new services. By deploying a service, users are actually deploying the entire infrastructure, including the virtual machines, network connections, and applications that are required to make the service work.

However, you can also use services to deploy only a single virtual machine without any specific purpose. Instead of deploying virtual machines in the historic way, you can now create a service that will deploy a virtual machine with, for example, Windows Server 2008 R2, and with several roles and features preinstalled and joined to domain. This simplifies the process of creating and later updating new virtual machines.

Deploying a new service requires a high level of automation and predefined components, and requires management software support. This is why VMM provides service templates. A service template is a template that encapsulates everything that is required to deploy and run a new instance of an application. Just as a private cloud user can create new virtual machines on demand, the user can also use service templates to install and start new applications on demand.

Process for Deploying a New Service

You use the following procedure when you use VMM service templates to deploy a new service or application:

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

- In terms of VMM clouds, a service is a set of one or more virtual machines that are deployed together and managed as a single entity
- A service template encapsulates all necessary components required to deploy and run a new instance of an application
- A service is deployed by an administrator or end user
- A service can contain several different components
- A service can be deployed to a private cloud or to a host group
- The administrator creates a service template in VMM
- The application owner deploys a service based on the service template
- The administrator uses App Controller or VMM Manager to deploy the service based on the template

2. The end-user application owner—for example, a developer who has to deploy the application environment—opens App Controller and requests a new service deployment based on the available service templates that the developer can access. The developer can deploy the service to a private cloud where a user has access. As an alternative to App Controller, the user can also use the VMM Manager console.
3. The Virtual Machine Manager server evaluates the submitted request. VMM searches for available resources in the private cloud, then calculates the user quota and verifies that the private cloud has enough resources for the requested service deployment.
4. Whereas the new service is created automatically, the virtual machines and applications (if any) are deployed on the host that is selected by VMM.
5. The user application owner gains control over service virtual machines through App Controller, or by Remote Desktop Protocol (RDP).
6. If you need manual approval for resource creation, you can use Microsoft System Center 2012 - Service Manager to create workflows for this purpose.

Information Included in the Service Template

The service template includes information about the virtual machines that are deployed as part of the service, which applications to install on the virtual machines, and the networking configuration needed for the service (including the use of an NLB). The service template can use existing virtual machine templates. While you can define the service without using any existing virtual machine templates, it is easier to build a template if you have already created virtual machine templates. After you create a service template, you configure it for deployment using the **Configure Deployment** option.

P2V and V2V Migrations

Many organizations have physical servers that they do not use fully. VMM can convert existing physical computers into virtual machines through a process known as P2V conversion. VMM simplifies P2V by providing a task-based wizard to automate much of the conversion process. Because the P2V process supports scripts, you can start large-scale P2V conversions through the Windows PowerShell command line interface.

VMM converts an operating system that is running on physical hardware to an operating system that is running in a Hyper-V virtual machine environment. VMM provides a conversion wizard, which automates much of the conversion process.

During a P2V conversion process, VMM generates disk images of the hard disks on the physical computer. It creates VHD files for the new virtual machine using the disk images as a basis. In addition, it creates a hardware configuration for the virtual machine similar to, or the same as, the hardware in the physical computer.

The new virtual machine has the same computer identity as the physical computer on which it is based. Because of this, as a best practice you should not use both a physical computer and its virtual replica concurrently. After the P2V conversion completes, you typically disconnect the physical computer from the network and decommission it.

The P2V process converts an operating system that is running on physical hardware to an operating system running inside a virtual machine

The V2V process converts existing VMware virtual machines to virtual machines running on Hyper-V

P2V conversion is finished in either online or offline mode. In online mode, the source operating system runs during the conversion process. In offline mode, the operating system does not run and conversion occurs through the Windows Preinstallation Environment (Windows PE). Later topics in this lesson describe these modes in further detail.

In addition to converting underused physical computers, VMM supports the management, migration, and conversions of other virtual machines that were created in the VMware environment. You can convert these virtual machines to Hyper-V virtual machines, place them on Hyper-V hosts, and then manage them under the Virtual Machine Manager Administrator Console. In addition, VMM and Hyper-V support migrating virtual machines from one host to another with minimal or zero downtime.

VMM allows you to convert existing VMware virtual machines to virtual machines running on the Hyper-V platform. This process is known as a *V2V conversion*. With V2V conversion, administrators can consolidate a virtual environment that is running various virtual platforms without moving data or rebuilding virtual machines from scratch.

VMM allows you to copy existing VMware virtual machines and create Hyper-V virtual machines. You can copy VMware virtual machines that are located on ESX server hosts, in Virtual Machine Manager libraries, or on Windows shares. Although V2V is called a conversion, V2V is a read-only operation that does not delete or affect the original source virtual machine. In addition, the term conversion is dedicated only to the process of converting VMware virtual machines. The term migration is used for virtual server machines.

During the conversion process, the VMM converts the VMware .vmdk files to .vhd files, and makes the operating system on the virtual machine compatible with Microsoft virtualization technologies. The virtual machine that the wizard creates matches VMware virtual machine properties, including name, description, memory, and disk-to-bus assignments.

Considerations for Deploying a Highly Available Virtual Machine Manager Server

VMM now supports a highly available Virtual Machine Manager server. You can use failover clustering to achieve high availability for VMM because VMM is now a cluster-aware application. However, you should consider several things before you deploy a VMM cluster.

Before installing a highly available VMM management server, ensure that:

- You have installed and configured a failover cluster that is running Windows Server 2008 R2, Windows Server 2008 R2 SP1, or Windows Server 2012.
- All computers on which you install the highly available Virtual Machine Manager server meet the minimum hardware requirements, and all prerequisite software is installed on all computers.
- You have created a domain account to be used by the VMM service. You must use a domain user account for a highly available Virtual Machine Manager server.
- You are prepared to use distributed key management to store encryption keys in AD DS. You must use distributed key management for a highly available Virtual Machine Manager server.

Considerations for deploying a highly available VMM server:

- VMM is cluster-aware and can be highly available
- When deploying VMM in a cluster, the service account must be the domain account
- Use distributed key management for encryption keys
- Make database and library servers highly available
- Do not install a Virtual Machine Manager Self Service Portal on a clustered Virtual Machine Manager server
- Use the Failover Cluster Manager to perform a planned failover

- You have a computer with a supported SQL Server version installed and running. Unlike VMM 2008 R2, VMM will not install a SQL Server Express Edition automatically.

Highly Available Databases and Library Servers

To achieve full redundancy, you should use a highly available SQL Server. Install a highly available SQL Server on a separate failover cluster from the failover cluster on which you are installing the highly available Virtual Machine Manager server. Similarly, you should also use a highly available file server for hosting your library shares.

Self Service Portal and Clustered Virtual Machine Manager Server

As a best practice, do not install the Virtual Machine Manager Self-Service Portal on the same computer as the highly available Virtual Machine Manager server. If your Virtual Machine Manager Self-Service Portal currently resides on the same computer as the Virtual Machine Manager server, we recommend that you uninstall the Virtual Machine Manager Self-Service Portal for VMM 2008 R2 SP1 before upgrading to VMM. We also recommend that you install the Virtual Machine Manager Self-Service Portal on a highly available web server to achieve redundancy and NLB.

Failover Cluster Manager

You cannot perform a planned failover—for example, to install a security update or perform maintenance on a cluster node—by using the Virtual Machine Manager Administrator Console. Instead, to perform a planned failover, use the Failover Cluster Manager.

During a planned failover, ensure that there are no tasks actively running on the Virtual Machine Manager server. Any tasks that are executing during a failover will be stopped and will not restart automatically. Any connections to a highly available Virtual Machine Manager server from the Virtual Machine Manager Administrator Console or the Virtual Machine Manager Self-Service Portal will also be lost during a failover. However, the Virtual Machine Manager Administrator Console can reconnect automatically to the highly available Virtual Machine Manager server after a failover if the console was open before you performed the failover.

Lab: Implementing Failover Clustering with Hyper-V

Scenario

The A. Datum Corporation's initial virtual machine deployment on Hyper-V has been successful. As a next step in the deployment, A. Datum is now considering ways to ensure that the services and applications that are deployed on the virtual machines are highly available. As part of the implementation, A. Datum is also considering options for making the virtual machines that run on Hyper-V highly available.

As one of the senior network administrators at A. Datum, you are responsible for integrating Hyper-V with failover clustering to ensure that the virtual machines that are deployed on Hyper-V are highly available. You are responsible for planning the virtual machine and storage configuration, and for implementing the virtual machines as highly available services on the failover cluster. In addition, you are considering some other techniques for virtual machine high availability, such as Hyper-V Replica.

Objectives

- Configure Hyper-V Replica.
- Configure a failover cluster for Hyper-V.
- Configure a highly available virtual machine.

Lab Setup

20412A-LON-DC1-B

20412A-LON-SVR1-B

20412A-LON-HOST1

20412A-LON-HOST2

Estimated time: **75 minutes**

Virtual Machine(s)	20412A-LON-DC1-B 20412A-LON-SVR1-B 20412A-LON-HOST1 20412A-LON-HOST2
User Name	Adatum\Administrator
Password	Pa\$\$w0rd

You should perform this lab with a partner. To perform this lab, you must boot the host computers to Windows Server 2012. Ensure that you and your partner have booted into different hosts (one should boot to 20412A-LON-HOST1 and the other should boot to 20412A-LON-HOST2) and then log on as **Adatum\Administrator** with the password of **Pa\$\$w0rd**. Once you have booted into the Windows Server 2012 environment, perform the following setup tasks:

1. On the host computer, in Server Manager, click **Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, click start the following virtual machines based upon your host:
 - For LON-HOST1, start **20412A-LON-DC1-B**.
 - For LON-HOST2, start **20412A-LON-SVR1-B**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on using the following credentials:
 - a. User name: **Adatum\Administrator**

- b. Password: **Pa\$\$w0rd**



Note: For this lab, verify that the classroom is configured so that only LON-HOST1 and LON-HOST2 can communicate. Each pair of host computers must be isolated from the rest of the classroom.

Exercise 1: Configuring Hyper-V Replicas

Scenario

Before you begin cluster deployment, you need to evaluate the new technology in Hyper-V 3.0 for replicating virtual machines between hosts. You want to be able to mount a copy of a virtual machine on another host manually if the active copy (or host) fails.

The main tasks for this exercise are as follows:

1. Boot the physical host machines from VHD.
2. Import the LON-CORE virtual machine on LON-HOST1.
3. Configure a replica on both host machines.
4. Configure replication for the LON-CORE virtual machine.
5. Validate a planned failover to the replica site.

► Task 1: Boot the physical host machines from VHD

1. Restart the classroom computer, and in the **Windows Boot Manager**, select either **20412A-LON-HOST1** or **20412A-LON-HOST2**.



Note: If you start LON-HOST1, your partner must start LON-HOST2.

2. Log on to the server as **Adatum\Administrator** with password **Pa\$\$w0rd**.
3. On LON-HOST1, make sure that virtual machine **20412A-LON-DC1** is running.
4. On LON-HOST2, make sure that virtual machine **20412A-LON-SVR1** is running.

► Task 2: Import the LON-CORE virtual machine on LON-HOST1

- On LON-HOST1, open Hyper-V Manager, and import the **20412A-LON-CORE** virtual machine using the following settings:
 - Path: **E:\Program Files\Microsoft Learning\20412\Drives\20412A-LON-CORE**
 - Accept default values



Note: The drive letter may differ based on the number of drives on the physical host machine.

► Task 3: Configure a replica on both host machines

1. On LON-HOST1 and LON-HOST2, configure each server to be a Hyper-V Replica server.
 - Authentication: **Kerberos (HTTP)**
 - **Allow replication from any authenticated server**
 - Create and use folder **E:\VMReplica** as a default location to store replica files
2. Enable the firewall rule named **Hyper-V Replica HTTP Listener (TCP-In)** on both hosts.

► **Task 4: Configure replication for the LON-CORE virtual machine**

1. On LON-HOST1, enable replication for the **20412A-LON-CORE** virtual machine.
 - Replica server: **LON-HOST2**
 - Authentication: **Kerberos authentication (HTTP)**
 - Configure Recovery History: **Only the latest recovery point**
 - **Start replication immediately**
2. Wait for initial replication to finish, and verify that the 20412A-LON-CORE virtual machine displays in the Hyper-V Manager console on LON-HOST2.

► **Task 5: Validate a planned failover to the replica site**

1. On LON-HOST2, view replication health for **20412A-LON-CORE**.
2. On LON-HOST1, perform planned failover to **LON-HOST2**. Verify that 20412A-LON-CORE is running on LON-HOST2.
3. On LON-HOST1, remove replication for **20412A-LON-CORE**.
4. On LON-HOST2, shut down 20412A-LON-CORE.

Results: After completing this exercise, you will have configured a Hyper-V Replica.

Exercise 2: Configuring a Failover Cluster for Hyper-V

Scenario

A. Datum has several virtual machines that are hosting important services that must be highly available. Because these services are not cluster-aware, A. Datum decided to implement Failover cluster on the Hyper-V host level. You plan to use iSCSI drives as storage for these virtual machines.

The main tasks for this exercise are as follows:

1. Connect to the iSCSI target from both host machines.
2. Configure failover clustering on both host machines.
3. Configure disks for the failover cluster.

► **Task 1: Connect to the iSCSI target from both host machines**

1. On LON-HOST1, start **iSCSI initiator**.
2. Use **172.16.0.21** for the address that will be used to discover and connect to the iSCSI target.
3. On LON-HOST2, start **iSCSI initiator**.
4. Use **172.16.0.21** for the address that will be used to discover and connect to the iSCSI target.
5. On LON-HOST2, navigate to **Disk Management**, and initialize and bring online all iSCSI drives:
 - Format the first drive, and name it **ClusterDisk**.
 - Format the second drive, and name it **ClusterVMs**.
 - Format the third drive, and name it **Quorum**.
6. On LON-HOST1, navigate to **Disk Management**, and bring online all three iSCSI drives.

► Task 2: Configure failover clustering on both host machines

1. On LON-HOST1 and LON-HOST2, install the failover clustering feature.
2. On LON-HOST1, create a failover cluster:
 - Add **LON-HOST1** and **LON-HOST2**
 - Name the cluster **VMCluster**
 - Assign the **172.16.0.126** address
 - Deselect the option to Add all eligible storage to the cluster

► Task 3: Configure disks for the failover cluster

1. On LON-HOST1, in the Failover Cluster Manager, add all three iSCSI disks to the cluster.
2. Verify that all three iSCSI disks appear available for cluster storage.
3. Add the disk named **ClusterVMs** to **Cluster Shared Volumes**.
4. From the **VMCluster.adatum.com** node, select **More Actions**, and then configure the **Cluster Quorum Settings** to use typical settings.

Results: After completing this exercise, you will have configured a failover cluster for Hyper-V.

Exercise 3: Configuring a Highly Available Virtual Machine

Scenario

After you configuring the Hyper-V failover cluster, you want to add virtual machines as highly available resources. Additionally, you want to evaluate live migration and test storage migration.

The main tasks for this exercise are as follows:

1. Move virtual machine storage to the iSCSI target.
2. Configure the virtual machine as highly available.
3. Perform live migration for the virtual machine.
4. Perform storage migration for the virtual machine.

► Task 1: Move virtual machine storage to the iSCSI target

1. In the Failover Cluster Manager, verify that **LON-HOST1** is the owner of the **ClusterVMs** disk. If it is not, move the ClusterVMs disk to LON-HOST1.
2. On LON-HOST1, open a Windows Explorer window, and browse to **E:\Program Files\Microsoft Learning\20412\Drives\20412A-LON-CORE\Virtual Hard Disks**.
3. Move **20412A-LON-CORE.vhd** to the **C:\ClusterStorage\Volume1** location.

► Task 2: Configure the virtual machine as highly available

1. On LON-HOST1, in Failover Cluster Manager, click **Roles**, and then start the New Virtual Machine Wizard.
2. Configure a virtual machine with the following settings:
 - Cluster node: **LON-HOST1**.
 - Computer name: **TestClusterVM**

- Store the file at **C:\ClusterStorage\Volume1**
 - RAM for TestClusterVM: **1536** MB
 - Connect machine to existing virtual hard disk drive **20412A-LON-CORE.vhd**, which is located at **C:\ClusterStorage\Volume1**.
3. From the **Roles** node, start the virtual machine.
- ▶ **Task 3: Perform live migration for the virtual machine**
1. On LON-HOST1, in the Failover Cluster Manager, start **Live Migration** failover of **TestClusterVM** from LON-HOST1 to LON-HOST2.
 2. Connect to **TestClusterVM**, and ensure that you can operate the virtual machine while it is migrating to another host.
- ▶ **Task 4: Perform storage migration for the virtual machine**
1. On LON-HOST2, open Hyper-V Manager.
 2. Move **20412A-LON-SVR1-B** from its current location to **C:\LON-SVR1**.
 3. Determine whether machine is operational during move process.
 4. When the migration completes, shut down all running virtual machines.

Results: After completing this exercise, you will have configured a highly available virtual machine.

▶ **To prepare for next module**

1. Restart LON-HOST1.
2. When you are prompted with the boot menu, select **Windows Server 2008 R2**, and then press Enter.
3. Log on to the host machine as directed by your instructor.
4. Repeat steps 1-3 on **LON-HOST2**.

Module Review and Takeaways

Question: In Windows Server 2008 R2, do you have to implement CSV in order to provide high availability for virtual machines in VMM?

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Virtual machine failover fails after implementing CSV and migrating the shared storage to CSV	
A virtual machine fails over to another node in the host cluster, but loses all network connectivity	
Four hours after restarting a Hyper-V host that is a member of a host cluster, there are still no virtual machines running on the host.	

Best Practice

- Develop standard configurations before you implement highly available virtual machines. The host computers should be configured as close to identical as possible. To ensure that you have a consistent Hyper-V platform, configure standard network names, and use consistent naming standards for CSV volumes.
- Implement VMM. VMM provides a management layer on top of Hyper-V and Failover Cluster Manager that can block you from making mistakes when you manage highly available virtual machines. For example, it blocks you from creating virtual machines on storage that is inaccessible from all nodes in the cluster.

Module 7

Implementing Disaster Recovery

Contents:

Module Overview	7-1
Lesson 1: Overview of Disaster Recovery	7-2
Lesson 2: Implementing Windows Server Backup	7-7
Lesson 3: Implementing Server and Data Recovery	7-16
Lab: Implementing Windows Server Backup and Restore	7-20
Module Review and Takeaways	7-25

Module Overview

Organizations are vulnerable to losing some of their data for reasons such as unintentional deletion of critical data, file system corruption, hardware failures, malicious users, and natural disasters. Because of this, organizations must have well-defined and tested recovery strategies that will help them to bring their servers and data back to a healthy and operational state, and in the fastest time possible.

In this module, you will learn how to identify security risks for your organization. You will also learn about disaster recovery, and disaster recovery requirements. You will also learn how to plan backup across your organization, and what steps you can take to recover data.

Objectives

After completing this module, you will be able to:

- Describe disaster recovery concepts.
- Implement Windows Server Backup.
- Implement server and data recovery.

Lesson 1

Overview of Disaster Recovery

Disaster recovery is a methodology that describes all the steps that you need to perform once a disaster has occurred, to bring data, services and servers back to an operational state. An effective disaster recovery plan addresses the organization's needs without providing an unnecessary level of coverage. While absolute protection may seem desirable, it is unlikely to be economically feasible. In creating a disaster recovery plan, you need to balance the cost to the organization of a particular disaster, with the cost to the organization of protection from that disaster.

Lesson Objectives

After completing this lesson, you will be able to:

- Identify disaster recovery requirements.
- Describe service level agreements.
- Describe enterprise disaster recovery strategies.
- Describe disaster mitigation strategies.
- Describe best practices for implementing a disaster recovery.

Identifying Disaster Recovery Requirements

Before developing a disaster recovery strategy, organizations must identify their disaster recovery requirements to ensure that they will provide appropriate protection for critical resources.

The following is a high-level list of steps that you can use to identify disaster recovery requirements:

1. Define organization critical resources. These resources include data, services, and the servers upon which the data and services run.
2. Identify risks associated with those critical resources. For example, data can be accidentally or intentionally deleted, and a hard drive or storage controller where data is stored might fail. Additionally, services that use critical data might fail due to many reasons such as network problems, and servers might fail because of hardware failures. Major power outages could also cause entire sites to shut down.
3. Identify the time needed to perform the recovery. Based on their business requirements, organizations should decide how much time is acceptable for recovering critical resources. Scenarios may vary from minutes to hours, or even a day.
4. Develop a recovery strategy. Based on the previous steps, organizations will define a service level agreement that will include information such as service levels and service hours. Organizations should develop a disaster recovery strategy that will help them minimize the risks, and at the same time, recover their critical resources within the minimum time acceptable for their business requirements.

Identify your disaster recovery options by:

1. Defining organization critical resources
2. Identifying risks associated with those critical resources
3. Identifying the time needed to complete the recovery
4. Developing a recovery strategy



Note: Organization will have differing disaster recovery requirements based on their business requirements and goals. Disaster recovery requirements should not be static, but they

should be evaluated and updated on a regular basis—for example once every few months. It is also important that administrators test the disaster recovery strategies on a regular basis. The testing should be performed in an isolated, non-production environment by using a copy of the production data.

What Are Service Level Agreements?

A *service level agreement* (SLA) is a document that describes the responsibilities of the IT department or IT service provider, with respect to a specific set of objectives. In terms of data protection SLAs, these agreements usually specify precisely which parts of the IT infrastructure and data will be protected, and how quickly they will return to service after a failure.

In some organizations, SLAs are formalized, and the performance of the IT department is measured against the objectives that are spelled out in the SLA. These metrics form part of the IT department's performance evaluation, and have a direct influence on items such as budgets and salaries. For managed services or cloud providers, SLAs are critical for billing purposes. In other organizations, SLAs are guidelines and are less formalized. The key to developing an SLA is that it needs to be realistic and achievable, rather than an unachievable standard, which may be impossible to reach.

Some of the elements of an SLA include:

- **Hours of operation.** Hours of operation defines how much time the data and services are available to users, and how much planned downtime there will be due to system maintenance.
- **Service availability.** Service availability is defined as a percentage of time per year that data and services will be available to users. For example, a service availability of 99.9 percent per year means that data and services will have unplanned downtime not more than 0.1 percent per year, or 8.75 hours per year on a 24 hours a day, seven days a week basis.
- **Recovery point objective (RPO).** An RPO sets a limit on how much data can be lost due to failure, measured as a unit of time. For example, if an organization sets an RPO of six hours, it would be necessary to take a backup every six hours, or to create a replication copy on different locations at six-hour intervals. In the event of a failure, it would be necessary to go back to the most recent backup, which, in the worst-case scenario, assuming that the failure occurred just before (or during) the next backup, would be six hours ago.

You can configure backup software to take backups every hour, offering a theoretical RPO of 60 minutes. When calculating RPO, it is also important to take into account the time it takes to perform the backup. For example, suppose it takes 15 minutes to perform a backup and you back up every hour. If a failure occurs during the backup process, your best possible RPO will be 1 hour and 15 minutes. A realistic RPO must always balance the desired recovery time with the realities of the network infrastructure. You should not aim for an RPO of 2 hours when a backup itself takes three hours to complete.

The RPO also depends on the backup software technology. For example, when you use the snapshot feature in Windows Server Backup, or other backup software that uses volume shadow copy service (VSS), you are backing up to the point in time when the backup was started.

- **Recovery time objective (RTO).** An RTO is the amount of time it takes to recover from failure. The RTO will vary depending on the type of failure. The loss of a motherboard on a critical server will have a

SLAs define responsibilities of the service provider

SLA components include:

- Hours of operation
- Service availability
- RPO and RTO
- Retention objectives
- System performance

different RTO than the loss of a disk on a critical server, because one of these components takes significantly longer to replace than the other.

- **Retention objectives.** Retention is a measure of the length of time you need to store backed-up data. For example, you may need to recover data quickly from up to a month ago, but need to store data in some form for several years. The speed at which you agree to recover data in your SLA will depend on the age of the data, with some data being quickly recoverable and other data needing to be recovered from the archives.
- **System performance.** Although not directly related to disaster recovery, system performance is also an important component of SLAs, because applications that are included in an SLA should be available, and they should also have acceptable response times to users' requests. If the system performance is slow, then business requirements will not be met.



Note: Each organization's data protection SLA depends on the components that are important to the organization.

Overview of Enterprise Disaster Recovery Strategies

When planning for backup for your enterprise, you need to develop strategies for recovering data, services, servers, and sites, and you need to make some provision for offsite backup.

Data Recovery Strategies

Data is the most commonly recovered category in an enterprise environment. This is because it is more likely that users will delete files accidentally, than it is for server hardware to fail or for applications to cause data corruption. Therefore, in developing an enterprise disaster recovery strategy, take into account small disasters, such as data deletion, in addition to big disasters, such as server or site failure.

You need strategies for recovering:

- Data
- Services
- Servers
- Sites
- Offsite backups

When considering data recovery strategies, backup is not the only technology for data recovery. You can address many file and folder recovery scenarios by implementing previous versions of file functionality on file shares. You can also replicate data in different physical locations, or to a public or private cloud. You could also use Microsoft® System Center 2012 - Data Protection Manager.

Service Recovery Strategies

The functionality of the network depends on the availability of certain critical network services. Although well-designed networks build redundancy into core services such as Domain Name System (DNS) and Active Directory® Domain Services (AD DS), even those services might have issues, such as when a major fault is replicated that requires a restore from backup. In addition, an enterprise backup solution must ensure that services such as Dynamic Host Configuration Protocol (DHCP) and Active Directory Certificate Services (AD CS), and important resources such as file shares can be restored in a timely and up-to-date manner.

Full Server Recovery Strategies

Developing a full-server recovery strategy involves determining which servers that you need to be able to recover, the RPO for critical servers, and the RTO for critical servers. Suppose that you have a site with two computers functioning as domain controllers. When developing your backup strategy, should you aim to

have both servers capable of full server recovery with a 15 minute RPO? Or is it only necessary for one server to be recovered quickly if it fails, given that either server will be able to provide the same network service and ensure business continuity?

In developing the full server recovery component of your organization’s enterprise backup plan, determine which servers are required to ensure business continuity, and ensure that they are regularly backed up.

Site Recovery Strategies

Most larger organizations have branch office sites. While it might be desirable to back up all the computers at those locations, it may not be economically feasible to do so. Developing a site recovery strategy involves determining which data, services, and servers at a specific site must be recoverable to ensure business continuity.

Offsite Backup Strategies

Many organizations that do not store offsite backups do not recover from a primary site disaster. If your organization’s head office site has a fire, is subject to a once-in-a-100-year flood, a cyclone, or a tornado, it will not matter what backups strategies you have in place if all those backups are stored at the location that was destroyed by the disaster.

A comprehensive enterprise data protection strategy involves moving backed-up data to a safe offsite location so that you can recover it no matter what kind of disaster occurs. This does not need to happen every day. The RPO for recovery at the offsite location—often called the *disaster recovery site*—is usually different from the RPO at the primary site.

Disaster Mitigation Strategies

No matter how prepared organizations are, they cannot prevent disasters from occurring. Therefore, organizations must also develop mitigation strategies that will minimize the impact of an unexpected loss of data, server, service, or sites. To prepare mitigation strategies, organizations must create risk assessments that analyze all possible disaster scenarios, and how to mitigate each of those scenarios.

The following table lists some of the risks associated with data or services loss, and the appropriate mitigation strategies.

Risk of disaster	Mitigation strategy
The media where a copy of the backup data is store becomes corrupted	Have at least two copies of your backup data
An administrator has accidentally deleted an OU that contains many user and computer objects	Protect Ous from accidental deletion, especially after migrations
A server in a branch office where important files are located has failed	Use DFS-R to replicate files from branch offices to central datacenters
The virtualization infrastructure where business servers are located is unavailable	Avoid deploying all critical servers, such as domain controllers, on the same virtual infrastructure
A major outage in a data center has occurred	Deploy a secondary data center that will contain replicas of most of he critical servers in your primary data center

Risk of disaster	Mitigation strategy
The media where a copy of the backup data is located becomes corrupted.	Have at least two copies of your backup data, and validate your backups on a regular basis.
An administrator has accidentally deleted an organizational unit (OU) that contains many user and computer objects.	Protect OUs from accidental deletion, especially after migrations.
A file server in a branch office where important files are located has failed.	Use Distributed File System Replication (DFS-R) to replicate files from branch offices to central data centers.

MIGT USE ONLY STUDENT USE PROHIBITED

Risk of disaster	Mitigation strategy
The virtualization infrastructure where business servers are located is unavailable.	Avoid deploying all critical servers, such as domain controllers, on the same virtual infrastructure.
A major outage in a data center has occurred.	Deploy a secondary data center that will contain replicas of the critical servers in your primary data center.

Best Practices for Implementing a Disaster Recovery

When implementing a disaster recovery strategy, organizations should follow these best practices:

- Perform a risk assessment plan first. This will help you identify all of the risks associated with the availability of your organization data, servers, services, and sites.
- Discuss the risks you evaluated with your business managers, and together decide which resources should be protected with the disaster recovery plan, and which resources should be protected with disaster mitigation, and at which level. The higher the requirements for disaster recovery are, the more expensive they are. You also want to have a low-level disaster recovery plan for resources that are protected with disaster mitigation.
- Each organization should have its own disaster recovery plan.
- Document in detail all of the steps that should be performed in a disaster scenario.
- Test your disaster recovery plan on a regular basis in an isolated, non-production environment.
- Evaluate your disaster recovery plan on a regular basis, and update your disaster recovery plan based on your evaluation.

To implement a disaster recovery strategy, you should:

- Perform a risk assessment plan first
- Discuss the risks you evaluated with your business managers, and create a disaster recovery strategy and disaster mitigation strategy
- Every organization should have its own disaster recovery plan
- Document all steps that should be performed in a disaster scenario
- Test your disaster recovery plan on regular basis, in an isolated, non-production environment.
- Evaluate your disaster recovery plan on a regular basis, and update your disaster recovery plan depending on your evaluation outcome

Lesson 2

Implementing Windows Server Backup

To protect critical data, every organization must perform regular backups. Having a well-defined and tested backup strategy ensures that companies can restore data if unexpected failures or data loss occur. This lesson describes the Windows Server Backup feature in Windows Server® 2012 and the Microsoft Online Backup Service for Windows Server 2012.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe data and service information that needs to be backed up in a Windows Server environment.
- Describe the backup types.
- Describe backup technologies.
- Describe backup capacity.
- Describe backup security.
- Describe Windows Server Backup.
- Explain how to configure a scheduled backup using Windows Server Backup.
- Describe the Windows Server 2012 online backup solution.
- Describe the considerations for an enterprise backup solution.
- Summarize the features available with System Center 2012 – Data Protection Manager.

What Needs to Be Backed Up?

When planning backups across your organization, ensure that you protect resources that are considered mission critical. Consider the following:

- Critical resources
- Backup verification
- Backup security
- Compliance and regulatory requirements

Determining Critical Resources to Back Up

In an ideal scenario, you would back up everything and instantly restore data as it existed at a particular point in time from any point in the last several years. In reality, such a backup strategy would produce expensive cost of ownership. Therefore, the first step in planning backup across the enterprise is to determine what exactly needs to be backed up.

For example, should you back up every domain controller in the domain, given that Active Directory information will be replicated back to a replacement domain controller as soon as it is promoted? Is it necessary to back up every file server in all file shares if every file is replicated to multiple servers through a distributed file system?

When planning your backup strategy, ensure that you:

- Determine the critical resources
- Verify your backups
- Confirm that backups are secure
- Ensure that compliance and regulatory responsibilities are met

You also need to distinguish between technical reasons and regulatory reasons for backing up data. Due to legal requirements, you may need to be able to provide your business with business-critical data for the past ten years or even longer.

To determining what to back up, consider the following:

- If the data is only stored in one place, ensure that it is backed up.
- If data is replicated, it may not be necessary to back up each replica. However, you must back up at least one location to ensure that the backup can be restored.
- Is the server or data a mission-critical component?
- If this server or disk failed, or if this data became corrupted, what steps would need to be taken to recover it?

Many organizations ensure the availability of critical services and data through redundancy. For example, Exchange Server 2010 provides continuous replication of mailbox databases to other servers through a technology called Database Availability Groups (DAGs). While DAGs do not mean that an organization should not back up its Exchange Server 2010 Mailbox servers, it does change how an organization should think about backing up its Mailbox servers or centralizing its backup strategies.

Verifying Your Backups

Performing a backup, and ensuring that the backup contains everything that you need, are two different tasks. You need to have a method for verifying that each backup has completed successfully. You also need to know when backups have failed. At a minimum, this will mean checking the logs on each server to determine whether a failure has occurred. If you have configured backups to occur on each server every six hours, how often should you check the logs? A better solution is to employ an alert mechanism—such as that which is available in System Center 2012 - Operations Manager—to alert you in the event that a backup fails. The point is to avoid discovering that your backups for a particular server have failed just when you need those backups to perform a recovery.

One way of verifying backups is to perform regular testing of the recovery procedures, in which you simulate a particular failure. This allows you to verify not only the integrity of the data that you are using to perform a recovery, but also that the recovery procedures that you have in place effectively resolve the failure. It is better to discover that you need to add steps to your recovery procedure during a test, rather than during an actual failure.

Confirming That Backups Are Secure

By definition, a good set of backups contains all of your organization's critical data. This data needs to be protected from unauthorized access. Although data might be protected by permissions and access controls while it is hosted on servers in a production environment, anyone who has access to the media that hosts that backup data can restore it. For example, some products, such as Windows Server Backup, do not allow administrators to encrypt backup data. This means that physical security is the only way that you can ensure that critical data does not fall into the hands of unauthorized users.

When developing an enterprise backup strategy, ensure that backup data is stored in a secure location.

You might also consider using backup software that allows you to split the backup and restore roles so that users who have permissions to back up data do not have permissions to restore that data, and users who have permissions to restore data do not have permissions to back it up.

Ensuring That Compliance and Regulatory Responsibilities Are Met

Systems administrators should be aware of what the organization's regulatory and compliance responsibilities are with respect to the archiving of data. For example, some jurisdictions require that business-relevant email message data be retained for a period of up to seven years. Unfortunately, regulatory requirements vary from country to country, and even from state to state. When developing

your organization's data protection strategy, you should schedule a meeting with your organization's legal team to determine precisely which data needs to be stored, and for how long.

Backup Types

In Windows Server 2012, you can perform the following types of backups:

- **Full backup.** A full backup is a block-level replica of all blocks on all the server's volumes. Rather than copying files and folders to backup media, the underlying blocks are copied across to the backup media.
- **Incremental backup.** An incremental backup is a copy of only those blocks that have changed since the last full or incremental backup. During an incremental backup, these blocks are copied across to the backup media.

When this process completes, the blocks are then marked as backed up. During recovery, the original set of blocks is restored. Then, each set of incremental blocks are applied, bringing the recovered data back to the appropriate state in a consistent manner.

- A full backup is a block-level replica of all blocks on all the server's volumes
- An incremental backup is a copy of only those blocks that have changed since the last full or incremental backup

Backup Technologies

Most backup products in use today use the VSS infrastructure that is present in Windows Server 2012. Some older applications, however, use streaming backup. It may be necessary to support such applications in complex heterogeneous environments.

One of the challenges of performing backups is ensuring the consistency of the data that you are backing up. Backups do not occur instantly, they can take seconds, minutes, or hours.

Unfortunately, servers are not static and the state of a server at the beginning of a backup might not be the same state that the server is in when the backup completes. If you do not take consistency account, this can cause problems during restoration because the configuration of the server may have changed during the backup.

VSS

VSS—a technology that Microsoft included with Windows Server 2003 R2, and which is present in all newer server operating systems—solves the consistency problem at the disk-block level by creating what is known as a shadow copy. A *shadow copy* is a collection of blocks on a volume that is frozen at a specific point in time. Changes can still be made to the disk, but when a backup occurs, the collection of frozen blocks are backed up, which means that any changes that might have occurred since the freeze are not backed up.

Creating a shadow copy tells the operating system first to put all files, such as DHCP databases and Active Directory database files, in a consistent state for a moment. Then the current state of the file system is

- The VSS backup technology solves data consistency issues by creating shadow copies
- You can also use streaming backups for older applications that are not VSS-aware

recorded at that specific point in time. After VSS creates the shadow copy, all write accesses that would overwrite data, store the previous data blocks first. Therefore, a shadow copy is small in the beginning, and it grows over the time as data changes. By default, the operating system is configured to reserve 12 percent of the volume for VSS data, and VSS automatically deletes older snapshots when this limit is reached. You can change this default value, and you can change the default location of the VSS data. This ensures that the backup has a snapshot of the system in a consistent state, no matter how long it actually takes to write the backup data to the backup storage device.

Streaming Backup

Streaming backup is often used by older applications that do not use VSS. You back up applications that are not VSS-aware by using a method known as a *streaming backup*. In contrast to VSS where the operating system ensures that data is kept in a consistent state and at a current point in time, when you use streaming backup, the application or the data protection application is responsible for ensuring that the data remains in a consistent state. In addition, after streaming backup completes, some files have the state they had in the beginning of the backup, while other files have the state of the end of the backup window.

Planning Backup Capacity

When you develop an enterprise recovery strategy, you need to determine how much storage capacity your organization will require for backups. The following factors affect the amount of space that is required to store backup data:

- Space requirements for a full backup
- Space requirements for an incremental backup
- Amount of time required to back up
- Backup frequency
- Backup retention

When planning for backup capacity, consider the following:

- Space requirements for a full backup
- Space requirements for an incremental backup
- Amount of time required to back up
- Backup frequency
- Backup retention

Full Backup Requirements

To calculate the space required for a full backup, determine how much space from all volumes you will need to back up. If the server has a dedicated drive for backups, you would not perform a backup on that drive.

With products that perform image-based backups, such as Windows Server Backup, this data is not compressed. On some types of servers, notably file servers, the amount of space required for a full backup grows over time. You can lessen this tendency by using file expiration policies such as those found in File Server Resource Manager (FSRM).

Incremental Backup Requirements

An incremental backup on Windows Server Backup stores all of the hard disk blocks that have changed since the last full or incremental backup. Incremental backups are substantially faster than full backups and require less space. The downside of incremental backups is that they can require greater recovery time.

Amount of Time Required to Back up

The amount of time required to write data from the server being backed up to the backup storage device can have an impact on projected RPO, because it is not recommended to begin a new backup operation prior to the completion of the current one.

Backup Frequency

Backup frequency is a measure of how often backups are taken. With incremental block-level backups, no substantial difference will exist between the amount of data written over the sum of four 30-minute sessions and one 2-hour incremental sessions on the same server. This is because over the two hours, the same number of blocks will have changed on the server as the four 30-minute sessions. However, the four 30-minute sessions have broken it up into smaller parts. When backups occur more frequently, they reduce the time required to perform the backup by splitting it into smaller parts. The overall total will be about the same.

Backup Retention

When attempting to determine the required backup capacity, you should determine precisely how long you need to retain backup data. For example, if you need to be able to recover to any backup point in the last 28 days, and you have recovery points generated every hour, you will need more space than if you have recovery points generated once a day and you only need to restore from the last 14 days.

Planning Backup Security

When planning your backup security, consider the following:

- Backups contain all organizational data. By nature, backups will contain all the data necessary to ensure your organization's continued ability to function in the event of failure. Because this data is likely to contain sensitive information, you should protect it with the same level of diligence as it is protected with when hosted on the server.
- Access to backup media means access to all data. If feasible, use administrative role separation to ensure that the users who back up the data are not the users who can restore it. In high security environments, ensure that backup and restore operations are properly audited so that you can track backups, and restore function activity.
- Windows Server Backup does not encrypt backups. Windows Server Backup writes backups in VHD format. This means that anyone who has access to Windows® 8 or Windows Server 2012 can mount those backups as volumes, and then extract data from them. An even more sophisticated attack might include booting into the backup VHD to impersonate the backed up system on the organizational network.
- Keep backup media in a secure location. At a minimum, backups should be kept locked up in a secure location. If your organization is backing up to disk drives that are attached to servers by USB cable, ensure that those disk drives are locked in place, even if they are located in a secure server room, and even if your organization's server room has a security camera.

When planning your backup security, consider the following:

- Backups contain all organizational data
- Access to backup media means access to all data
- Windows Server Backup does not encrypt backups
- Keep backup media in a secure location

What Is Windows Server Backup?

The Windows Server Backup feature in Windows Server 2012 consists of a Microsoft Management Console (MMC) snap-in, the command `wbadmin`, and Windows PowerShell® commands. You can use wizards in the Windows Server Backup feature to guide you through running backups and recoveries.

You can use Windows Server Backup 2012 to back up:

- Full server (all volumes).
- Selected volumes.
- Select specific items for backup, such as specific folders or the system state.

You can use Windows Server Backup to:

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

In addition, Windows Server Backup 2012 allows you to:

- Perform a bare-metal restore. A bare-metal backup contains at least all critical volumes, and allows you to restore without first installing an operating system. You do this by using the product media on a DVD or USB key, and the Windows Recovery Environment (Windows RE). You can use this backup type together with the Windows RE to recover from a hard disk failure, or if you have to recover the whole computer image to new hardware.
- Use system state. The backup contains all information to roll back a server to a specific point in time. However, you need an operating system installed prior to recovering the system state.
- Recover individual files and folders or volumes. The **Individual files and folders** option enables you to select to back up and restore specific files, folders, or volumes, or you can add specific files, folders, or volumes to the backup when you use an option such as critical volume or system state.
- Exclude selected files or file types. For example, you can exclude temporary files from the backup.
- Select from more storage locations. You can store backups on remote shares or non-dedicated volumes.
- Use the Microsoft Online Backup Service. The Microsoft Online Backup Service is a cloud-based backup solution for Windows Server 2012 that enables files and folders to be backed up and recovered from the public or private cloud to provide off-site backup.

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

Demonstration: Configuring a Scheduled Backup

In this demonstration, you will see how to configure Windows Server 2012 to perform a scheduled backup of specific folders, with a filter to exclude specific file types.

Demonstration Steps

1. On LON-SVR1, start Windows Server Backup.
2. Configure the backup schedule with the following options:
 - Backup Configuration: Custom and **C:\HR Data** folder is backed up, with the exception of **C:\HR Data\Old HR file.txt**

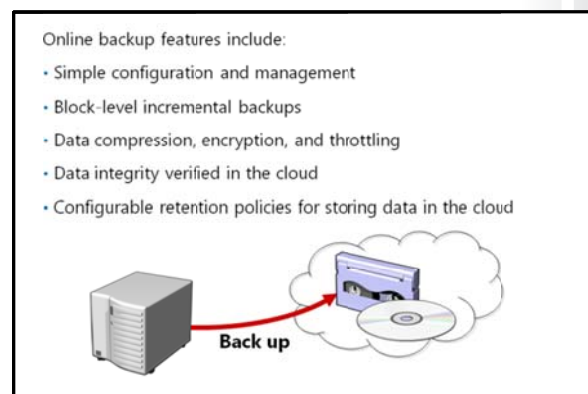
- Backup Time: **Once a day, 1:00 AM**
 - Destination Type: **Back up to a shared network folder**
 - Remote Shared Folder: **\\LON-DC1\Backup:**
 - Register Backup Schedule: Username: **Administrator**
 - Password: **Pa\$\$w0rd**
3. Run the Backup Once Wizard using the scheduled backup options.
 4. Close Windows Server Backup.

What Is Online Backup?

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

Microsoft Online Backup Service is built on the Windows Azure® platform, and uses Windows Azure blob storage for storing customer data.

Windows Server 2012 uses the downloadable Microsoft Online Backup Service Agent to transfer file and folder data securely to the Microsoft Online Backup Service. After you install the Microsoft Online Backup Service Agent, the agent integrates its functionality through the Windows Server Backup interface.



Key Features

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

- Simple configuration and management. Integration with the Windows Server Backup tool provides a seamless backup and recovery experience to a local disk, or to a cloud platform. Other features include:
 - Simple user interface to configure and monitor backups.
 - Integrated recovery experience to recover files and folders from local disk or from a cloud platform.
 - Easy data recoverability for data that was backed up onto any server of your choice.
 - Scripting capability that is provided by the Windows PowerShell command-line interface.
- Block-level incremental backups. The Microsoft Online Backup Agent performs incremental backups by tracking file and block-level changes, and only transferring the changed blocks, which reduces the storage and bandwidth usage. Different point-in-time versions of the backups use storage efficiently by only storing the changed blocks between these versions.
- Data compression, encryption, and throttling. The Microsoft Online Backup Service Agent ensures that data is compressed and encrypted on the server before it is sent to the Microsoft Online Backup Service on the network. Therefore, the Microsoft Online Backup Service only stores encrypted data in cloud storage. The encryption passphrase is not available to the Microsoft Online Backup Service, and

therefore, the data is never decrypted in the cloud. In addition, users can set up throttling and configure how the Microsoft Online Backup Service uses the network bandwidth when backing up or restoring information.

- Data integrity verified in the cloud. In addition to the secure backups, the backed up data is also checked automatically for integrity after the backup completes. Therefore, any corruptions that may arise because of data transfer can be easily identified. These corruptions are fixed automatically in the next backup.
- Configurable retention policies for storing data in the cloud. The Microsoft Online Backup Service accepts and implements retention policies to recycle backups that exceed the desired retention range, thereby meeting business policies and managing backup costs.



Reference Links: You can find out more about Windows Azure at: <http://www.windowsazure.com/en-us/home/features/storage/>

Considerations for an Enterprise Backup Solution

Windows Server Backup is a single-server backup solution. When planning backup for an enterprise, consider the following points:

- Maximum amount of data lost. What is the theoretical RPO of the product? Products that offer restoration closer to the point of the failure are likely to cost more than products that offer 15 minute or 30 minute RPOs. You need to determine your organization's needs. Does your organization need to be able to recover to the last SQL Server transaction, or is a 15-minute recovery window an acceptable compromise?
- How quick is RTO recovery? How long does it take to go from failure to restored functionality? Being able to restore to the last SQL Server transaction is the optimal solution, but if it takes two days to recover to that point, the solution is not looking as good.
- Does the solution provide centralized backup? Does the product allow you to centralize your backup solution on one server, or must backups be performed directly on each server in the organization?
- Is the solution supported by vendors? Some vendors use undocumented application programming interfaces (APIs) to back up and recover specific products, or to back up files without ensuring that the service is at a consistent state.
- Is the backup solution compatible with your applications? For example, a new update to a product makes the backup solution incompatible. Check with the application vendor to determine whether the enterprise backup solution is supported.
- Recovery point capacity. Determine what the recovery point capacity of the product is. How many restore points does the enterprise data protection solution offer, and is this adequate for your organization's needs?

Considerations for an enterprise backup solution are:

- What is the theoretical RPO of the product?
- How quick is RTO recovery?
- Does the solution provide centralized backup?
- Is the solution supported by vendors?
- What is the recovery point capacity?

What Is Data Protection Manager?

DPM is a Microsoft enterprise data protection and recovery product with the following features:

- Backup centralization. DPM uses a client/server architecture, where the client software is installed on all the computers that are to be backed up. Those clients stream backup data to the DPM server. This allows each DPM server to support entire small to medium-sized organizations. You can also manage multiple DPM servers from one centralized DPM console.
- 15-minute RPO. DPM allows 15-minute snapshots of supported products. This includes most of the Microsoft enterprise suite of products, including Windows Server with its roles and services, Exchange Server, Hyper-V®, and SQL Server.
- Supports Microsoft workloads. DPM was designed specifically by Microsoft to support Microsoft applications such as Exchange Server, SQL Server, and Hyper-V. However, DPM has not been specifically designed to support non-Microsoft server applications that do not have consistent states on disk, or that do not support VSS.
- Disk-based backup. DPM can perform scheduled backups to disk arrays and storage area networks (SANs). You can also configure DPM to export specific backup data to tape for retention and compliance related tasks.
- Remote site backup. DPM uses an architecture that allows it to back up clients that are located in remote sites. This means that a DPM server that is located in a head office site can perform backups of servers and clients that are located across wide area network (WAN) links.
- Supports Backup to Cloud strategies. DPM supports backup of DPM servers. This means that a DPM server at a cloud-based hosting facility can be used to back up the contents of a head office DPM server. For disaster redundancy, you can also configure DPM servers to back up each other.

DPM:

- Allows you to centralize backups
- Offers 15-minute snapshots of servers and clients
- Can store backup data on SANs and export to tape
- Can back up remote sites
- Can be used as part of a backup-to-cloud strategy
- Supports Microsoft products

Lesson 3

Implementing Server and Data Recovery

Recovering servers and data requires well-defined and documented procedures that administrators can follow when failures occur. The recovery process also requires knowledge of the backup and restore hardware and software, such as DPM, and tape library devices.

This lesson describes how to restore data and servers by using the Windows Server Backup feature in Windows Server 2012, and Microsoft Online Backup Service in Windows Server 2012.

Options for Server Recovery

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

- **Files and folders.** You can back up individual files or folders as long as the backup is on separate volume or in a remote shared folder.
- **Applications and data.** You can recover applications and data if the application has a VSS writer, and is registered with Windows Server Backup.
- **Volumes.** Restoring a volume always restores all the contents of the volume. When you choose to restore a volume, you cannot restore individual files or folders.
- **Operating system.** You can recover the operating system through Windows RE, the product DVD, or a USB flash drive.
- **Full server.** You can recover the full server through Windows RE.
- **System state.** *System state* creates a point-in-time backup that you can use to restore a server to a previous working state.

The options for server recovery include:

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

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

- **Recovery Destination.** Under Recovery Destination, you can select any one of the following options:
 - **Original location.** The original location restores the data to the location to which it was backed up originally.
 - **Another location.** Another location restores the data to a different location.
- **Conflict Resolution.** Restoring data from a backup frequently conflicts with existing versions of the data. Conflict resolution allows you to determine how to handle those conflicts. When these conflicts occur, you have the following options:
 - **Create copies and retain both versions.**
 - **Overwrite existing version with recovered version.**
 - **Do not recover items if they already exist in the recovery location.**
- **Security Settings.** Use this option to restore permissions to the data that is being recovered.

Options for Server Restore

You perform server restore by starting the computer from the Windows Server 2012 installation media, selecting the computer repair option, and then selecting the full server restore option. Alternatively, you can use the installation media on a USB flash drive, or using Windows RE.

When you perform full server restore, consider the following:

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

The server restore locations include:

- Original host: bare-metal restore
- New host: bare -metal restore
- Hyper-V: virtual machine restore
- Alternate boot-to-VHD

Options for Data Recovery

Data is the most frequently recovered component of an IT infrastructure. This is due to users accidentally deleting data, and needing you to recover it. There are several strategies that you can pursue when you are developing a data recovery procedure. You can:

- Allow users to recover their own data.
- Perform a recovery to an alternative location.
- Perform a recovery to the original location.
- Perform a full volume recovery.

The four options for recovering data include:

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

Users Recover Their Own Data

The most common form of data recovery performed by IT departments is the recovery of files and folders that users have deleted, lost, or in some way corrupted. The **Previous Versions of Files** functionality that was introduced in Windows Server 2003, (which you can also enable on all computers running Windows Server 2012,) lets users recover their own files using the file or folder properties right from their workstation. After end-users are trained how to do this, the IT department spends less time recovering user data, which allows them to focus on more valuable tasks.

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

Recover Data to an Alternative Location

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

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

Recover Data to the Original Location

During some types of failures, such as data corruption or deletion, you will have to restore data to the original location. This is the case when applications or users who access the data are preconfigured with information about where the data is located.

Recover a Volume

If a disk fails, the quickest way to recover the data could be to perform a volume recovery, instead of a selective recovery of files and folders. When you perform a volume recovery, you must check whether any shared folders are configured for the disks, and whether the quotas and FSRM management policies are still in effect.



Note: During the restore process, you should copy event logs before you start the restore process. If you overwrite the event log files—for example with a system recovery—you will be not able to read event log information that occurred before the restore started. That event log data could lead you to information about what caused the issue.

Demonstration: Using Windows Server Backup to Restore a Folder

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

Demonstration Steps

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

Restoring with an Online Backup Solution

You can use the Microsoft Online Backup Service to back up only Windows Server 2012 servers. However, you do not have to restore data on to the same server from which you backed it up.

You can recover files and folders by using both Microsoft Online Backup MMC in Server Manager, or by using the Windows PowerShell® command-line interface. To use the Microsoft Online Backup MMC, perform the following steps:

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

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



Lab: Implementing Windows Server Backup and Restore

Scenario

Much of the data that is stored on the A. Datum Corporation's network is extremely valuable to the organization. Losing this data would be a significant loss to the organization. Additionally, many of the servers that are running on the network provide extremely valuable services for the organization, which means that losing these servers for a significant period of time would also result in losses to the organization. Because of the significance of the data and services, it is critical that they can be restored in the event of disaster.

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

As one of the senior network administrators at A. Datum, you are responsible for planning and implementing a disaster recovery solution that will ensure that critical data and services can be recovered in the event of any type of failure. You need to implement a backup and restore process that can recover lost data and services.

Objectives

- Back up data on a Windows Server 2012 server.
- Restore files using Windows Server Backup.
- Implement Microsoft Online Backup and Restore.

Lab Setup

20412A-LON-DC1

20412A-LON-SVR1

MSL-TMG1

Estimated time: **60 minutes**

Virtual Machine(s)	20412A-LON-DC1 20412A-LON-SVR1 MSL-TMG1
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 **20412A-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:
 - o User name: **Adatum\Administrator**
 - o Password: **Pa\$\$w0rd**
5. Repeat steps 2-4 for **20412A-LON-SVR1**.
6. Repeat step 2 for **MSL-TMG1**.

Exercise 1: Backing Up Data on a Windows Server 2012 Server

Scenario

The LON-SVR1 server contains financial data that must be backed up on a regular basis. This data is critical to the organization. You decided to use Windows Server Backup to back up critical data. You will to install this feature and configure scheduled backups.

The main tasks for this exercise are as follows:

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

► Task 1: Install Windows Server Backup

1. Switch to LON-SVR1.
2. From Server Manager, install the Windows Server Backup feature. Accept the default values on the Add Roles and Features Wizard.

► Task 2: Configure a scheduled backup

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

► Task 3: Complete an on-demand backup

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

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

Exercise 2: Restoring Files Using Windows Server Backup

Scenario

To ensure that the financial data can be restored, you must validate the procedure for restoring the data to an alternate location.

The main tasks for this exercise are as follows:

1. Delete a file from the server.
2. Restore a file from backup.

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

- On LON-SVR1, open Windows Explorer and then delete the **C:\Financial Data** folder.

► **Task 2: Restore a file from backup**

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

Results: After completing this exercise, you will have tested and validated the procedure for restoring a file from backup

Exercise 3: Implementing Microsoft Online Backup and Restore

Scenario

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

The main tasks for this exercise are as follows:

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

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

1. On LON-SVR1, in drive E, locate the installation file of the Microsoft Online Backup Agent, **OBSInstaller.exe**.
2. Start installing Microsoft Online Backup Agent by double-clicking the installation file **OBSInstaller.exe**.
3. Complete the setup by specifying the following information:
 - Installation Folder: **C:\Program Files**
 - Cache Location: **C:\Program Files\Microsoft Online Backup Service Agent**
 - Microsoft Update Opt-In: **I don't want to use Microsoft Update**
4. Verify the installation and ensure that you receive the following message: Microsoft Online Backup Service Agent installation has completed successfully.

5. Clear the **Check for newer updates** check box, and then click **Finish**.
6. On the **Start** screen, verify the installation by clicking **Microsoft Online Backup Service** and **Microsoft Online Backup Service Shell**.

► **Task 2: Register the server with Microsoft Online Backup Service**

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

1. In the Server Manager window, rename LON-SVR1 as *YOURCITYNAME-YOURNAME*, and then restart *YOURCITYNAME-YOURNAME*.
2. Wait until *YOURCITYNAME-YOURNAME* has restarted, and then log on as **Adatum\Administrator** with password **Pa\$\$w0rd**.
3. In the Microsoft Online Backup Service console, register LON-SVR1 by specifying the following information:
 - Username: holuser@onlinebackupservice.onmicrosoft.com
 - Password: Pa\$\$w0rd



Note: In a real-life scenario, you would type the username and password of your Microsoft Online Backup Service subscription account.

- Enter passphrase: **Pa\$\$w0rdPa\$\$w0rd** Confirm passphrase: **Pa\$\$w0rdPa\$\$w0rd**
4. Verify that you receive the following message: Microsoft Online Backup Service is now available for this server.

► **Task 3: Configure an online backup and start a backup**

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

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

1. On LON-SVR1, open Windows Explorer and delete **C:\Financial Data**.
2. Switch to the Microsoft Online Backup Service console.
3. Restore files and folders by using the **Recover Data** option, and specify the following information:
 - Identify the server on which the backup was originally created: **This server**
 - Select Recovery Mode: **Browse for files**
 - Select Volume and Date: **C:** and **date and time of the latest backup**
 - Select Items to Recover: **C:\Financial Data**
 - Specify Recovery Options: **Original location** and **Create copies so that you have both versions**
4. In Windows Explorer, expand drive **C:**, and ensure that the **Financial Data** folder is restored to drive C.

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

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

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

► **To prepare for the next module**

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20412A-LON-DC1**, and then click **Revert**.
3. In the Revert Virtual Machine dialog box, click Revert.
4. Repeat steps 2 and 3 for **20412A-LON-SVR1**, and **MSL-TMG1**.

Module Review and Takeaways

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

Question: How frequently should you perform backup on critical data?

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
The server has suffered a major failure on its components.	

Real-world Issues and Scenarios

If a failure were to occur, your organization needs information about which data to back up, how frequently to back up different types of data and technologies, where to store backed up data (onsite or in the cloud), and how fast it can restore backed-up data. How would you improve your organization's ability to restore data efficiently when it is necessary?

Answer: Your company should develop backup and restore strategies based on multiple parameters, such as business continuity needs, risk assessment procedures, and resource and critical data identification. You must develop strategies that should be evaluated and tested. These strategies should take into consideration the dynamic changes occurring with new technologies, and changes that occur with the organization's growth.

Best Practice:

- Analyze your important infrastructure resources and mission-critical and business-critical data. Based on that analysis, create a backup strategy that will protect the company's critical infrastructure resources and business data.
- Identify with the organization's business managers the minimum recovery time for business-critical data. Based on that information, create an optimal restore strategy.
- Always test backup and restore procedures regularly. Perform testing in a non-production and isolated environment.

Tools

Tool	Use	Where to find it
Windows Server Backup	Performing on demand or scheduled backup and restoring data and servers	Server Manager - Tools
Microsoft Online Backup Service	Performing on-demand or scheduled backup to the cloud, and restoring data from the backup located in the cloud	Server Manager - Tools

MCT USE ONLY. STUDENT USE PROHIBITED

Module 8

Implementing Distributed Active Directory Domain Services Deployments

Contents:

Module Overview	8-1
Lesson 1: Overview of Distributed AD DS Deployments	8-2
Lesson 2: Deploying a Distributed AD DS Environment	8-9
Lesson 3: Configuring AD DS Trusts	8-18
Lab: Implementing Complex AD DS Deployments	8-23
Module Review and Takeaways	8-27

Module Overview

For most organizations, the Active Directory® Domain Services (AD DS) deployment may be the single most important component in the IT infrastructure. When organizations deploy AD DS or any of the other Active Directory–linked services within the Windows Server® 2012 operating system, they are deploying a central authentication and authorization service that provides Single Sign On (SSO) access to many other network services in the organization. AD DS provides the primary security mechanism for authentication and authorization within most organizations, and enables policy-based management for user and computer accounts. With other AD DS services, you can extend some of this functionality to users who are external to the organization.

This module will describe the key components of a complex AD DS environment, and how to install and configure a highly complex AD DS deployment.

Objectives

After completing this module, you will be able to:

- Describe the components of distributed AD DS deployments.
- Explain how to deploy a distributed AD DS deployment.
- Explain how to configure AD DS trusts.
- Explain how to implement complex AD DS deployments.

Lesson 1

Overview of Distributed AD DS Deployments

Before starting to configure a complex AD DS deployment, it is important to know the components that comprise the AD DS structure, and how they interact with each other to help provide a scalable and more secure IT environment. The lesson starts by examining the various components of an AD DS environment, in particular, domains, trees and forests.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the components of an AD DS environment.
- Explain how AD DS domains and forests form boundaries for security and administration.
- Explain reasons for having more than one domain in an AD DS environment.
- Explain reasons for having more than one forest in an AD DS environment.
- Explain the importance of Domain Name System (DNS) in a complex AD DS structure.
- Outline the options for upgrade and coexistence with previous AD DS versions.

Discussion: Overview of AD DS Components

An AD DS environment has various components, and it is important for you to understand the purpose of each component, and how they interact with each other. Some of the AD DS environment components are domains, trees, and forests. There is one global catalog in each forest, and there are trust relationships. It is also important to understand the purpose and benefit of these components.

Question: What is an AD DS domain?

Question: What is an AD DS domain tree?

Question: What is an AD DS forest?

Question: What are trust relationships?

Question: What is a global catalog?

- What is an AD DS domain?
- What is an AD DS tree?
- What is an AD DS forest?
- What is a trust relationship?
- What is the global catalog?

Overview of Domain and Forest Boundaries in an AD DS Structure

As already discussed, domains and forests provide boundaries within the AD DS namespace. An understanding of the different types of boundaries is essential to managing a complex AD DS environment.

Boundaries and Limits in AD DS Domains and Forests

The AD DS domain forms boundaries and limits for several items:

AD DS container	Replication and extent
Domain	All AD DS domain objects
	Direct local management of access to resources
	Security boundary: account settings such as password settings
	Security boundary and extent: audit policies
	Group Policy inheritance
	AD DS domain-integrated DNS zones and records
Forest	Schema partition
	Configuration partition
	Global catalog
	AD DS forest integrated DNS zones and records

- All AD DS objects that exist in a single domain are stored in the AD DS database on each domain controller in the domain. The replication process ensures that all originating updates are replicated to all of the other domain controllers in the same domain. In large organizations where there are a large number of changes, the replication traffic can have a noticeable effect on network bandwidth between the domain controllers. This is one of many factors that you must consider when designing an AD DS domain structure.
- The management of access to resources is usually straightforward within a single AD DS domain. Although you can grant users access to resources throughout the AD DS forest, it is simpler to manage permission within the AD DS domain boundary. This way, the administrators for the domain have the permission to do this in their own AD DS domain.
- Auditing is centrally managed by using GPOs. The maximum scope of these settings is at the AD DS domain level. It is possible to have the same audit settings in different AD DS domains, but then they must be managed separately in each domain.
- Group Policies can be linked at the following levels: local, site, domain, and organizational unit (OU). Apart from site-level Group Policies, the scope of Group Policies is the AD DS—domain—there is no inheritance of Group Policies from one AD DS domain to another, even if one AD DS domain is lower than another in the DNS namespace.
- The DNS services work best to support an AD DS environment when it is Active Directory–Integrated. This means that instead of the DNS records being stored locally on each DNS Server in text files, they are stored and replicated in the AD DS database. Because it is the database for the AD DS domain, it becomes the limit of replication of those records. The administrator can then decide whether to replicate the DNS information to all domain controllers in the domain (regardless of whether they are DNS servers), to all domain controllers that are DNS servers in the domain, or to all domain controllers that are DNS servers in the forest. For the last two options, separate replication partitions (domainDnsZones and forestDnsZones) exist. Alternatively, it is possible to create a custom partition where the administrator has to select manually which domain controllers participate in its replication, but this method is not often used.

The AD DS forest acts as a boundary for certain replication and management areas:

- The schema partition contains the rules and syntax for the AD DS database. This is replicated to all the domain controllers in the AD DS forest. Because the schema may have to be modified to support certain applications that are integrated with the AD DS database, there will have to be careful control over the schema modifications that are allowed, particularly to make sure that none of the updates adversely affect the operation of other applications or the AD DS database itself.
- The configuration partition contains the details of the AD DS domain layout, including: domains, domain controllers, replication partners, site and subnet information, and Dynamic Host Configuration Protocol (DHCP) authorization or the configuration of Dynamic Access Control. The

configuration partition also contains information about applications that are integrated with the AD DS database. An example of one of these applications is Exchange Server 2010.

- The global catalog is the read-only list containing every object in the entire AD DS forest. To keep it to a manageable size the global catalog contains only some attributes for each object, but it can still grow very large depending on the extent of the organization. The size of the global catalog and the number of ongoing changes to it are important factors in the allocation of which AD DS domain controllers will hold a copy of the global catalog. In addition, network bandwidth is also an important factor to take into account.
- In an AD DS forest with multiple AD DS domains, there is a DNS zone with forest-wide replication. This enables clients to locate records for AD DS domain controllers in other AD DS domains. There will be some DNS records that must be available to clients in every domain—for example, domain controllers that store a copy of the global catalog. When AD DS domain controllers start up, they register several server (SRV) resource records in the DNS database. Some of these records must be replicated to every DNS server in the AD DS forest, not just restricted to the AD DS domain—which would be the case if they were added to a regular Active Directory–integrated DNS zone. For this reason, a special forest-level DNS zone named forestDnsZones is created, and the replication of this is to all DNS servers in the AD DS forest, rather than to every DNS server or AD DS domain controller in the domain.

Why Implement Multiple Domains?

Many organizations can function adequately with a single AD DS domain. However, some entities require multiple domains for several reasons:

- The organization is decentralized and cannot support the numbers of users by using a centralized AD DS model: In this case, AD DS replication over network links may put an undo strain on the network connections. For this reason, it might be better to install a separate AD DS domain for the remote location. In practice, you would not do this unless there were a large number of accounts that needed to access the AD DS domain controllers.
- It may be necessary to divide a large AD DS database into more manageable sections. By doing so, you can also reduce the impact of AD DS replication traffic on the network.
- There is a requirement for different DNS namespaces: Sometimes there is a requirement to have more than one DNS namespace in an AD DS forest. This is typically the case when one company acquires another company, or merges with another organization, and there is need to preserve the domain names from the existing environment.
- Security: There may exist security or political requirements to have different parts of the AD DS database in different domains. If a company is setting up a facility in a foreign country and there are political or legal reasons why the new organization has to have a very distinct security base, they may need to implement separate AD DS domains.
- Dedicated root domain: It is best practice to separate the AD DS forest root domain from the day-to-day AD DS server usage. This model is sometimes referred to as an empty root domain or a dedicated root domain. Other variants are the peer-root domain, and the designated domain.

<p>Decentralized organization</p> <ul style="list-style-type: none"> • Large numbers of accounts • Slow or unreliable network links <p>Separate DNS namespaces required</p> <ul style="list-style-type: none"> • Merger or acquisition? <p>Security - political</p> <ul style="list-style-type: none"> • Legal or political constraints <p>Security: separate forest root</p> <ul style="list-style-type: none"> • Implement the empty root domain <p>Very high security</p> <ul style="list-style-type: none"> • Required for secret level security <p>Resource domains</p> <ul style="list-style-type: none"> • Company uses different domains

The AD DS forest root domain has two groups—the Schema Admins group and the Enterprise Admins group—that do not exist in any other domain in the AD DS forest. Because these groups have far-reaching rights in the AD DS forest, you may want to restrict the use of these groups by only using the AD DS forest root domain to store them. In early implementations of Active Directory, this model was often referred to as the empty root domain model.

- **Compliance:** It may be desirable to have all active domains at the same level in the namespace, so that your organization can utilize the empty root domain model. In certain organizations, it may be unacceptable to have different divisions in the same AD DS domain. The reason for this is that the Domain Admins in any AD DS domain have full control over every object in the domain, and this may violate certain corporate security policies. For example, different departments within an organization may be required for legal compliance reasons to not be in the same AD DS domain. In that case, there is the need to at least to create a separate AD DS domain for each department, if not a separate AD DS forest.
- **Resource domains:** For companies that have made the decision to utilize multiple domains, it might be more appropriate to provide separate resource domains for resources shared across the other domains.

Why Implement Multiple Forests?

Organizations may sometimes require that their AD DS design comprises more than one forest. There are several reasons why one AD DS forest may not be sufficient:

- **Security:** If your organization requires isolated security, then you should implement a separate AD DS forest. The AD DS forest root domain has the Schema Admins and Enterprise Admins groups—which can affect all the domains in the AD DS forest. Separate AD DS forests are often deployed by government defense contractors and other organizations where the isolation of security is a requirement.
- **Schema modifications:** Within your enterprise, there may be organizational groups that need separate control of their Active Directory schema. Because the schema is shared between the domains, multiple entities within a common forest must agree to those changes, which might take a large amount of time, or not be possible. Therefore, you may need to deploy different forests for those groups.
- **Security boundaries:** If two or more independent organizations want to share resources, but are not in a position where they are prepared to trust the domain administrators of the partner's organizations, then separate forests are needed. Having an AD DS structure with multiple forests provides security boundaries. Although domains in different forests can share resources, they rely on manually implemented trust relationships and additional administration. Each forest maintains its own isolated security databases and rules.
- **Politics:** Some countries have strict controls over the ownership of enterprises within the country. Having a separate AD DS forest may provide the administrative isolation to meet that need.

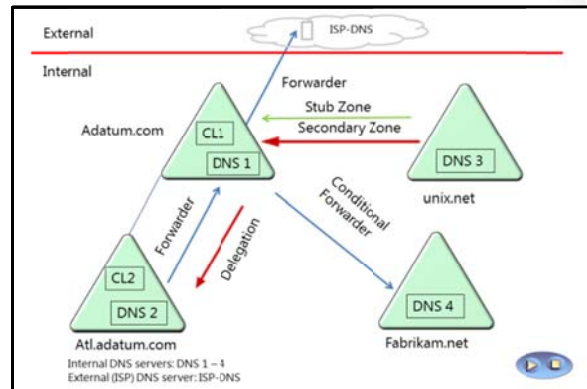
- * **Secret Security:** Isolate the AD DS forest root domain
- * **Schema conflicts:** Separate conflicting schema modifications
- * **Shared Security:** Isolated security model
- * **Politics:** Restraint in foreign territory

Note: The global catalog is available only within a single forest, so that when there is resource sharing between more than one forest, there will be directory lookups in two or more global catalogs.

Best Practice: As a best practice, choose the simplest design that achieves the required goal, as it will be less costly to implement and more straightforward to administer.

DNS Requirements for Complex AD DS Environments

For an AD DS domain to function, it requires DNS. There are many operations that rely on DNS lookups in order to take place. Name resolution is one of the original uses for Domain Name System. With name resolution, clients can connect to service points by using the DNS name, which resolves to an IP address. However, because so many operations involve connecting to an AD DS domain controller that offers a particular service. AD DS requires service (SRV) resource records. If a user wants to log on to their AD DS user account, then their system uses DNS lookups to find SRV records for a domain controller that is running the Kerberos service. When AD DS domain controllers need to communicate with each other, they use SRV records in DNS. If the AD DS forest contains more than one domain, then it is important to ensure that DNS lookups are successful for resources that are in a different domain from the DNS resolver (the client performing the DNS lookup).



There are several important configuration areas that you need to address when deploying a DNS structure to support a complex AD DS environment:


- **DNS Client configuration:** Configure all computers in the AD DS domain with at least two addresses of functional DNS servers. All computers must have good network connectivity with DNS servers. Servers will usually be configured statically, thus you should monitor their network configuration and reconfigure as necessary to meet any changes in the infrastructure.
- **IP Address Management (IPAM) monitoring:** A recommended option is to use IPAM to monitor the IP addressing, and the correct functioning and availability of DNS and DHCP servers in the AD DS forest. IPAM is new in Windows Server 2012, and it allows the central monitoring, reporting, and administration of decentralized DHCP and DNS services. When you are installing IPAM, you cannot install it on a domain controller, only on a domain member server.
- **Event escalation options:** If you have a monitoring solution such as Microsoft® System Center 2012 - Operations Manager, ensure that the events raised by IPAM are escalated in your global monitoring solution. IPAM will not provide the same escalation options as a monitoring infrastructure (for example, notifications).
- **Verification:** Verify that all of your computers, including domain controllers, are able to perform successful DNS lookups for key resources. Apart from routine name resolution for host to IP resolutions, all computers must be able to locate the SRV records for domain controllers in the AD DS domain and the AD DS forest.

Some of the ways to enable successful DNS lookups within a multi-domain AD DS environment are:

- CL1, the DNS resolver in the adatum.com DNS domain can perform successful queries for DNS records in the DNS domain adatum.com by contacting DNS 1. This is because DNS 1 stores the zone file for the adatum.com DNS domain. If DNS 1 receives a DNS query for a node outside of the local

DNS domain, and if DNS 1 does not have the answer already cached in memory, it will by default contact a DNS server from the Internet root DNS domain by using the root level hints configured by default on the DNS server.

- In order to speed up this process, you can configure DNS 1 to forward any queries that it cannot resolve by itself to one or more specific DNS servers. In this case, it could forward any query that it cannot resolve by itself to the DNS server for the Internet Service Provider (ISP). DNS 1 will utilize the fact that DNS-ISP will have cached a large number of DNS lookup replies and can answer (non-authoritatively) out of memory.
- In a multi domain setting, you can configure DNS servers that receive DNS queries that they are unable to resolve themselves, to forward these queries to other DNS server in the network. By using forwarding, DNS servers can forward all their Internet DNS queries through a central DNS server, which streamlines the process and speeds up the resolution time. This can greatly reduce the DNS resolution traffic through the firewall.
- If there are one or more separate DNS namespaces within your organization than can only be resolved by internal corporate DNS servers, you can facilitate this process by configuring conditional forwarders. The slide shows a separate DNS namespace for fabrikam.net, with its own DNS server, DNS 4. DNS 1 can be configured with a conditional forwarder so that queries for records in the fabrikam.net DNS namespace are directed to DNS 4, and all other queries are still forwarded to the ISP-DNS server.
- When there is a DNS domain that is lower in the DNS namespace (for instance the atl.adatum.com domain), you must configure the DNS servers for the parent DNS domain to enable DNS resolution for DNS records in the child domain. By default, DNS 1 has no knowledge of DNS 2. A delegated domain record in DNS creates a special subdomain in the adatum.com DNS domain that lists one or more DNS servers that store DNS records for the atl.adatum.com DNS domain. In this case, it will enable DNS 1 to pass DNS queries for atl.adatum.com to DNS 2
- Another option to allow DNS resolution in a disjointed DNS domain environment is to use a secondary zone. In this example, DNS 1 will store a complete read-only copy of all the records in the unix.net DNS zone. Because the records in the unix.net zone are updated regularly, the secondary zone will contain an up-to-date copy of all of the records in the unix.net DNS zone. In a large organization, this may involve a large amount of replication traffic, and if that traffic has a detrimental effect on network connections, then it may not be the optimum solution. Note that in this scenario, there will be virtually no DNS lookup traffic over the network to DNS 3, but there will be regular zone transfer traffic.
- The stub zone is another option for this case. The stub zone is a special type of secondary zone that only stores read-only records for the DNS servers in the remote DNS domain. However, like a standard secondary zone, the records are updated on a regular basis. The stub zone contains only the DNS records for the DNS server names and their IP addresses. This solution will often be the preferred option, because although there will be DNS lookup traffic over the network link, the stub zone update traffic will be low, and it may be a better solution.

 **Note:** Forwarders, conditional forwarders, and delegation are set up by an administrator, and point to IP addresses of one or more DNS servers. These are entered manually, but the DNS servers to which they refer may have changes and these will not update the DNS records automatically. If you decide to use delegation, forwarding, or conditional forwarding, then there will need to be a system for regularly checking that the IP addresses entered for those server referrals are still valid. This would not be necessary if you use stub zones as the solution, because they are regularly updated with fresh information.

Options for Upgrading and Coexistence with Previous AD DS Versions

If your current AD DS environment is running on Windows Server 2003 or Windows Server 2008–level AD DS domain controllers, you may want to consider upgrading to Windows Server 2012. AD DS has new features that are available only in AD DS domains running at Windows Server 2012 level.

If you decide to upgrade your existing AD DS domain controllers, you can upgrade them in place if they are running Windows Server 2008 or Windows Server 2008 R2. However, you cannot upgrade previous versions in place.

- Upgrade the existing AD DS domain controllers to Windows Server 2012
- Join one or more Windows Server 2012 servers to the AD DS domain, and promoting them to be AD DS domain controllers
- Introduce one or more AD DS domains running at Windows Server 2012 level

There are three options for upgrading your Active Directory domain controllers:

1. Upgrading the existing AD DS domain controllers to Windows Server 2012. This option involves performing an in-place upgrade. The existing AD DS domain controllers will be upgraded directly to Windows Server 2012. If the AD DS forest functional level is lower than Windows Server 2012 level then some schema upgrades will need to be made before starting the upgrade of the operating systems. The version of Server Manager that comes with Windows Server 2012 is able to detect the schema updates that are necessary, and will update the schema as part of the Server Manager AD DS Installation wizard.
2. Joining one or more Windows Server 2012 servers to the AD DS domain, and promoting them to be AD DS domain controllers. The Windows Server 2012 servers will be able to join the domain, but before they can be promoted to AD DS domain controllers, there will be some schema updates that need to be performed. Again, the Server Manager AD DS Installation wizard will perform automatically. The AD DS domain and forest functional levels must be at least Windows Server 2003 Native mode.
3. This third option does not immediately raise the AD DS domain to Windows Server 2012, but relies on introducing one or more AD DS domains that are running Windows Server 2012. In this scenario, one or more AD DS domains from another part of the forest or even a different forest will be connected to the original domain with trust relationships. This will allow the different AD DS domains to coexist and share resources. At some point in the future, they can be consolidated into one or more AD DS domains running at the Windows Server 2012 level.

Lesson 2

Deploying a Distributed AD DS Environment

This lesson outlines different ways to install an AD DS domain, and describes the different functional levels of AD DS domains and forests. In this lesson, you will learn about some of the important points that you must address when deploying a complex AD DS environment, and you will see how to upgrade from a previous version of AD DS.

Lesson Objectives

After completing this lesson, you will be able to:

- Explain how to install a domain controller in a new domain in a forest.
- Describe AD DS domain functional levels.
- Describe AD DS forest functional levels.
- Explain how to upgrade a previous version of AD DS to a Windows Server 2012 version.
- Explain how to migrate to Windows Server 2012 AD DS from a previous version.
- Describe some important considerations for implementing a complex AD DS environment.

Demonstration: Installing a Domain Controller in a New Domain in a Forest

In this demonstration, you will see how to install a domain controller in a new domain in a forest.

Demonstration Steps

Configure LON-SVR1 as an AD DS Domain Controller in atl.adatum.com

1. Log on to **LON-DC1** as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. On LON-DC1, in Server Manager, use the AD DS Installation Wizard to remotely install AD DS on LON-SVR1.
3. Use the AD DS Installation Wizard to install and configure LON-SVR1 as an AD DS domain controller in a new domain, atl.adatum.com.


Access LON-SVR1 as Adatum\Administrator


1. Select options to install DNS and global catalog, and set the password for the Directory Services Restore Mode administrator account.
2. Reboot and log on as **Adatum\Administrator** with the password **Pa\$\$w0rd**, on the newly created AD DS domain controller LON-SVR1.

AD DS Domain Functional Levels


AD DS domains can run at different functional levels. Generally, upgrading the domain to a higher functional level will introduce additional features. Some of the domain functional levels are listed in the following table.

- New functionality requires that domain controllers are running a particular version of Windows
 - Windows Server 2003
 - Windows Server 2008
 - Windows Server 2008 R2
 - Windows Server 2012
- Cannot raise functional level while domain controllers are running previous Windows versions
- Cannot add domain controllers running previous Windows versions after raising functional level



Domain functional levels	Features
Windows 2000 Server native	<ul style="list-style-type: none"> • Universal groups • Group nesting • Security identifier (SID) history <p>Install Domain Controllers from Media</p> <p>Install from Media allows the installation of an AD DS domain controller without impacting the network connection, because most of the new AD DS database is restored locally from an ntdsutil backup on a USB drive or DVD. Install from Media could also be accessed over the network when the network is not busy. Once the new AD DS domain controller is installed and rebooted, it will use the network to retrieve AD DS originating updates that were made since the ntdsutil backup was made. This will be a small amount of replication, unless there have been many originating updates to the AD DS database, which the new AD DS domain controller needs to bring up to date.</p> <p> Note: Windows Server 2012 domain controllers cannot be installed in a domain running at Windows 2000 Server native level.</p>
Windows Server 2003	<ul style="list-style-type: none"> • LastLogonTimestamp attribute remembers time of last domain logon for users, and replicates this to other AD DS domain controllers in the AD DS domain. • Constrained Delegation makes it possible for applications to take advantage of the secure delegation of user credentials by using Kerberos-based authentication. • Selective authentication allows you to specify the users and groups that are allowed to authenticate to specific resource servers in a trusting forest. • You can store DNS zones in application partitions, which allows them to be replicated on domain controllers that are also DNS servers in the domain, or even across the forest. • Group attributes and other multi-valued attributes are replicated at the attribute level, instead of the object level. In previous versions of AD DS, group membership was considered part of the object, and the group would be replicated as a single object. This meant that if two administrators changed the membership of the same group in the same replication period, the last write would win. The first changes made would be lost, because the new version of the group would replace the previous

Domain functional levels	Features
	<p>one entirely. With multivalued replication, group membership is treated at the attribute level, and therefore all originating updates are merged together. This also greatly reduces the replication traffic that would occur. An additional benefit from this is the removal of the previous group membership restriction that limited the maximum number of members to 5,000.</p>
Windows Server 2008	<ul style="list-style-type: none"> • Distributed File System – Replication (DFS-R) is available as a more efficient and robust file replication service for the SYSVOL folders. DFS-R can replace the file replication service NT File Replication Service (NTFRS). • A large amount of interactive logon information is stored for each user, instead of just last logon time. • Fine-grained password settings allow account policies to be set for users and groups, which replaces the default domain settings for those users or group members. • Personal virtual desktops are available for users to connect to, by using RemoteApp and Remote Desktop. • Advanced Encryption Services (AES 128 and 256) support for Kerberos is available. • Read-only domain controllers (RODCs) provide a secure and economic way to provide AD DS logon services in remote sites, without storing confidential information such as passwords in untrusted environments. • Group and other multivalue attributes are replicated on a per-value level, instead of being replicated together (which removed the limit of 5,000 users per group).
Windows Server 2008 R2	<ul style="list-style-type: none"> • Authentication mechanism assurance, which packages information about a user's logon method, can be used in conjunction with application authentication—for example, with Active Directory Federation Services (AD FS). In another example, a user logging on by using a smart card can be granted access to more resources than when they log on with a username and password. • Managed services accounts allow account passwords to be managed by the Windows operating system, and provide service principal name (SPN) management.
Windows Server 2012	<ul style="list-style-type: none"> • Instead of the Windows PowerShell® command-line interface, you can use Server Manager for setting up and managing the AD DS Recycle Bin. In Windows Server 2008, fine-grained password settings were complicated to set up and deploy. In Windows Server 2012, you can use Server Manager to deploy and manage these settings more conveniently. • Support for Dynamic Access Control and Kerberos armoring

 **Note:** Generally, you cannot roll back AD DS domain functional levels. However, in Windows Server 2012 and Windows Server 2008 R2, you are able to roll back to a minimum of Windows Server 2008, as long as you do not have optional features (such as the Recycle Bin) enabled. If you have implemented a feature that is only available in a higher domain functional level, you cannot rollback to an earlier state.



Additional Reading: To learn more about the AD DS domain functional levels, refer to the following link:

[http://technet.microsoft.com/en-us/library/understanding-active-directory-functional-levels\(v=ws.10\).aspx](http://technet.microsoft.com/en-us/library/understanding-active-directory-functional-levels(v=ws.10).aspx)

AD DS Forest Functional Levels

The AD DS forest can run at different functional levels, and sometimes raising the AD DS forest functional level makes additional features available. The most noticeable additional features come with the upgrade to a Windows Server 2003 forest functional level. Additional features that are made available with Windows Server 2003 include:

- Forest Trusts: AD DS forests can have trusts set up between them, which enables resource sharing. There are full trusts and selective trusts.
- Linked-value replication: This feature improved Windows 2000 Server replication, and improved how group membership was handled.
- Improved AD DS replication calculation algorithms: Knowledge Consistency Checker (KCC) and intersite topology generator (ISTG) use improved algorithms to speed up the calculation of the AD DS replication infrastructure, and provide much faster site link calculations.
- Support for Read Only Domain Controllers (RODCs). RODCs are supported at the Windows Server 2003 forest functional level. The RODC must be running Windows Server 2008 or later.
- Conversion of inetOrgPerson objects to user objects. You can convert an instance of an inetOrgPerson object, used for compatibility with certain non-Microsoft directory services, into an instance of class user. You can also convert a user object to an inetOrgPerson object.
- Deactivation and redefinition of attributes and object classes. Although you cannot delete an attribute or object class in the schema at the Windows Server 2003 functional level, you can deactivate or redefine attributes or object classes.

- Windows Server 2003
 - Forest trusts
 - Domain rename
 - Linked-value replication
 - Support for RODCs
 - Improved KCC
 - Conversion of inetOrgPerson objects to user objects
 - Deactivation and redefinition of attributes and object classes
- Windows Server 2008
 - No new features; sets minimum level for all new domains
- Windows Server 2008 R2
 - Active Directory Recycle Bin
- Windows Server 2012
 - No new features; sets minimum level for all new domains

The Windows Server 2008 forest functional level does not add new forest-wide features. The Windows Server 2008 R2 forest functional level adds the Active Directory Recycle Bin feature. This feature allows the ability to restore deleted Active Directory objects.

Although the Windows Server 2008 R2 AD DS forest functional level introduced AD DS Recycle Bin, the Recycle Bin had to be managed with Windows PowerShell. However, the version of Remote Server Administration Tools (RSAT) that comes with Windows Server 2012 has the ability to manage the AD DS Recycle Bin by using GUI tools.

When you raise the forest functional level, you limit possible domain functional levels for domains that you add to the forest. For example, if you raise the forest functional level to Windows Server 2012, you cannot add a new domain running at Windows Server 2008 R2 domain functional level.

Upgrading a Previous Version of AD DS to Windows Server 2012


To upgrade a previous versions of AD DS to Windows Server 2012 AD DS , you can use either of the following two methods:

- Upgrade the operating system on the existing domain controllers to Windows Server 2012.
- Introduce Windows Server 2012 servers as domain controllers in the existing domain. You can then decommission AD DS domain controllers running earlier versions of AD DS.

Of these two methods, the second is preferred, because there will be no old or disused code and files remaining. Instead, you will have a clean installation of the Windows Server 2012 operating system and AD DS database.

Options to upgrade AD DS to Windows Server 2012:


- In place upgrade (from Windows Server 2008 or Windows Server 2008 R2)
 - Benefit: Apart from prerequisite checks, all files and programs stay in place and there is no additional work required
 - Watch for: Leftover accumulated files and DLLs
- Introduce a new Windows Server 2012 server into the domain and promote it to be a domain controller
 - This option is the usually the preferred choice
 - Good: Provides a new server with no accumulated dross
 - Watch for: May require additional work to migrate users' files and profile settings, and to clean up metadata and obsolete DNS records
- Both options require that the schema is at the Windows Server 2012 level



Upgrading to Windows Server 2012

To upgrade an AD DS domain from Windows Server 2008 functional level to Windows Server 2012 functional level, you must first upgrade all the domain controllers from the Windows Server 2008 operating system to the Windows Server 2012 operating system. You can achieve this by upgrading all of the existing domain controllers to Windows Server 2012, or by introducing new domain controllers running Windows Server 2012, and then phasing out the existing domain controllers.

There is no reason to prevent Windows Server 2012 servers from being part of a Windows Server 2008 domain. However, before you can install the first domain controller that is running Windows Server 2012, you must upgrade the schema. In versions of AD DS prior to Windows Server 2012, you would run the adprep.exe tool to perform the schema upgrades. In a Windows Server 2012 environment, the Active Directory Domain Services Installation Wizard that is included in Server Manager incorporates the commands necessary to upgrade the AD DS forest schema.

 **Note:** Windows Server 2012 still provides a 64-bit version of ADPrep, so you can run Adprep.exe separately. For example, if the administrator installing the first Windows Server 2012 domain controller is not a member of the Enterprise Admins group, then you might need to run the command separately. You only have to run adprep.exe if you are planning to do an in-place upgrade for the first Windows Server 2012 domain controller in the domain.

The Upgrade Process

To upgrade the operating system of a Windows Server 2008 domain controller to Windows Server 2012:

1. Insert the installation disk for Windows Server 2012, and run **Setup**.
2. After the language selection page, select **Install now**.
3. After the operating system selection window and the license acceptance page, on the Which type of installation do you want? window, choose **Upgrade: Install Windows and keep files, settings, and apps**.

With this type of upgrade, AD DS on the domain controller is upgraded to Windows Server 2012 AD DS. . As a best practice, you should check for hardware and software compatibility before doing an upgrade. Following the operating system upgrade, remember to update your drivers and other services (such as monitoring agents), and check for updates for both Microsoft applications and non-Microsoft software.

The Clean Installation Process

To introduce a clean install of Windows Server 2012 as a domain member:

1. Deploy and configure a new installation of Windows Server 2012, and then join it to the domain.
2. Promote the new server to be a domain controller in the domain by using Server Manager.



Note: You can upgrade directly from Windows Server 2008 and Windows Server 2008 R2 to Windows Server 2012. To upgrade servers that are running a version of Windows Server that is older than Windows Server 2008, you must either perform an interim upgrade to Windows Server 2008 or Windows Server 2008 R2, or perform a clean install. Note that Windows Server 2012 AD DS domain controllers are able to coexist as domain controllers in the same domain as Windows Server 2003 domain controllers or newer.

Migrating to Windows Server 2012 AD DS from a Previous Version

As part of deploying AD DS, you might choose to restructure your environment for the following reasons:

- To optimize the arrangement of elements within the logical Active Directory structure.
- To assist in completing a business merger, acquisition, or divestiture.

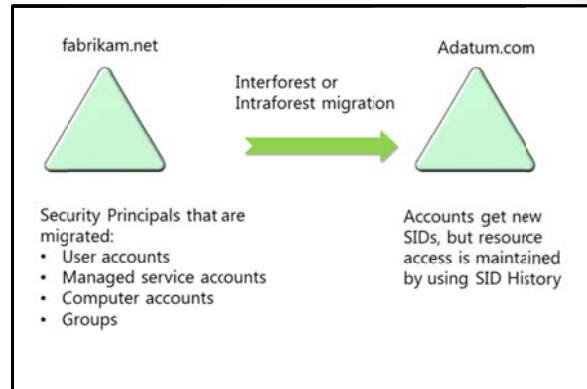
Restructuring involves the migration of resources between AD DS domains in either the same forest or in different forests. After you deploy AD DS, you might decide to further reduce the complexity of your environment by either restructuring AD DS domains between AD DS forests or restructuring domains within a single AD DS forest.

You can use the latest version of the Active Directory Migration Tool to perform object migrations and security translation as necessary, so that users can maintain access to network resources during the migration process.

Pre-Migration Steps

Before performing the migration, you must perform several tasks to prepare the source and target domains. These tasks include:

- For domain member computers that are pre-Windows Vista® Service Pack 1 (SP1) or Windows Server 2008 R2, configure a registry on the target AD DS domain controller to allow cryptography algorithms that are compatible with the Microsoft Windows NT® Server 4.0 operating system.
- Enable firewall rules on source and target AD DS domain controllers to allow file and printer sharing.
- Prepare the source and target AD DS domains to manage how the users, groups and user profiles will be handled.
- Create a rollback plan.
- Establish trust relationships that are required for the migration.
- Configure source and target AD DS domains to enable SID History migration.
- Specify service accounts for the migration.



- Perform a test migration, and fix any errors that are reported.


Migration Steps


When you are confident that all pre-migration steps have been completed, you can complete the migration. During this process, you will migrate user accounts, group accounts and computer accounts to the new domain. You will also assign permissions to network resources using the accounts in the new domain.

Post-Migration Steps

The new accounts (SIDs) will still have access to resources, because they retain an attribute called **SID History**. For example, a user uses a new user account to attempt to access a resource to which they had access with their old account. Instead of the user being denied because their SID is not referenced in the permissions, they can present their previous SID by using the **SID History** attribute. The migration tools can re-permit the resources so the new SIDs are entered into the permissions on the resource, and the old SID can be removed.

Once you have tested everything and verified that it is working in the new AD DS domain, you can decommission the old domain controllers.

 **Note:** The entire process may involve running the migration several times. For example, the user accounts would be migrated early in the process, but the user accounts profile migration would be accomplished in another pass of the migration tool.

 **Additional Reading:** Download the Active Directory Migration Tool version 3.2 from <http://www.microsoft.com/en-us/download/details.aspx?id=8377>. You can download the Active Directory Migration Tool Guide from <http://www.microsoft.com/en-us/download/details.aspx?id=19188>.

Considerations for Implementing a Complex AD DS Environment

Before implementing a complex AD DS environment, it is important to consider implications of the design if the deployment is to be successful. With proper planning, you can design an AD DS model to provide the administrative and security requirements for your organization.

- More than one AD DS forest?
- More than one AD DS tree?
- Number of AD DS domains
- DNS namespace design
- DNS resolution for host records and SRV records
- OUs
- Number and location of AD DS domain controllers
- Sites and replication topology

Some of the key points for consideration are listed in the following table.

Scenario	Key Points to Consider
More than one AD DS forest?	Security, politics, multiple schemas, administrative separation
More than one AD DS tree?	Multiple namespaces, acquisition, merger
Number of AD DS domains	Security, politics, administrative separation—for example, different departments in a governmental organization
DNS namespace design	Must support the proposed AD DS domain structure
DNS resolution for host records and SRV records	Must support resolution throughout the organization
OUs	Structure to support administrative delegation and Group Policy deployment
Number and location of AD DS domain controllers	Number of active accounts, network bandwidth, AD DS services availability, AD DS replication traffic
Sites and replication topology	AD DS replication requirements, network bandwidth, slow links, application support



Note: Details are discussed earlier in this module, and more information is available in Module 9.

AD DS Forest Root Domain

Each new AD DS forest starts with an AD DS forest root domain. This root domain has some unique features that do not exist in any other AD DS domain in the AD DS forest, including the following:

- The Schema Operations Master
- The Domain Naming Master
- The Schema Admins group
- The Enterprise Admins group

For this reason, the AD DS forest root domain must be treated with extra caution, particularly as the Enterprise Admins group and the Domain Admins group in the AD DS forest root domain have full control over every AD DS domain and object in the entire AD DS forest.

DNS Services

You should configure DNS services at the beginning of the deployment process, as all subsequent operations will depend on DNS functioning correctly. This also means that all computers should be configured with the IP addresses of at least two DNS servers so that they can successfully perform DNS lookups. In a complex environment, it will be necessary to decide how to make DNS records accessible to DNS resolvers (client computers).

Trust Relationships

You will also need to consider trust relationships for several reasons, such as:

- Enabling authentication between AD DS domain and external domains or realms.

- Enabling authentication between AD DS forests (forest trusts—complete or selective).
- Facilitating fast and reliable authentication traffic between AD DS domains in the same forest (shortcut trusts).

Multiple UPN Suffixes

Multiple User Principal Name (UPN) suffixes may be required to allow users to log on to their user accounts using an email account name from a different DNS namespace. For example, a user named Holly in the adatum.com AD DS domain could log on to that user account by using a UPN such as holly@fabrikam.com.

Lesson 3

Configuring AD DS Trusts

This lesson examines trust relationships and how they provide functionality for accessing resources and logging on to the domain. There are several types of trust relationships, and this lesson describes them in turn.

When an AD DS multi-domain forest is created, then trusts are automatically created to link all of the AD DS domains in the AD DS forest. In addition, there are other trusts that you can establish to provide additional functionality. These trusts include shortcut, external, realm, and forest trusts.

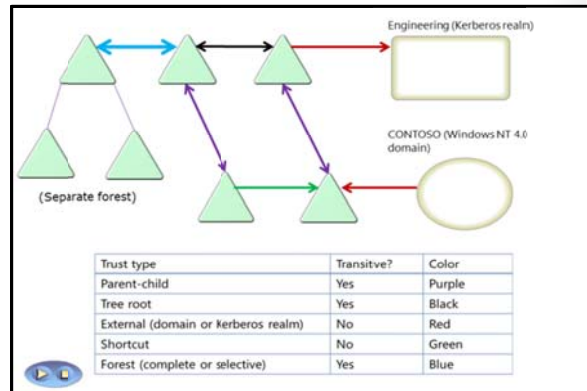
Lesson Objectives

After completing this lesson you will be able to:

- Describe the types of trusts that can be configured in a Windows Server 2012 environment.
- Explain how trusts work within an AD DS forest.
- Explain how trusts work between AD DS forests.
- Describe how to configure advanced trust settings.
- Describe how to configure a forest trust.

Overview of Different AD DS Trust Types

In a multi-domain AD DS forest, two-way transitive trust relationships are generated automatically between the AD DS domains, so that there is a path of trust between all of the AD DS domains. These trusts are called parent-child trusts. The trusts that are automatically created in the forest are all transitive trusts. That means that if A trusts B, and B trusts C, then A trusts C. However, this may not be the most efficient way to provide an authentication connection between all of the AD DS domains, and you can improve performance by setting up shortcut trusts.



There are other types of trust that you can deploy. For example, you can set up a realm trust with a non-Microsoft organization that is running Kerberos V5 and Windows NT 4.0 domains can be connected by using an external trust. The following table shows the main trust types.

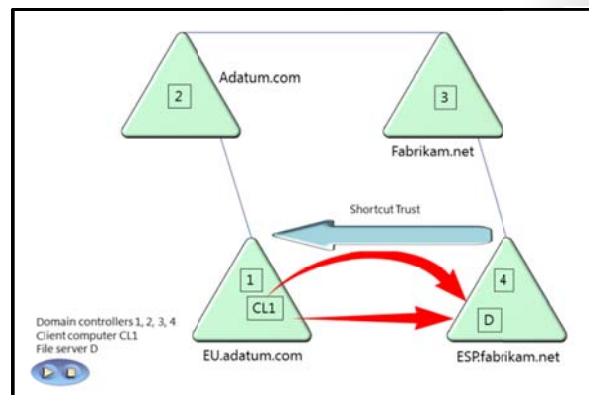
Trust type	Transitivity	Direction	Description
Parent and child	Transitive	Two-way	When a new AD DS domain is added to an existing AD DS tree, new parent and child trusts are created.
Tree-root	Transitive	Two-way	When a new AD DS tree is created in an existing AD DS forest, a new tree-root trust is created.
External	Non-transitive	One-way or two-way	External trusts enable resource access to be granted with a Windows NT 4.0 domain or an AD DS domain in another

Trust type	Transitivity	Direction	Description
			forest. These may also be set up to provide a framework for a migration.
Realm	Transitive or non-transitive	One-way or two-way	Realm trusts establish an authentication path between a Windows Server AD DS domain and a Kerberos V5 realm.
Forest (Complete or Selective)	Transitive	One-way or two-way	Trusts between AD DS forests allow two forests to share resources.
Shortcut	Transitive	One-way or two-way	Shortcut trusts improve authentication times between AD DS domains that are in different parts of an AD DS forest.

How Trusts Work Within a Forest

When you set up trusts between domains either within the same forest, across forests, or with an external realm, information about these trusts is stored in AD DS so you can retrieve it when necessary. A trusted domain object stores this information.

The trusted domain object stores information about the trust such as the trust transitivity and type. Whenever you create a trust, a new trusted domain object is created and stored in the System container in the trust's domain.



How Trusts Enable Users to Access Resources in a Forest

When a user attempts to access a resource in another domain, the Kerberos authentication protocol must determine whether the trusting domain has a trust relationship with the trusted domain.

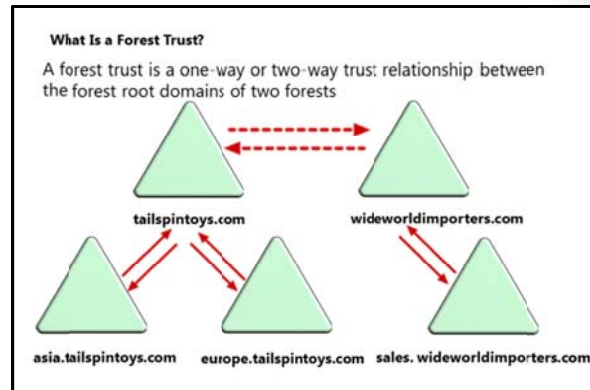
To determine this relationship, the Kerberos V5 protocol travels the trust path, utilizing trust information and DNS lookups to obtain a referral to the target domain's domain controller. The target domain controller issues a service ticket for the requested service. The trust path is the shortest path in the trust hierarchy.

When the user in the trusted domain attempts to access the resource in the other domain, the user's computer first contacts the domain controller in its domain to get authentication to the resource. If the resource is not in the user's domain, the domain controller uses the trust relationship with its parent, and refers the user's computer to a domain controller in its parent domain. This attempt to locate a resource continues up the trust hierarchy, possibly to the forest root domain, and down the trust hierarchy, until contact occurs with a domain controller in the domain where the resource is located.

How Trusts Work Between Forests

If the AD DS environment contains more than one forest, then it is possible to set up trust relationships between the AD DS forest roots. These forest trusts can be either complete trusts or selective trusts. Forest trusts can be one-way or two-way.

A single forest trust relationship allows users who are authenticated by a domain in one forest to access resources that are in the other forest, as long as they have been granted access rights. If the forest trust is one-way, domain controllers in the trusting forest can authenticate users in any domain in the trusted forest. Forest trusts are significantly easier to establish, maintain, and administer than separate trust relationships between each of the domains in the forests.



Forest trusts are particularly useful in scenarios that involve cross-organization collaboration or mergers and acquisitions, or within a single organization that has more than one forest in which to isolate Active Directory data and services. Forest trusts are also useful for application service providers, for collaborative business extranets, and for companies seeking a solution for administrative autonomy.

Forest trusts provide the following benefits:

- Simplified management of resources across two Windows Server 2008 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 V5 protocol and NTLM authentication protocols to improve the trustworthiness of authorization data that is transferred between forests.
- Flexibility of administration. Administrative tasks can be unique to each forest.

In AD DS in Windows Server 2008, you can link two Windows Server 2003 or Windows Server 2008 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.

You can create a forest trust only between two AD DS forests, and you cannot extend the trust implicitly to a third forest. This means that, if you create a forest trust between Forest 1 and Forest 2, and you create a forest trust between Forest 2 and Forest 3, Forest 1 does not have an implicit trust with Forest 3. Forest trusts are not transitive.

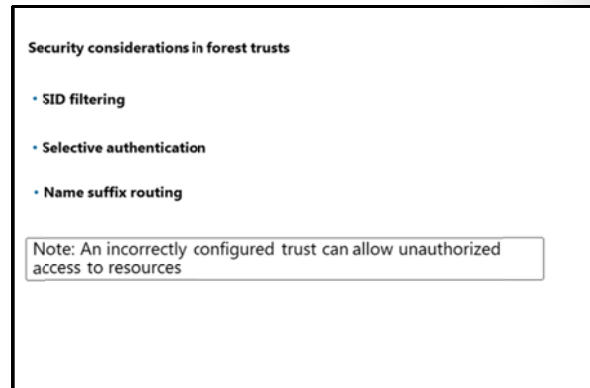
You must address several requirements before you can implement a forest trust, including that the forest functional level must be Windows Server 2003 or newer, and you must have DNS name resolution between the forests.

Configuring Advanced AD DS Trust Settings

In some cases, trusts can present security issues. Additionally, if you do not configure a trust properly, users who belong to another domain can gain unwanted access to some resources. There are several technologies that you can use to help control and manage security in a trust.

SID Filtering

By default, when you establish a forest or domain trust, you enable a domain quarantine, which is also known as SID filtering. When a user authenticates in a trusted domain, the user presents authorization data that includes the SIDs of all of the groups to which the user belongs. Additionally, the user's authorization data includes SIDs from other attributes of the user and the user's groups.



AD DS sets SID filtering by default to prevent malicious users who have access at the domain or enterprise administrator level in a trusted forest or domain, from granting (to themselves or to other user accounts in their forest or domain) elevated user rights to a trusting forest or domain. SID filtering prevents misuse of the attributes that contain SIDs on security principals (including InetOrgPerson objects) in the trusted forest or domain. One common example of an attribute that contains a SID is the SID history attribute (**SIDHistory**) on a user account object. Domain administrators typically use the SID history attribute to seamlessly migrate the user and group accounts that are held by a security principal from one domain to another. All SID filtering activities occur in the background, and administrators do not need to explicitly configure anything unless they want to disable SID filtering.

When security principals are created in a domain, the SID of the principal includes the domain SID, so that you can identify in which domain it was created. 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 user can access.

Authentication

When you create an external trust or a forest trust, you can manage the scope of authentication of trusted security principals. There are two modes of authentication for an external or forest trust:

- 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, this enables all trusted users to authenticate for services and access on all computers in the trusting domain. Trusted users can, therefore, 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 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, they 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 Windows Server 2008 forests and Windows Server 2003 forests that are joined by forest trusts. To simplify the administration of authentication requests, when you create a forest trust, AD DS routes all unique name suffixes by default. A *unique name suffix* is a name suffix within a forest, such as a UPN suffix, SPN suffix, or DNS forest or domain tree name that is not subordinate to any other name suffix. For example, the DNS forest name fabrikam.com is a unique name suffix within the fabrikam.com forest.

AD DS routes all names that are subordinate to unique name suffixes implicitly. For example, if your forest uses fabrikam.com as a unique name suffix, authentication requests for all child domains of fabrikam.com (childdomain.fabrikam.com) are routed, because the child domains are part of the fabrikam.com name suffix. Child names appear in the Active Directory Domains and Trusts snap-in. If you want to exclude members of a child domain from authenticating in the specified forest, you can disable name suffix routing for that name. You also can disable routing for the forest name itself.

Demonstration: Configuring a Forest Trust

In this demonstration, you will see how to configure DNS name resolution by using a conditional forwarder. You will also see how to configure a two-way selective forest trust.

Demonstration Steps

Configure DNS name resolution by using a conditional forwarder

- Configure DNS name resolution between adatum.com and treyresearch.net by creating a conditional forwarder so that LON-DC1 has a referral to MUN-DC1 as the DNS server for the DNS domain treyresearch.net.

Configure a two-way selective forest trust

- On LON-DC1, in Active Directory Domains and Trusts, create a two-way selective forest trust between adatum.com and treyresearch.net, by supplying the credentials of the treyresearch.net domain **Administrator** account.

Lab: Implementing Complex AD DS Deployments

Scenario

A. Datum Corporation has deployed a single AD DS domain with all the domain controllers located in its London data center. As the company has grown and added branch offices with large numbers of users, it is becoming increasingly apparent that the current AD DS environment is not meeting company requirements. The network team is concerned about the amount of AD DS-related network traffic that is crossing WAN links, which are becoming highly utilized.

The company has also become increasingly integrated with partner organizations, some of whom need access to shared resources and applications that are located on the A. Datum internal network. The security department at A. Datum wants to ensure that the access for these external users is as secure as possible.

As one of the senior network administrators at A. Datum, you are responsible for implementing an AD DS infrastructure that will meet the company requirements. You are responsible for planning an AD DS domain and forest deployment that will provide optimal services for both internal and external users, while addressing the security requirements at A. Datum.

Objectives

- Implement child domains in AD DS.
- Implement forest trusts in AD DS.

Lab Setup

Estimated Time: 45 minutes

- 20412A-LON-DC1
- 20412A-TOR-DC1
- 20412A-LON-SVR1
- 20412A-MUN-DC1

User name: **Adatum\Administrator**

Password: **Pa\$\$w0rd**

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

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, click **20412A-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on using the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
5. Repeat steps 2-4 for **20412A-LON-SVR1** and **20412A-TOR-DC1**.
6. Start **20412A-MUN-DC1** and log on as **Treyresearch\Administrator** with the password of **Pa\$\$w0rd**.

Exercise 1: Implementing Child Domains in AD DS

Scenario

A. Datum has decided to deploy a new domain in the adatum.com forest for the North American region. The first domain controller will be deployed in Toronto, and the domain name will be na.adatum.com. You need to configure and install the new domain controller.

The main tasks for this exercise are as follows:

1. Configure Domain Name System (DNS) for domain delegation.
2. Install a domain controller in a child domain.
3. Verify the default trust configuration.

► Task 1: Configure Domain Name System (DNS) for domain delegation


- On LON-DC1, open DNS Manager and configure a delegated zone record for na.adatum.com. Specify TOR-DC1 as the authoritative DNS server.

► Task 2: Install a domain controller in a child domain

1. On TOR-DC1, use Server Manager to install AD DS.
2. When the AD DS binaries have installed, use the **Active Directory Domain Services Configuration Wizard** to install and configure **TOR-DC1** as an AD DS domain controller for a new child domain named **na.adatum.com**.
3. When prompted, use **Pa\$\$w0rd** as the **Directory Services Restore Mode (DSRM)** password.

► Task 3: Verify the default trust configuration

1. Log on to **TOR-DC1** as **NA\Administrator** using the password **Pa\$\$w0rd**.
2. When Server Manager opens, click **Local Server**. Verify that **Windows Firewall** shows **Domain: On**. If it does not, then next to **Local Area Connection** click **172.16.0.25, IPv6 enabled**. Right-click **Local Area Connection** and then click **Disable**. Right-click **Local Area Connection** and then click **Enable**. The Local Area Connection should now show **Adatum.com**.
3. From Server Manager, launch the **Active Directory Domains and Trusts** management console and verify the parent child trusts.

 **Note:** If you receive a message that the trust cannot be validated, or that the secure channel (SC) verification has failed, ensure that you have completed step 2 and then wait for at least 10-15 minutes. You can continue with the lab and come back later to verify this step.

Results: After completing this exercise, you will have implemented child domains in AD DS.

Exercise 2: Implementing Forest Trusts

Scenario

A. Datum is working on several high-priority projects with a partner organization named Trey Research. To simplify the process of enabling access to resources located in the two organizations, they have deployed a dedicated wide area network (WAN) between London and Munich, where Trey Research is located. You now need to implement and validate a forest trust between the two forests, and configure the trust to allow access to only selected servers in London.

The main tasks for this exercise are as follows:

1. Configure stub zones for DNS name resolution.
2. Configure a forest trust with selective authentication.
3. Configure a server for selective authentication.

► **Task 1: Configure stub zones for DNS name resolution**

1. Log on to **LON-DC1** as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
2. Using the DNS management console, configure a DNS stub zone for treyresearch.net.
3. Use **172.16.10.10** as the Master DNS server.
4. Close DNS Manager.
5. Log on to **MUN-DC1** as **TreyResearch\Administrator** with the password **Pa\$\$w0rd**.
6. Using the DNS management console, configure a DNS stub zone for adatum.com.
7. Use **172.16.0.10** as the Master DNS server.
8. Close DNS Manager.

► **Task 2: Configure a forest trust with selective authentication**

1. On LON-DC1, create a one-way: outgoing trust between the treyresearch.net AD DS forest and the adatum.com forest. Configure the trust to use Selective authentication.
2. Confirm and validate the trust from treyresearch.net.
3. Close Active Directory Domains and Trusts.

► **Task 3: Configure a server for selective authentication**

1. On LON-DC1, from Server Manager open **Active Directory Users and Computers**.
2. On LON-SVR1, configure the members of **treyresearch.com\it** group with the **Allowed to authenticate** permission. If you are prompted for credentials, type **Treyresearch\administrator** with the password of **Pa\$\$w0rd**.
3. On LON-SVR1, create a shared folder **IT-Data** and grant access to members of the treyresearch.net\it group. If you are prompted for credentials, type **Treyresearch\administrator** with the password of **Pa\$\$w0rd**.
4. Log off of MUN-DC1.
5. Log on to **MUN-DC1** as **treyresearch\alice**, and access the shared folder on LON-SVR1.

Results: After completing this exercise, you will have implemented forest trusts.

► **To prepare for the next module**

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

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

Lab Review

Question: Why did you configure a delegated subdomain record in DNS on LON-DC1 before adding the child domain na.adatum.com?

Question: What are the alternatives to creating a delegated subdomain record in Q1?

Question: When you are creating a forest trust, why would you create a selective trust instead of a complete trust?

Module Review and Takeaways

To design and implement a reliable and efficient AD DS environment, it is important to have an understanding of which components are required and how they interact. This module covered the constituent parts of an AD DS design and also demonstrated the different ways to deploy AD DS in a complex scenario.

The Domain Name Service is crucial to the satisfactory functioning of an AD DS system, and students saw different ways to provide a robust DNS record resolution process in a multi-domain situation. The different DNS resolution methods were discussed, including forwarders, conditional forwarders, delegation, secondary zones, and stub zones. Students also saw how trust relationships can provide an effectual authentication mechanism in various environments.

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
You receive error messages such as: DNS lookup failure, RPC server unavailable, domain does not exist, domain controller could not be found.	
User cannot be authenticated to access resources on another AD DS domain or Kerberos realm.	

MCT USE ONLY. STUDENT USE PROHIBITED

Module 9

Implementing Active Directory Domain Services Sites and Replication

Contents:

Module Overview	9-1
Lesson 1: Overview of AD DS Replication	9-2
Lesson 2: Configuring AD DS Sites	9-10
Lesson 3: Configuring and Monitoring AD DS Replication	9-16
Lab: Implementing AD DS Sites and Replication	9-22
Module Review and Takeaways	9-26

Module Overview

When you deploy Active Directory® Domain Services (AD DS), it is important to provide an efficient logon infrastructure and a highly available directory service. Implementing multiple domain controllers throughout the infrastructure helps you meet both of these goals. However, you must ensure that AD DS replicates Active Directory information between each domain controller in the forest. In this module, you will learn how AD DS replicates information between domain controllers within a single site and throughout multiple sites. You also will learn how to create multiple sites and monitor replication to help optimize AD DS replication and authentication traffic.

Objectives

After completing this module, you will be able to:

- Describe how AD DS replication works.
- Configure AD DS sites to help optimize authentication and replication traffic.
- Configure and monitor AD DS replication.

Lesson 1

Overview of AD DS Replication

Within an AD DS infrastructure, standard domain controllers replicate Active Directory information by using a multimaster replication model. This means that if a change is made on one domain controller, that change then replicates to all other domain controllers in the domain, and potentially to all domain controllers throughout the entire forest. This lesson provides an overview of how AD DS replicates information between both standard and read-only domain controllers (RODC).

Lesson Objectives

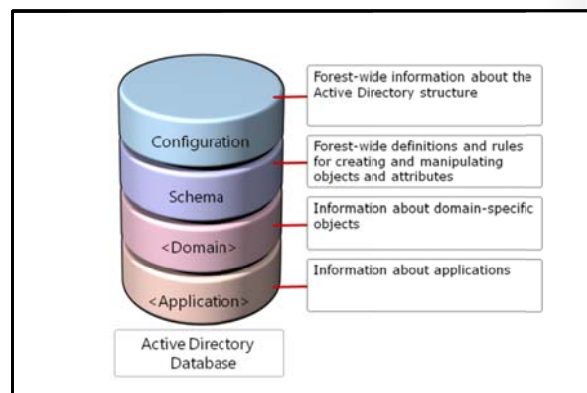
After completing this lesson, you will be able to:

- Describe the AD DS partitions.
- Describe the characteristics of AD DS replication.
- Describe the replication process within a single site.
- Describe how AD DS resolves replication conflicts.
- Describe how you generate the replication topology.
- Describe how read-only domain controller replication works.
- Describe System Volume (SYSVOL) replication.

What Are AD DS Partitions?

The Active Directory data store contains information that AD DS distributes to all domain controllers throughout the forest infrastructure. Much of the information that the data store contains is distributed within a single domain. However, some information may be related to, and replicated throughout, the entire forest, regardless of the domain boundaries.

To help provide replication efficiency and scalability between domain controllers, the Active Directory data is separated logically into several partitions. Each partition is a unit of replication, and each partition has its own replication topology. The default partitions include the following:



- **Configuration partition.** The configuration partition is created automatically when you create the first domain of a forest. The configuration partition contains information about the forest-wide AD DS structure, including which domains and sites exist and which domain controllers exist in each domain. The configuration partition also stores information about forest-wide services such as DHCP authorization and certificate templates. This partition replicates to all domain controllers in the forest.
- **Schema partition.** The schema partition contains definitions of all the objects and attributes that you can create in the data store, and the rules for creating and manipulating them. Schema information replicates to all domain controllers in the forest. Therefore, all objects must comply with the schema object and attribute definition rules. AD DS contains a default set of classes and attributes that you cannot modify. However, if you have Schema Admins credentials, you can extend the schema by adding new attributes and classes to represent application-specific classes. Many applications such as Microsoft Exchange Server and Microsoft System Center Configuration Manager may extend the

schema to provide application-specific configuration enhancements. These changes target the domain controller that contains the forest's schema master role. Only the schema master is permitted to make additions to classes and attributes.

- **Domain partition.** When you create a new domain, AD DS automatically creates and replicates an instance of the domain partition to all of the domain's domain controllers. The domain partition contains information about all domain-specific objects, including users, groups, computers, organizational units (OUs), and domain-related system settings. All objects in every domain partition in a forest are stored in the global catalog, with only a subset of their attribute values.
- **Application partition.** The application partition stores nondomain, application-related information that may have a tendency to be updated frequently or have a specified lifetime. An application is typically programmed to determine how it stores, categorizes, and uses application-specific information stored in the Active Directory database. To prevent unnecessary replication of an application partition, you can designate which domain controllers in a forest will host the specific application's partition. Unlike a domain partition, an application partition does not store security principal objects, such as user accounts. Additionally, the global catalogue does not store data contained in application partitions.



Note: You can use ADSI Edit to connect to and view the partitions.

Characteristics of AD DS Replication

An effective AD DS replication design ensures that each partition on a domain controller is consistent with the replicas of that partition hosted on other domain controllers. Typically, not all domain controllers have exactly the same information in their replicas at any one moment because changes are occurring to the directory constantly. However, Active Directory replication ensures that all changes to a partition are transferred to all replicas of the partition. Active Directory replication balances accuracy (or integrity) and consistency (called *convergence*) with performance (keeping replication traffic to a reasonable level).

- Multimaster replication ensures:
 - Accuracy (integrity)
 - Consistency (convergence)
 - Performance (keeping replication traffic to a reasonable level)
- Key characteristics of Active Directory replication include:
 - Multimaster replication
 - Pull replication
 - Store-and-forward
 - Partitions
 - Automatic generation of an efficient, robust replication topology
 - Attribute-level replication
 - Distinct control of intrasite and intersite replication
 - Collision detection and remediation

The key characteristics of Active Directory replication are:

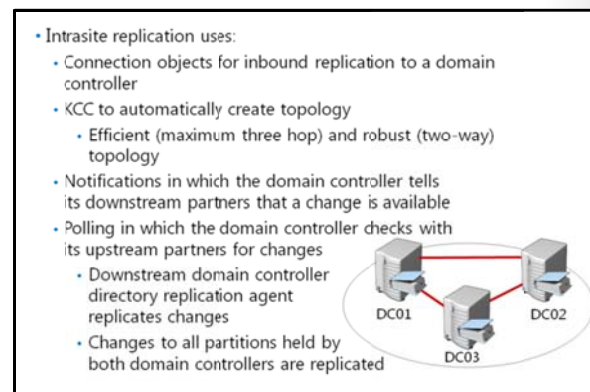
- **Multimaster replication.** Any domain controller except RODCs can initiate and commit a change to AD DS. This provides fault tolerance, and eliminates dependency on a single domain controller to maintain the operations of the directory store.
- **Pull replication.** A domain controller requests, or *pulls*, changes from other domain controllers. Even though a domain controller can notify its replication partners that it has changes to the directory, or poll its partners to see if they have changes to the directory, in the end, the target domain controller requests and pulls the changes themselves.
- **Store-and-forward replication.** A domain controller can pull changes from one partner, and then make those changes available to another partner. For example, domain controller B can pull changes initiated by domain controller A. Then, domain controller C can pull the changes from domain controller B. This helps balance the replication load for domains that contain several domain controllers.

- Data store partitioning. A domain's domain controllers only host the domain-naming context for their domains, which helps minimize replication, particularly in multidomain forests. By default, other data, including application directory partitions and the partial attribute set (global catalog), do not replicate to every domain controller in the forest.
- Automatic generation of an efficient and robust replication topology. By default, AD DS configures an effective, two-way replication topology so that the loss of one domain controller does not impede replication. AD DS automatically updates this topology as domain controllers are added, removed, or moved between sites.
- Attribute-level replication. When an attribute of an object changes, only that attribute, and minimal metadata that describes that attribute, replicates. The entire object does not replicate, except upon its initial creation.
- Distinct control of intrasite replication and intersite replication. You can control replication within a single site *and* between sites.
- Collision detection and management. On rare occasions, you can modify an attribute on two different domain controllers during a single replication window. If this occurs, you must reconcile the two changes. AD DS has resolution algorithms that satisfy almost all scenarios.

How AD DS Replication Works Within a Site

AD DS replication within a single site is called *intrasite replication*, which takes place automatically. However, you can configure it to occur manually, as necessary. The following concepts are related to intrasite replication:

- Connection objects
- The knowledge consistency checker (KCC)
- Notification
- Polling



Connection Objects

A domain controller that replicates changes from another domain controller is called a replication partner. Replication partners are linked by connection objects. A connection object represents a replication path from one domain controller to another. Connection objects are one-way, representing inbound-only pull replication.

To view and configure connection objects, open Active Directory Sites and Services, and then select the NTDS Settings container of a domain controller's server object. You can force replication between two domain controllers by right-clicking the connection object, and then selecting **Replicate Now**. Note that replication is inbound-only, so if you want to replicate both domain controllers, you need to replicate the inbound connection object of each domain controller.

The Knowledge Consistency Checker

The replication paths built between domain controllers by connection objects create the forest's replication topology. You do not have to create the replication topology manually. By default, AD DS creates a topology that ensures effective replication. The topology is two-way, which means that if any one domain controller fails, replication continues uninterrupted. The topology also ensures that there are no more than three hops between any two domain controllers.

On each domain controller, a component of AD DS called the knowledge consistency checker (KCC) helps generate and optimize the replication automatically between domain controllers within a site. The KCC evaluates the domain controllers in a site, and then creates connection objects to build the two-way, three-hop topology described earlier. If you add or remove a domain controller, or if a domain controller is not responsive, the KCC rearranges the topology dynamically, adding and deleting connection objects to rebuild an effective replication topology. The KCC runs at specified intervals (every 15 minutes by default) and designates replication routes between domain controllers that are the most favorable connections available at the time.

You can create connection objects manually to specify replication paths that should persist. However, creating a connection object manually is not typically required or recommended because the KCC does not verify or use the manual connection object for failover. The KCC will also not remove manual connection objects, which means that you must delete connection objects that you create manually.

Notification

When a change is made to an Active Directory partition on a domain controller, the domain controller queues the change for replication to its partners. By default, the source server waits 15 seconds to notify its first replication partner of the change. Notification is the process by which an upstream partner informs its downstream partners that a change is available. By default, the source domain controller then waits three seconds between notifications to additional partners. These delays, called the *initial notification delay* and the *subsequent notification delay*, are designed to stagger the network traffic that intrasite replication can cause.

Upon receiving the notification, the downstream partner requests the changes from the source domain controller, and the directory replication agent pulls the changes from the source domain controller. For example, suppose domain controller DC01 makes an initial change to AD DS. It is the originating domain controller, and the change that it makes, that originates the change. When DC02 receives the change from DC01, it makes the change to its directory. DC02 then queues the change for replication to its own downstream partners.

Then, suppose DC03 is a downstream replication partner of DC02. After 15 seconds, DC02 notifies DC03 that it has a change. DC03 makes the replicated change to its directory, and then notifies its downstream partners. The change has made two hops, from DC01 to DC02, and then from DC02 to DC03. The replication topology ensures that no more than three hops occur before all domain controllers in the site receive the change. At approximately 15 seconds per hop, the change fully replicates in the site within one minute.

Polling

At times, a domain controller may not make any changes to its replicas for an extended time, particularly during off hours. Suppose this is the case with DC01. This means that DC02, its downstream replication partner, will not receive notifications from DC01. DC01 also might be offline, which would prevent it from sending notifications to DC02.

It is important for DC02 to know that its upstream partner is online and simply does not have any changes. This is achieved through a process called polling. During polling, the downstream replication partner contacts the upstream replication partner with queries as to whether any changes are queued for replication. By default, the polling interval for intrasite replication is once per hour. You can configure the polling frequency from a connection object's properties by clicking **Change Schedule**, although we do not recommend it. If an upstream partner fails to respond to repeated polling queries, the downstream partner launches the KCC to check the replication topology. If the upstream server is indeed offline, the KCC rearranges the site's replication topology to accommodate the change.

Resolving Replication Conflicts

Because AD DS supports a multimaster replication model, replication conflicts may occur. Typically, there are three types of replication conflicts that may occur in AD DS:

- Simultaneously modifying the same attribute value of the same object on two domain controllers.
- Adding or modifying the same object on one domain controller at the same time that the container object for the object is deleted on another domain controller.

- In multimaster replication models, replication conflicts arise when:
 - The same attribute is changed on two domain controllers simultaneously
 - An object is moved or added to a deleted container on another domain controller
 - Two objects with the same relative distinguished name are added to the same container on two different domain controllers
- To resolve replication conflicts, AD DS uses:
 - Version number
 - Time stamp
 - Server GUID

- Adding objects with the same relative distinguished name into the same container.

To help minimize conflicts, all domain controllers in the forest record and replicate object changes at the attribute level rather than at the object level. Therefore, changes to two different attributes of an object, such as the user's password and postal code, do not cause a conflict even if you change them at the same time from different locations.

When an originating update is applied to a domain controller, a stamp is created that travels with the update as it replicates to other domain controllers. The stamp contains the following components:

- Version number. The version number starts at one for each object attribute, and increases by one for each update. When performing an originating update, the version of the updated attribute is one number higher than the version of the attribute that is being overwritten.
- Timestamp. The timestamp is the update's originating time and date according to the system clock of the domain controller where the change is made.
- Server globally unique identifier (GUID). The server GUID identifies the domain controller that performed the originating update.

Resolving Replication Conflicts

The table below outlines several conflicts and how AD DS resolves the issue:

Conflict	Resolution
Attribute value	If the version number value is the same, but the attribute value is different, then the timestamp is evaluated. The update operation that has the higher stamp value replaces the attribute value of the update operation with the lower stamp value.
Add or move under a deleted container object, or the deletion of a container object	After resolution occurs at all replicas, AD DS deletes the container object, and the leaf object is made a child of the folder's special LostAndFound container. Stamps are not involved in this resolution.
Adding objects with the same relative distinguished name	The object with the larger stamp keeps the relative distinguished name. AD DS assigns the sibling object a unique relative distinguished name by the domain controller. The name assignment is the relative distinguished name + CNF: + a reserved character (the asterisk,) + the object's GUID. This name assignment ensures that the generated name does not conflict with any other object's name.

How Replication Topology Is Generated

Replication topology is the route by which replication data travels through a network. To create a replication topology, AD DS must determine which domain controllers replicate data with other domain controllers. AD DS creates a replication topology based on the information that AD DS contains. Because each AD DS partition may be replicated to different domain controllers in a site, the replication topology can differ for schema, configuration, domain, and application partitions.

Because all domain controllers within a forest share schema and configuration partitions, AD DS replicates schema and configuration partitions to all domain controllers. Domain controllers in the same domain also replicate the domain partition. Additionally, domain controllers that host an application partition also replicate the application partition. To optimize replication traffic, a domain controller may have several replication partners for different partitions. In a single site, the replication topology will be fault tolerant and redundant. This means that if the site contains more than two domain controllers, each domain controller will have at least two replication partners for each AD DS partition.

How the Schema and Configuration Partitions Are Replicated

Replication of the schema and configuration partitions follows the same process as all other directory partitions. However, because these partitions are forest-wide rather than domain-wide, connection objects for these partitions may exist between any two domain controllers regardless of the domain controller's domain. Furthermore, the replication topology for these partitions includes all domain controllers in the forest.

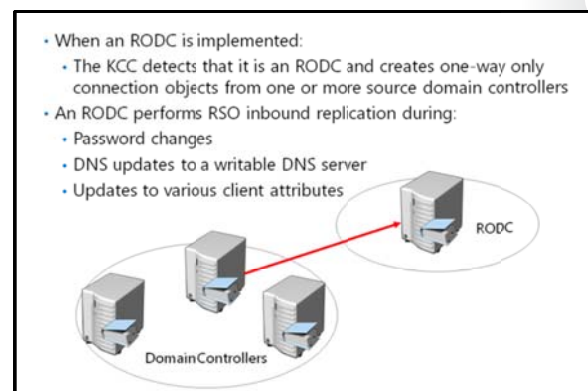
How the Global Catalog Affects Replication

The configuration partition contains information about the site topology and other global data for all domains that are members of the forest. AD DS replicates the configuration partition to all domain controllers through normal forest-wide replication. Each global catalog server obtains domain information by contacting a domain controller for that domain and obtaining the partial replica information. The configuration partition also provides the domain controllers with a list of the forest's global catalog servers.

Global catalog servers register DNS service records in the DNS zone that corresponds to the forest root domain. These records, which are registered only in the Forest Root DNS zone, help clients and servers locate global catalog servers throughout the forest to provide client logon services.

How RODC Replication Works

As previously mentioned, domain controllers replicate data by pulling changes from other originating domain controllers. A RODC does not allow any non-replicated changes to be written to its database and never replicates any information out to other domain controllers. Since changes are never written to an RODC directly, other domain controllers do not have to pull directory changes from an RODC. Restricting RODCs from originating changes prevents any changes or corruption that a malicious user or application might make from replicating to the rest of the forest.



When a user or application attempts to perform a write request to a RODC, one of the following actions typically occurs:

- The RODC forwards the write request to a writable domain controller, which is then replicated back to the RODC. Examples of this type of request includes password changes, service principal name (SPN) updates, and computer\domain member attribute changes.
- The RODC responds to the client and provides a referral to a writable domain controller. The application can then communicate directly with a writable domain controller. Lightweight Directory Access Protocol (LDAP) and DNS record updates are examples of acceptable RODC referrals.
- The write operation fails because it is not referred or forwarded to a writable domain controller. Remote procedure call (RPC) writes are an example of communication that may be prohibited from referrals or forwarded to another domain controller.

When you implement an RODC, the KCC detects that the domain controller is configured with a read-only replica of all applicable domain partitions. Because of this, the KCC creates one-way only connection objects from one or more source Windows Server 2008 or higher domain controllers to the RODC.

For some tasks, an RODC performs inbound replication using a replicate-single-object (RSO) operation. This is initiated on-demand outside of the standard replication schedule. These tasks include:

- Password changes.
- DNS updates when a client is referred to a writable DNS server by the RODC. The RODC then attempts to pull the changes back using an RSO operation. This only occurs for Active Directory-integrated DNS zones.
- Updates for various client attributes including **client name**, **DnsHostName**, **OsName**, **OsVersionInfo**, supported encryption types, and the **LastLogonTimeStamp** attribute.

How SYSVOL Replication Works

The SYSVOL is a collection of files and folders on each domain controller that is linked to %SystemRoot%\SYSVOL location. SYSVOL contains logon scripts and objects related to Group Policy such as Group Policy templates. The contents of the SYSVOL folder replicate to every domain controller in the domain using the connection object topology and schedule that the KCC creates.

Depending on the domain controller operating system version, domain's functional level, and migration status of SYSVOL, the File Replication Service or Distributed File System Replication replicates SYSVOL changes between domain controllers. The File Replication Service was used primarily in Windows Server 2003 R2 and older domain structures. The File Replication Service has limitations in both capacity and performance which has led to the adoption of Distributed File System Replication.

In Windows Server 2008 and newer domains, you can use Distributed File System Replication to replicate the contents of SYSVOL. Distributed File System Replication supports replication scheduling and bandwidth throttling, and it uses a compression algorithm known as Remote Differential Compression (RDC). Using RDC, Distributed File System Replication replicates only the differences (or changes within files) between the two servers, resulting in lower bandwidth use during replication.

- SYSVOL contains logon scripts, Group Policy templates, and Group Policy Objects with their content
- SYSVOL replication can take place using:
 - File Replication Service: primarily used in Windows Server 2003 and older domain structures
 - Distributed File System Replication : used in Windows Server 2008 and newer domains
- To migrate SYSVOL replication from the File Replication Service to Distributed File System Replication:
 - The domain functional level must be at least Windows Server 2008
 - Use the Dfsrmig.exe tool to perform the migration



Note: You can use the `dfsrmig.exe` tool to migrate SYSVOL replication from the File Replication Service to Distributed File System Replication . For the migration to succeed, the domain functional level must be at least Windows Server 2008.

Lesson 2

Configuring AD DS Sites

Within a single site, AD DS replication occurs automatically without regard for network utilization. However, some organizations have multiple locations that are connected by wide area network (WAN) connections. If this is the case, you must ensure that AD DS replication does not impact network utilization negatively between locations. You also may need to localize network services to a specific location. For example, you may want users at a branch office to authenticate to a domain controller located in their local office, rather than over the WAN connection to a domain controller located in the main office. You can implement AD DS sites to help manage bandwidth over slow or unreliable network connections, and to assist in service localization for authentication as well as many other site-aware services on the network.

Lesson Objectives

After completing this lesson, you will be able to:

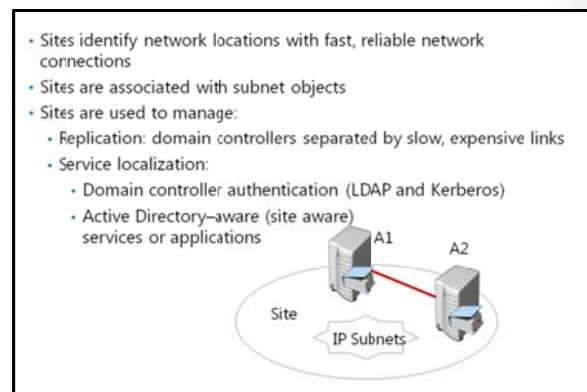
- Describe AD DS sites.
- Explain why organizations might implement additional sites.
- Configure additional AD DS sites.
- Describe how AD DS replication works between sites.
- Describe the Inter-site Topology Generator.
- Describe how Service Locator (SRV) records are used to locate domain controllers.
- Describe how client computers locate domain controllers.

What Are AD DS Sites?

To most administrators, a site is a physical location, an office, or a city typically separated by a WAN connection. These sites are physically connected by network links that might be as basic as dial-up connections or as sophisticated as fiber links. Together, the physical locations and links make up the physical network infrastructure.

AD DS represents the physical network infrastructure with objects called *sites*. AD DS site objects are stored in the Configuration container (CN=Sites, CN=Configuration, DC=*forest root domain*) and are used to achieve two primary service management tasks:

- Manage replication traffic. Typically, there are two types of network LAN connections within an enterprise environment: highly connected and less highly connected. Conceptually, a change made to AD DS should replicate immediately to other domain controllers within the highly connected network in which the change was made. However, you might not want the change to replicate immediately over a slower, more expensive, or less reliable link to another site. Instead, you might want to optimize performance, reduce costs, and manage bandwidth, you can manage replication over less highly connected segments of your enterprise. An Active Directory site represents a highly connected portion of your enterprise. When you define a site, the domain controllers within the site replicate




changes almost instantly. However, you can manage and schedule replication between sites as needed.

- Provide service localization. Active Directory sites help you localize services, including those provided by domain controllers. During logon, Windows clients are automatically directed to domain controllers in their sites. If domain controllers are not available in their sites, they are directed to domain controllers in the nearest site that can authenticate the client efficiently. Many other services such as replicated Distributed File System (DFS) resources are also site-aware to ensure that users are directed to a local copy of the resource.

What Are Subnet Objects?

Subnet objects identify the network addresses that map computers to AD DS sites. A subnet is a segment of a TCP/IP network to which a set of logical IP addresses are assigned. Because the subnet objects map to the physical network, so do the sites. A site can consist of one or more subnets. For example, if your network has three subnets in New York and two in London, you can create a site in New York and one in London, respectively, and then add the subnets to the respective sites.

 **Note:** When designing your AD DS site configuration, it is critical that you correctly map IP subnets to sites. Likewise, if the underlying network configuration changes, you must ensure that these changes are updated to reflect the current IP subnet to site mapping. Domain controllers use the IP subnet information in AD DS to map client computers and servers to the correct AD DS site. If this mapping is not accurate, AD DS operations such as logon traffic and applying Group Policies are likely to happen across WAN links and may be disrupted.

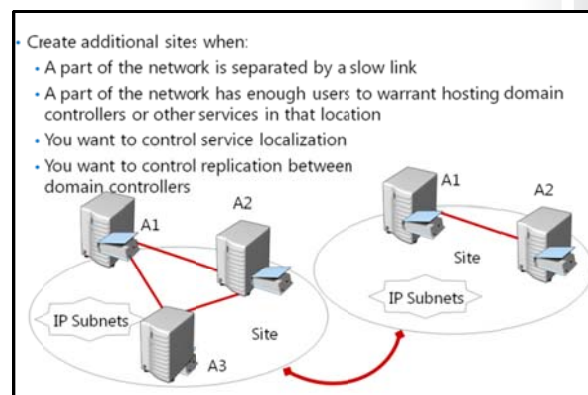
Default First Site

AD DS creates a default site when you install a forest's first domain controller. By default, this site is called *Default-First-Site-Name*. You can rename this site to a more descriptive name. When you install the forest's first domain controller, AD DS places it in the default site automatically. If you have a single site, it is not necessary to configure subnets or additional sites since all machines will be covered by the default-first-site-name default site. However, multiple sites need to have subnets associated to them as needed.

Why Implement Additional Sites?

Every Active Directory forest includes at least one site. You should create additional sites when:

- A slow link separates part of the network. As previously mentioned, a site is characterized by a location with fast, reliable, inexpensive connectivity. If two locations are connected by a slow link, you should configure each location as a separate AD DS site. A slow link typically is one that has a connection of less than 512 kilobits per second (Kbps).
- A part of the network has enough users to warrant hosting domain controllers or other services in that location. Concentrations of users can also influence your site design. If a network location has a sufficient number of users for whom the inability to authenticate would be problematic, place a domain controller in the location to support authentication within the location. After you place a domain controller or other distributed service in a location that will support those



users, you might want to manage Active Directory replication to the location or localize service use by configuring an Active Directory site to represent the location.

- You want to control service localization. By establishing AD DS sites, you can ensure that clients use domain controllers that are nearest to them for authentication, which reduces authentication latency and traffic on WAN connections. In most scenarios, each site will contain a domain controller. However, you might configure sites to localize services other than authentication, such as Distributed File System, BranchCache, and Exchange Server services. In this case, some sites might be configured without a domain controller present in the site.
- You want to control replication between domain controllers. There may be scenarios in which two well-connected domain controllers are allowed to communicate only at certain times of the day. Creating sites allows you to control how and when replication takes place between domain controllers.

Demonstration: Configuring AD DS Sites

In this demonstration, you will see how to configure AD DS sites.

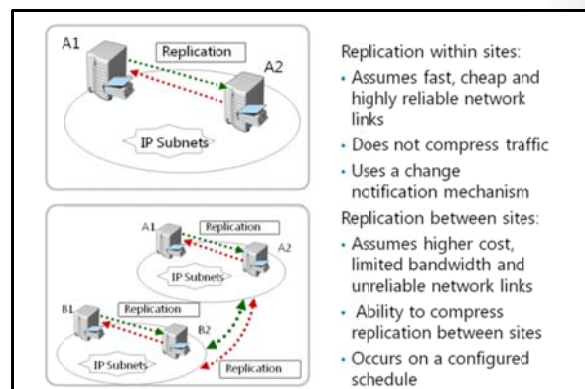
Demonstration Steps

1. From Server Manager, open Active Directory Sites and Services.
2. Rename the Default-First-Site-Name site, as needed.
3. Right-click the **Sites** node, and then click **New Site**. Specify a name, and then associate the new site with the default site link.
4. Create additional sites, as needed.
5. In the navigation pane, right-click **Subnets**, and then click **New Subnet**.
6. Provide the prefix, and then associate the IP prefix to an available site object.
7. If required, move a domain controller to the new site.

How Replication Works Between Sites

The main characteristics or assumptions about replication within sites are:

- The network connections within a site are both reliable, cheap, and have sufficient available bandwidth.
- Replication traffic within a site is not compressed, because a site assumes fast, highly reliable network connections. Not compressing replication traffic helps reduce the processing load on the domain controllers. However, uncompressed traffic may increase the network bandwidth.
- A change notification process initiates replication within a site.



- Replication within sites:
- Assumes fast, cheap and highly reliable network links
 - Does not compress traffic
 - Uses a change notification mechanism
- Replication between sites:
- Assumes higher cost, limited bandwidth and unreliable network links
 - Ability to compress replication between sites
 - Occurs on a configured schedule

The main characteristics or assumptions about replication between sites are:

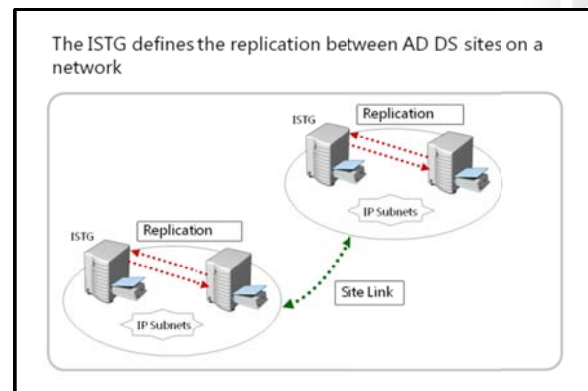
- The network links between sites have limited available bandwidth, may have a higher cost, and may not be reliable.
- Replication traffic between sites can be designed to optimize bandwidth by compressing all replication traffic. Replication traffic is compressed to 10 to 15 percent of its original size before it is transmitted. Although compression optimizes network bandwidth, it imposes an additional processing load on domain controllers, when it compresses and decompresses replication data.
- Replication between sites occurs automatically after you have defined configurable values, such as a schedule or a replication interval. You can schedule replication for inexpensive or off-peak hours. By default, changes are replicated between sites according to a schedule that you define, and not according to when changes occur. The schedule determines when replication can occur. The interval specifies how often domain controllers check for changes during the time that replication can occur.

What Is the Inter-Site Topology Generator?

When you configure multiple sites, the KCC on one domain controller in each site is designated as the site's Inter-Site Topology Generator (ISTG). There is only one ISTG per site, regardless of how many domains or other directory partitions the site has. ISTG is responsible for calculating the site's ideal replication topology.

When you add a new site to the forest, each site's ISTG determines which directory partitions are present in the new site. The ISTG then calculates how many new connection objects are necessary to replicate the new site's required information. In some networks, you might want to specify that only certain domain controllers are responsible for intersite replication. You can do this by specifying bridgehead servers. The bridgehead servers are responsible for all replication into, and out of, the site. ISTG creates the required connection agreement in its directory, and this information is then replicated to the bridgehead server. The bridgehead server then creates a replication connection with the bridgehead server in the remote site, and replication begins. If a replication partner becomes unavailable, the ISTG automatically selects another domain controller, if possible. If bridgehead servers have been manually assigned, and they become unavailable, ISTG will not automatically select other servers.

The ISTG selects bridgehead servers automatically, and creates the intersite replication topology to ensure that changes replicate effectively between bridgeheads sharing a site link. Bridgeheads are selected per partition, so it is possible that one domain controller in a site might be the bridgehead server for the schema, while another is for the configuration. However, you usually will find that one domain controller is the bridgehead server for all partitions in a site, unless there are domain controllers from other domains or application directory partitions. In this scenario, bridgeheads will be chosen for those partitions.



Overview of SRV Records for Domain Controllers

When you add a domain controller to a domain, the domain controller advertises its services by creating SRV records (also known as *locator records*) in DNS. Unlike host A records, which map host names to IP addresses, SRV records map services to host names. For example, to publish its ability to provide authentication and directory access, a domain controller registers Kerberos version 5 protocol and LDAP SRV records. These SRV records are added to several folders within the forest's DNS zones.

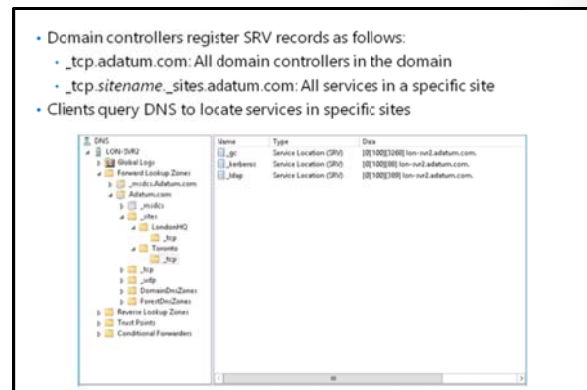
Under the domain zone, a folder exists that is named *name_tcp*. This folder contains the SRV records for all domain controllers in the domain. Additionally, under the domain zone exists a folder called *name_sites*, which contains subfolders for each site configured in the domain. Each site-specific folder contains SRV records that represent services available in the site. For example, if a domain controller is located in a site, a SRV record will be located at the path *_sites\sitename_tcp*, where *sitename* is the name of the site.

A typical SRV record contains the following information:

- The service name and port. This portion of the SRV record indicates a service with a fixed port. It does not have to be a well-known port. SRV records in Windows Server 2012 include LDAP (port 389), Kerberos (port 88), Kerberos Password protocol (KPASSWD, port 464), and global catalog services (port 3268).
- Protocol. The TCP or UDP is indicated as a transport protocol for the service. The same service can use both protocols in separate SRV records. Kerberos records, for example, are registered for both TCP and UDP. Microsoft clients use only TCP, but UNIX clients can use both UDP and TCP.
- Host name. The host name corresponds to the A record for the server hosting the service. When a client queries for a service, the DNS server returns the SRV record and associated A records, so the client does not need to submit a separate query to resolve the IP address of a service.

The service name in an SRV record follows the standard DNS hierarchy with components separated by dots. For example, a domain controller's Kerberos service is registered as: *kerberos._tcp.sitename._sites.domainname*, where:

- *domainName*: The domain or zone, for example contoso.com
- *_sites*: All sites registered with DNS
- *sitename*: The site of the domain controller registering the service
- *_tcp*: Any TCP-based services in the site
- *kerberos*: A Kerberos Key Distribution Center (KDC) that uses TCP as its transport protocol



How Client Computers Locate Domain Controllers Within Sites

When you join a Windows client to a domain and restart it, it goes through a domain controller location and registration process. The goal of this registration process is to locate the domain controller with the most efficient and closest location to the client's location based on IP subnet information.

The process for locating a domain controller is:

1. The new client queries for all domain controllers in the domain. As the new domain client restarts, it receives an IP address from a DHCP server, and is ready to authenticate to the domain. However, the client does not know where to find a domain controller. Therefore, the client queries for a domain controller by querying the `_tcp` folder, which contains the SRV records for all domain controllers in the domain.
2. The client attempts an LDAP ping to all domain controllers in a sequence. DNS returns a list of all matching domain controllers, and the client attempts to contact all of them on its first startup.
3. The first domain controller responds. The first domain controller that responds to the client examines the client's IP address, cross-references that address with subnet objects, and informs the client of the site to which the client belongs. The client stores the site name in its registry, and then queries for domain controllers in the site-specific `_tcp` folder.
4. The client queries for all domain controllers in the site. DNS returns a list of all domain controllers in the site.
5. The client attempts LDAP ping sequentially to all domain controllers in the site. The domain controller that responds first authenticates the client.
6. The client forms an affinity. The client forms an affinity with this domain controller, and then attempts to authenticate with the same domain controller in the future. If the domain controller is unavailable, the client queries the site's `_tcp` folder again, and again attempts to bind with the first domain controller that responds in the site.

Locating a domain controller occurs as follows:

- 1 New client queries for all domain controllers in the domain
- 2 Client attempts LDAP ping to find all domain controllers
- 3 First domain controller responds
- 4 Client queries for all domain controllers in the site
- 5 Client attempts LDAP ping to find all domain controllers in the site
- 6 Client stores domain controller and site name for further use
- 7 Domain controller is used for the full logon process including authentication, building the token, and building the list of GPOs to apply
 - Domain controller offline? Client queries for domain controllers in registry stored site
 - Client moved to another site? Domain controller refers client to another site

If the client moves to another site, such as the case for a mobile computer, the client attempts to authenticate to its preferred domain controller. The domain controller notices that the client's IP address is associated with a different site, and then refers the client to the new site. The client then queries DNS for domain controllers in the local site

Automatic Site Coverage

As mentioned previously, you can configure sites to direct users to local copies of replicated resources, such as shared folders replicated within a DFS namespace. There may be scenarios in which you only require service localization with no need for a domain controller located within the site. In this case, a nearby domain controller will register its SRV records in the site by using a process called site coverage.

A site without a domain controller generally is covered by a domain controller in a site with the lowest site-link cost to the site that requires coverage. You also can configure site coverage and SRV record priority manually if you want to control authentication in sites without domain controllers.



Additional Reading: For more information about how site coverage is evaluated see: <http://go.microsoft.com/fwlink/?LinkId=168550>.

Lesson 3

Configuring and Monitoring AD DS Replication

After you configure the sites that represent your network infrastructure, the next step is to determine if any additional site links are necessary to help control AD DS replication. AD DS provides several options that you can configure to control how replication occurs over site links. You also need to understand the tools that you can use to monitor and manage replication in an AD DS network environment.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe AD DS site links.
- Explain the concept of site link bridging.
- Describe Universal Group Membership caching.
- Describe how to control intersite replication.
- Configure AD DS intersite replication.
- Describe options for configuring password replication policies for RODCs.
- Configure password replication policies.
- Describe tools used for monitoring and managing replication.

What Are AD DS Site Links?

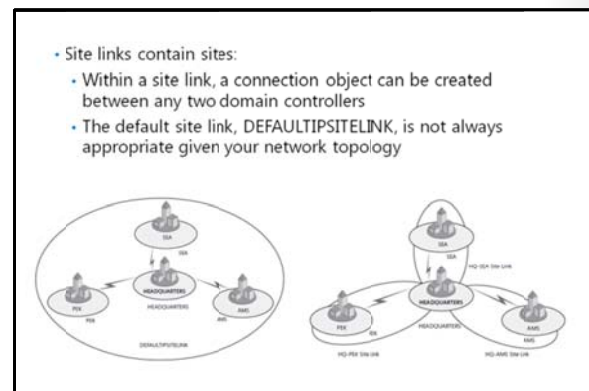
For two sites to exchange replication data, a site link must connect them. A site link is a logical path that the KCC\ISTG uses to establish replication between sites. When you create additional sites, you must select at least one site link that will connect the new site to an existing site. Unless a site link is in place, the KCC cannot make connections between computers at different sites, nor can replication occur between sites.

The important thing to remember about a site link is that it represents an available path for replication. A single site link does not control the

network routes that are used. When you create a site link and add sites to it, you are telling AD DS that it can replicate between any of the sites associated with the site link. The ISTG creates connection objects, and those objects will determine the actual replication path. Although the replication topology that the ISTG builds does replicate AD DS effectively, it might not be efficient, given your network topology.

To better understand this concept, consider the following example. When you create a forest, one site link object is created: DEFAULTIPSITELINK. By default, each new site that you add is associated with the DEFAULTIPSITELINK. Consider an organization with a data center at the headquarters and three branch offices. The three branch offices are each connected to the data center with a dedicated link. You create sites for each branch office: Seattle (SEA), Amsterdam (AMS), and Beijing (PEK). Each of the sites, including headquarters, is associated with the DEFAULTIPSITELINK site link object

Because all four sites are on the same site link, you are instructing AD DS that all four sites can replicate with each other. That means that Seattle may replicate changes from Amsterdam; Amsterdam may



replicate changes from Beijing; and Beijing may replicate changes from the headquarters, which in turn replicates changes from Seattle. In several of these replication paths, the replication traffic on the network flows from one branch through the headquarters on its way to another branch. With a single site link, you do not create a hub-and-spoke replication topology even though your network topology is hub-and-spoke.

To align your network topology with Active Directory replication, you must create specific site links. That is, you can manually create site links that reflect your intended replication topology. Continuing the preceding example, you would create three site links as follows:

- HQ-AMS includes the Headquarters and Amsterdam sites.
- HQ-SEA includes the Headquarters and Seattle sites.
- HQ-PEK includes the Headquarters and Beijing sites.

After you create site links, the ISTG will use the topology to build an intersite replication topology that connects each site, and then creates connection objects automatically to configure the replication paths. As a best practice, you should set up your site topology correctly and avoid creating connection objects manually.

What Is Site Link Bridging?

After you have created site links and the ISTG generates connection objects to replicate partitions between domain controllers that share a site link, your work might be complete. In many environments, particularly those with straightforward network topologies, site links might be sufficient to manage intersite replication. In more complex networks, however, you can configure additional components and replication properties.

Automatic Site Link Bridging

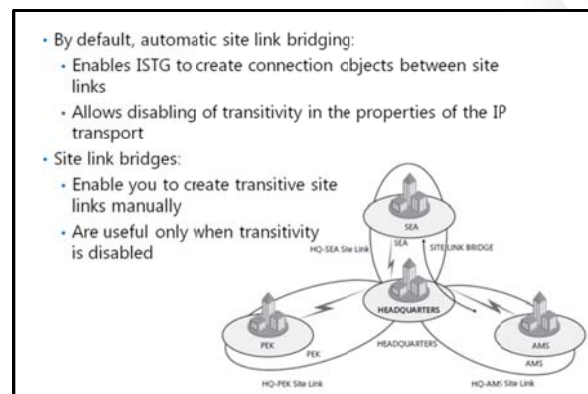
By default, all site links are bridged. For example, if the Amsterdam and Headquarters sites are linked, and the Headquarters and Seattle sites are linked, Amsterdam and Seattle are linked with a higher cost. This means, theoretically, that the ISTG could create a connection object directly between a domain controller in Seattle and a domain controller in Amsterdam when a domain controller is not available at the headquarters for replication, again working around the hub-and-spoke network topology.

You can disable automatic site-link bridging by opening the properties of the IP transport in the Inter-Site Transports container, and then clearing the **Bridge All Site Links** check box. Before you do this in a production environment, read the technical resources about replication in the Windows Server technical libraries on Microsoft TechNet at <http://technet.microsoft.com>.

Site Link Bridges

A site link bridge connects two or more site links in a way that creates a transitive link. Site link bridges are necessary only when you have cleared the **Bridge All Site Links** check box for the transport protocol. Remember that automatic site-link bridging is enabled by default, in which case, site link bridges are not required.

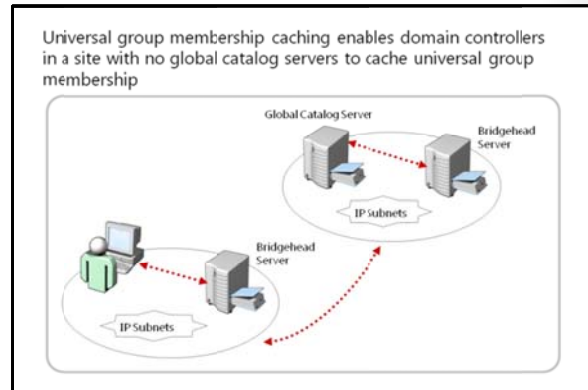
The figure on the slide illustrates the use of a site link bridge in a forest in which automatic site-link bridging has been disabled. By creating a site link bridge, AMS-HQ-SEA, that includes the HQ-AMS and



HQ-SEA site links, those two site links become transitive, so a replication connection can be made between a domain controller in Amsterdam and a domain controller in Seattle.

What Is Universal Group Membership Caching?

One of the issues that you may need to address when configuring AD DS replication is whether to deploy global catalog servers in each site. Because global catalog servers are required when users log on to the domain, deploying a global catalog server in each site optimizes the user experience. However, deploying a global catalog server in a site might result in additional replication traffic, which may be an issue if the network connection between AD DS sites has limited bandwidth. In these scenarios, you can deploy domain controllers running Windows Server 2008 or newer, and then enable universal group membership caching for the site.



How Universal Group Membership Caching Works

A domain controller in a site that has enabled universal group membership caching, stores the universal group information locally after a user attempts to log on for the first time. The domain controller obtains the user's universal group membership information from a global catalog server in another site, it then caches the information indefinitely and periodically refreshes it. The next time that the user tries to log on, the domain controller obtains the universal group membership information from its local cache without contacting a global catalog server.

By default, the universal group membership information contained by each domain controller's cache is refreshed every eight hours. To refresh the cache, domain controllers send a universal group membership confirmation request to a designated global catalog server.

You can configure universal group membership caching from the properties of the **NTDS Site Settings** node.

Controlling Intersite Replication

When you create a site link, you have a number of configuration options that you can use to help control inter-site replication. These options include:

- **Site Link Costs.** Site link costs manage the flow of replication traffic when there is more than one route for replication traffic. You can configure site link costs to indicate that a link is faster, more reliable, or is preferred. Higher costs are used for slow links, and lower costs are used for fast links. AD DS replicates by using the connection with the lowest cost. By default, all site links are configured with a cost of 100.

- Site link costs
 - Replication uses the connections with the lowest cost
- Replication
 - Polling: Downstream bridgehead polls upstream partners
 - Default: 3 hours
 - Minimum: 15 minutes
 - Recommended: 15 minutes
 - Replication schedules
 - 24 hours a day
 - Can be scheduled

- **Replication Frequency.** Intersite replication is based only on polling. By default, every three hours a replication partner polls its upstream replication partners to determine whether changes are available. This replication interval may be too long for organizations that want changes to the directory to replicate more quickly. You can change the polling interval by accessing the properties of the site link object. The minimum polling interval is 15 minutes.
- **Replication Schedules.** By default, replication occurs 24 hours a day. However, you can restrict intersite replication to specific times by changing the schedule attributes of a site link.

Demonstration: Configuring AD DS Intersite Replication

In this demonstration, you will see how to configure AD DS intersite replication.

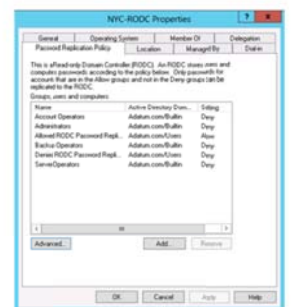
Demonstration Steps

1. From Server Manager, open Active Directory Sites and Services.
2. Rename the DEFAULTIPSITELINK as needed.
3. Right-click the site link, and then click **Properties**.
4. Modify the Cost, Replication interval, and Schedule as needed.
5. If necessary, open the properties of the IP node, and then modify the **Bridge all site links** option.

Options for Configuring Password Replication Policies for RODCs

RODCs have unique AD DS replication requirements related to cached user's credentials. They use password replication policies to determine which users' credentials might be cached on the server. If a password replication policy allows an RODC to cache a user's credentials, the RODC can process that user's authentication and service-ticket activities. If a user's credentials are not allowed to be cached on the RODC, the RODC refers the authentication and service-ticket activities to a writable domain controller.

- Password replication policies are:
- Used to determine which users' credentials should be cached on the RODC
- Determined by the Allowed List and the Denied List



To access the password replication policy, open the properties of the RODC in the Domain Controllers OU, and then click the **Password Replication Policy** tab. An RODC's password replication policy is determined by two multivalued attributes of the RODC's computer account. These attributes are known commonly as the Allowed List and the Denied List. If a user's account is on the Allowed List, the user's credentials are cached. You can include groups on the Allowed List, in which case all users who belong to the group can have their credentials cached on the RODC. If the user is on the Allowed List *and* the Denied List, the RODC does not cache the user's credentials. The Denied List takes precedence.

To facilitate the management of password replication policy, two domain local security groups are created in the Users container of AD DS. The first one, the Allowed RODC Password Replication Group, is added to the Allowed List for each new RODC. By default, the group has no members. Therefore, by default, a new RODC will not cache any user's credentials. If there are users whose credentials you want to be cached by all domain RODCs, add those users to the Allowed RODC Password Replication Group. As a best practice, you can create one Allow List per site, and configure only the users assigned to that site in the Allow List.

The second group, the Denied RODC Password Replication Group, is added to the Denied List for each new RODC. If you want to ensure that domain RODCs never cache certain users' credentials, you can add those users to the Denied RODC Password Replication Group. By default, this group contains security-sensitive accounts that are members of groups including Domain Admins, Enterprise Admins, Schema Admins, Cert Publishers, and Group Policy Creator Owners.

Demonstration: Configuring Password Replication Policies

In this demonstration, you will see how to configure password replication policies.

Demonstration Steps

1. Run Active Directory Users and Computers.
2. Precreate an RODC computer object named LON-RODC1.
3. In the **Domain Controllers** OU, open the properties of **LON-RODC1**.
4. Click the **Password Replication Policy** tab, and view the default policy.
5. Close the LON-RODC1 Properties.
6. In the Active Directory Users and Computers console tree, click the **Users** container.
7. Double-click **Allowed RODC Password Replication Group**, and then go to the **Members** tab and examine the default membership of **Allowed RODC Password Replication Group**. There should be no members by default.
8. Click **OK**.
9. Double-click Denied RODC Password Replication Group, and then go to the Members tab.
10. Click Cancel to close the Denied RODC Password Replication Group properties.

Tools for Monitoring and Managing Replication

After you have implemented your replication configuration, you must be able to monitor replication for ongoing support, optimization, and troubleshooting. Two tools are particularly useful for reporting and analyzing replication: the Replication Diagnostics tool (Repadmin.exe) and the Directory Server Diagnosis (Dcdiag.exe) tool.

The Repadmin.exe Tool

The Replication Diagnostics tool, Repadmin.exe, is a command-line tool that enables you to report the status of replication on each domain controller. The information that Repadmin.exe produces can help you spot a potential problem with replication in the forest. You can view levels of detail down to the replication metadata for specific objects and attributes, enabling you to identify where and when a problematic change was made to AD DS. You can even use Repadmin.exe to create the replication topology and force replication between domain controllers.

Repadmin.exe supports a number of commands that perform specific tasks. You can learn about each command by typing **repadmin /?:command**. Most commands require arguments. Many commands take

```

• RepAdmin.exe examples
• repadmin /showrepl Lon-dc1.adatum.com
• repadmin /showconn Lon-dc1.adatum.com
• repadmin /showobjmeta Lon-dc1 "cn=Linda Miller,ou=..."
• repadmin /kcc
• repadmin /replicate Tor-dc1 Lon-dc1 dc=adatum,dc=com
• repadmin /syncall Lon-dc1.adatum.com /A /e

• DCDiag /test.testName
• FrsEvent or DFSREvent
• Intersite
• KccEvent
• Replications
• Topology

```

a *DC_LIST* parameter, which is simply a network label (DNS, NetBIOS name, or IP address) of a domain controller. Some of the replication monitoring tasks you can perform by using Repadmin are:

- Display the replication partners for a domain controller. To display the replication connections of a domain controller, type **repadmin /showrepl *DC_LIST***. By default, Repadmin.exe shows only intersite connections. Add the **/repsto** argument to see intersite connections, as well.
- Display connection objects for a domain controller. Type **repadmin /showconn *DC_LIST*** to show the connection objects for a domain controller.
- Display metadata about an object, its attributes, and replication. You can learn a lot about replication by examining an object on two different domain controllers to find out which attributes have or have not replicated. Type **repadmin /showobjmeta *DC_LIST Object***, where *DC_LIST* indicates the domain controller(s) to query. (You can use an asterisk [*] to indicate all domain controllers.) *Object* is a unique identifier for the object, its distinguished name or GUID, for example.

You can also make changes to your replication infrastructure by using Repadmin. Some of the management tasks you can perform are:

- Launching the KCC. Type **repadmin /kcc** to force the KCC to recalculate the inbound replication topology for the server.
- Forcing replication between two partners. You can use Repadmin to force replication of a partition between a source and a target domain controller. Type **repadmin /replicate *Destination_DC_LIST Source_DC_Name Naming_Context***.
- Synchronizing a domain controller with all replication partners. Type **repadmin /syncall *DC/A /e*** to synchronize a domain controller with all its partners, including those in other sites.

The Dcdiag.exe Tool

The Directory Service Diagnosis tool, Dcdiag.exe, performs a number of tests and reports on the overall health of replication and security for AD DS. Run by itself, dcdiag.exe performs summary tests and reports the results. On the other extreme, **dcdiag.exe /c** performs almost every test. The output of tests can be redirected to files of various types, including XML. Type **dcdiag /?** for full usage information.

You can also specify one or more tests to perform using the **/test:*Test Name*** parameter. Tests that are directly related to replication include:

- FrsEvent. Reports any operation errors in the File Replication System.
- DFSREvent. Reports any operation errors in the Distributed File System Replication system.
- Intersite. Checks for failures that would prevent or delay intersite replication.
- KccEvent. Identifies errors in the knowledge consistency checker.
- Replications. Checks for timely replication between domain controllers.
- Topology. Checks that the replication topology is connected fully for all domain controllers.
- VerifyReplicas. Verifies that all application directory partitions are instantiated fully on all domain controllers hosting replicas.

Lab: Implementing AD DS Sites and Replication

Scenario

A. Datum has deployed a single AD DS domain with all the domain controllers located in the London data center. As the company has grown and added branch offices with large numbers of users, it has become apparent that the current AD DS environment is not meeting the company requirements. Users in some of the branch offices report that it can take a long time for them to log on to their computers. Access to network resources such as the company's Microsoft Exchange 2010 servers and the Microsoft SharePoint® servers can be slow, and they sporadically fail.

As one of the senior network administrators, you are responsible for planning and implementing an AD DS infrastructure that will help address the business requirements for the organization. You are responsible for configuring AD DS sites and replication to optimize the user experience and network utilization within the organization.

Objectives

- Configure the default site created in AD DS.
- Create and configure additional sites in AD DS.
- Configure and monitor replication between AD DS sites.

Lab Setup

20412A-LON-DC1

20412A-TOR-DC1

Estimated time: **60 minutes**

Virtual machines	20412A-LON-DC1 20412A-TOR-DC1
User Name	Adatum\Administrator
Password	Pa\$\$w0rd

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

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V® Manager, click **20412A-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on using the following credentials:
 - a. User name: **Adatum\Administrator**
 - b. Password: **Pa\$\$w0rd**
5. Repeat steps 2 through 4 for **20412A-TOR-DC1**.

Exercise 1: Modifying the Default Site

Scenario

A. Datum has decided to implement additional AD DS sites to optimize the network utilization for AD DS network traffic. The first step in implementing the new environment is to install a new domain controller for the Toronto site. You will then reconfigure the default site and assign appropriate IP address subnets

to the site. You have been asked to change the name of the default site to LondonHQ and associate it with the IP subnet 172.16.0.0/24, which is the subnet range used for the London head office.

The main tasks for this exercise are as follows:

1. Install the Toronto domain controller
2. Rename the default site
3. Configure IP subnets associated with the default site

► **Task 1: Install the Toronto domain controller**

1. On TOR-DC1, use Server Manager to install **Active Directory Domain Services**.
2. When the AD DS binaries have installed, use the **Active Directory Domain Services Configuration Wizard** to install and configure TOR-DC1 as an additional domain controller for Adatum.com.
3. After the server restarts, log on as **Adatum\Administrator** with the password of **Pa\$\$w0rd**.

► **Task 2: Rename the default site**

1. If necessary, on LON-DC1, open the Server Manager console.
2. Open Active Directory Sites and Services, and then rename the **Default-First-Site-Name** site to **LondonHQ**.
3. Verify that both LON-DC1 and TOR-DC1 are members of the LondonHQ site.

► **Task 3: Configure IP subnets associated with the default site**

1. If necessary, on LON-DC1, open the Server Manager console, and then open Active Directory Sites and Services.
2. Create a new subnet with the following configuration:
 - Prefix: 172.16.0.0/24
 - Site object: **LondonHQ**

Results: After completing this exercise, you will have reconfigured the default site and assigned IP address subnets to the site.

Exercise 2: Creating Additional Sites and Subnets

Scenario

The next step in implementing the AD DS site design is to configure the new AD DS site. The first site that you need to implement is the Toronto site for the North American data center. The network team in Toronto would also like to dedicate a site called TestSite in the Toronto data center. You have been instructed that the Toronto IP subnet address is 172.16.1.0/24. The test network IP subnet address is 172.16.100.0/24.

The main tasks for this exercise are as follows:

1. Create the AD DS sites for Toronto
2. Create IP subnets associated with the Toronto sites

► **Task 1: Create the AD DS sites for Toronto**

1. If necessary, on LON-DC1 open the Server Manager console, and then open Active Directory Sites and Services.

2. Create a new site with the following configuration:
 - o Name: Toronto
 - o Site link object: **DEFAULTIPSITELINK**
3. Create another new site with the following configuration:
 - o Name: TestSite
 - o Site link object: **DEFAULTIPSITELINK**

► **Task 2: Create IP subnets associated with the Toronto sites**

1. If necessary, on LON-DC1 open Active Directory Sites and Services.
2. Create a new subnet with the following configuration:
 - o Prefix: 172.16.1.0/24
 - o Site object: **Toronto**
3. Create another new subnet with the following configuration:
 - o Prefix: 172.16.100.0/24
 - o Site object: **TestSite**
4. In the navigation pane, click the **Subnets** folder. Verify that the three subnets were created and associated with their appropriate site as displayed in the details pane.

Results: After this exercise, you will have created two additional sites representing the IP subnet addresses located in Toronto.

Exercise 3: Configuring AD DS Replication

Scenario

Now that the AD DS sites have been configured for Toronto, the next step is to configure the site links to manage replication between the sites, and then to move the TOR-DC1 domain controller to the Toronto site. Currently all sites belong to DEFAULTIPSITELINK. You need to modify site linking so that LondonHQ and Toronto belong to one common site link called LON-TOR. You should configure this link to replicate every hour. Additionally, you should link the TestSite site only to the Toronto site using a site link named TOR-TEST.

Replication should not be available from the Toronto site to the TestSite during the working hours of 9 A.M. and 3 P.M. You then will use tools to monitor replication between the sites.

The main tasks for this exercise are as follows:

1. Configure site links between AD DS sites
2. Move TOR-DC1 to the Toronto site
3. Monitor AD DS site replication

► **Task 1: Configure site links between AD DS sites**

1. If necessary, on LON-DC1, open Active Directory Sites and Services.
2. Create a new IP-based site link with the following configuration:
 - o Name: TOR-TEST

- Sites: Toronto, TestSite
 - Modify the schedule to only allow replication from **Monday 9am to Friday 3pm**
3. Rename DEFAULTIPSITELINK and configure it with the following settings:
 - Name: LON-TOR
 - Sites: LondonHQ, Toronto
 - Replication: Every **60** minutes

► **Task 2: Move TOR-DC1 to the Toronto site**

1. If necessary, on LON-DC1 open Active Directory Sites and Services.
2. Move **TOR-DC1** from the **LondonHQ** site to the **Toronto** site.
3. Verify that TOR-DC1 is located under the Servers node in the Toronto site.

► **Task 3: Monitor AD DS site replication**

1. On LON-DC1, on the taskbar, click the **Windows PowerShell** button.
2. Use the following commands to monitor site replication:

```
Repadmin /kcc
```

This command recalculates the inbound replication topology for the server:

```
Repadmin /showrep1
```

Verify that the last replication with TOR-DC1 was successful:

```
Repadmin /bridgeheads
```

This command displays the bridgehead servers for the site topology:

```
Repadmin /rep1summary
```

This command displays a summary of replication tasks. Verify that no errors appear:

```
DCDiag /test:replications
```

Verify that all connectivity and replication tests pass successfully.

3. Switch to TOR-DC1, and then repeat the commands to view information from the TOR-DC1 perspective.

Results: After this exercise, you will have configured site links and monitored replication.

► **To prepare for the next module**

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

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

Module Review and Takeaways

Question: Why is it important that all subnets are identified and associated with a site in a multisite enterprise?

Question: What are the advantages and disadvantages of reducing the intersite replication interval?

Question: What is the purpose of a bridgehead server?

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Client cannot locate domain controller in its site.	
Replication between sites does not work.	
Replication between two domain controllers in the same site does not work.	

Best Practice

You should implement the following best practices when you manage Active Directory sites and replication in your environment:

- Always provide at least one or more global catalog servers per site.
- Ensure that all sites have appropriate subnets associated.
- Do not setup long intervals without replication when you configure replication schedules for intersite replication.
- Avoid using SMTP as a protocol for replication.

Module 10

Implementing Active Directory Certificate Services

Contents:

Module Overview	10-1
Lesson 1: PKI Overview	10-2
Lesson 2: Deploying CAs	10-10
Lesson 3: Deploying and Managing Certificate Templates	10-16
Lesson 4: Implementing Certificate Distribution and Revocation	10-21
Lesson 5: Managing Certificate Recovery	10-29
Lab: Implementing Active Directory Certificate Services	10-33
Module Review and Takeaways	10-41

Module Overview

Public key infrastructure (PKI) consists of several components that help you secure corporate communications and transactions. One such component is the Certification Authority (CA). You can use CAs to manage, distribute, and validate digital certificates that are used to secure information. You can install Active Directory® Certificate Services (AD CS) as a root CA or a subordinate CA in your organization. In this module, you will learn about implementing AD CS server role and certificates.

Objectives

After completing this module, you will be able to:

- Describe PKI.
- Deploy CAs.
- Deploy and manage certificate templates.
- Implement certificate distribution and revocation.
- Manage certificate recovery.

Lesson 1

PKI Overview

PKI helps you verify and authenticate the identity of each party involved in an electronic transaction. It also helps you establish trust between computers and the corresponding applications that are hosted on application servers. A common example includes the use of PKI technology to secure websites. Digital certificates are key PKI components that contain electronic credentials, which are used to authenticate users or computers. Moreover, certificates can be validated using certificate discovery, path validation, and revocation checking processes. Windows Server® 2012 supports building a certificate services infrastructure in your organization using AD CS components.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe PKI.
- Describe components of a PKI solution.
- Describe CAs.
- Describe the AD CS server role in Windows Server 2012.
- Describe new features in AD CS in Windows Server 2012.
- Explain the difference between public and private CAs.
- Describe cross-certification hierarchy.

What Is PKI?

PKI is a combination of software, encryption technologies, processes, and services that assist an organization with securing its communications and business transactions. It is a system of digital certificates, certification authorities, and other registration authorities. When an electronic transaction takes place, PKI verifies and authenticates the validity of each party involved. PKI standards are still evolving, but they are widely implemented as an essential component of electronic commerce.

PKI :

- Is a standard approach to security-based tools, technologies, processes, and services used to enhance the security of communications, applications, and business transactions
- Relies on the exchange of digital certificates between authenticated users and trusted resources

PKI provides:

- Confidentiality
- Integrity
- Authenticity
- Non-repudiation

General concepts of PKI

In general, a PKI solution relies on several technologies and components. When you plan to implement PKI, you should consider and understand the following:

- Infrastructure: The meaning in this context is the same as in any other context, such as electricity, transportation, or water supply. Each of these elements does a specific job, and has requirements that must be met for it to function efficiently. The sum of these elements allows for the efficient and safe use of PKI. Some of the elements that make up a PKI are the following:
 - A CA
 - A certificate repository
 - A registration authority

- An ability to revoke certificates
- An ability to back up, recover, and update keys
- An ability to regulate and track time
- Client-side processing

Most of these components will be discussed in later topics and lessons of this module.

- Public/Private Keys: In general, there are two methods for encrypting and decrypting data:
 - Symmetric encryption: The methods to encrypt and decrypt data are identical, or mirrors of each other. Data is encrypted by using a particular method or key. To decrypt the data, you must have the same, identical method or key. Therefore, anyone who has the key can decrypt the data. The key must remain private to maintain the integrity of the encryption.
 - Asymmetric encryption: In this case, the methods to encrypt and decrypt data are not identical or mirrors of each other. Data is encrypted by using a particular method or key. However, a different key is used to decrypt data. This is achieved by using a pair of keys. Each person gets a key pair, which consists of a public key and a private key. These keys are unique, and data that the public key encrypts can be decrypted by using the private key, and vice versa. In this situation, the keys are sufficiently different and knowing or possessing one does not allow you to determine the other. Therefore, one of the keys (public) can be made publicly available without reducing the security of the data, as long as the other key (private) remains private—hence the name Public Key Infrastructure.

Algorithms that use symmetric encryption are fast and efficient for large amount of data. However, because they use a symmetric key, they are not considered secure enough, because you always must transport the key to the other party. Alternatively, algorithms that use asymmetric encryption are secure, but very slow. Because of this, it is common to use hybrid approach, which means that data is encrypted by using symmetric encryption, while the symmetric encryption key is protected with asymmetric encryption.

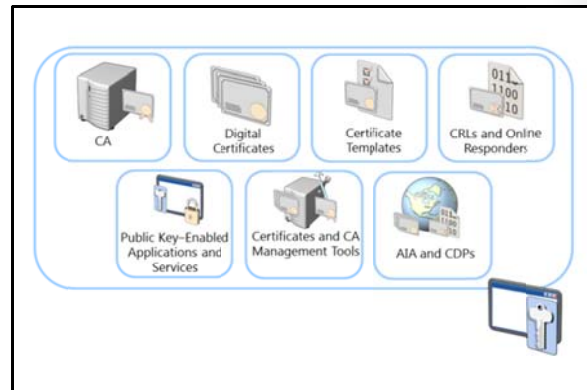
When you implement a PKI solution, your entire system, especially the security aspect, can benefit. The benefits of using PKI include:

- Confidentiality: A PKI solution enables you to encrypt both stored and transmitted data.
- Integrity: You can use PKI to sign data digitally. A digital signature identifies whether any data was modified while information was transmitted.
- Authenticity and non-repudiation: Authentication data passes through hash algorithms such as Secure Hash Algorithm 1 (SHA-1) to produce a message digest. The message digest is then digitally signed using the sender's private key to prove that the message digest was produced by the sender. Non-repudiation is digitally signed data in which the digital signature provides both proof of the integrity of signed data, and proof of the origin of data.
- Standards-based approach: PKI is standards-based, which means that multiple technology vendors are compelled to support PKI-based security infrastructures. It is based on industry standards defined in RFC 2527, "Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework."


Components of a PKI Solution

There are many components that are required to work together to provide a complete PKI solution. The PKI components in Windows Server 2012 are:

- CA: CA issues and manages digital certificates for users, services, and computers. By deploying CA, you establish the PKI in your organization.
- Digital certificates: Digital certificates are similar in function to an electronic passport. A digital certificate is used to prove the identity of the user (or other entity). Digital certificates contain the electronic credentials that are associated with a public key and a private key, which are used to authenticate users and other devices such as Web servers and mail servers. Digital certificates also ensure that software or code is run from a trusted source. Digital certificates contain various fields, such as **Subject**, **Issuer**, and **Common Name**. These fields are used to determine the specific use of the certificate. For example, a Web server certificate might contain the **Common Name** field of **web01.contoso.com**, which would make that certificate valid only for that web server. If an attempt were made to use that certificate on a web server named **web02.contoso.com**, the user of that server would receive a warning.
- Certificate templates: This component describes the content and purpose of a digital certificate. When requesting a certificate from an AD CS enterprise CA, the certificate requestor will, depending on his or her access rights, be able to select from a variety of certificate types based on certificate templates, such as User and Code Signing. The certificate template saves users from low-level, technical decisions about the type of certificate they need. In addition, they allow administrators to distinguish who might request which certificates.
- CRLs and Online Responders:
 - Certificate revocation lists (CRLs) are complete, digitally signed lists of certificates that have been revoked. These lists are published periodically and can be retrieved and cached by clients (based on the configured lifetime of the CRL). The lists are used to verify a certificate's revocation status.
 - Online Responders are part of the Online Certificate Status Protocol (OCSP) role service in Windows Server 2008 and Windows Server 2012. An Online Responder can receive a request to check for revocation of a certificate without requiring the client to download the entire CRL. This speeds up certificate revocation checking, and reduces the network bandwidth. It also increases scalability and fault tolerance, by allowing for array configuration of Online Responders.
- Public key–based applications and services: This relates to applications or services that support public key encryption. In other words, the application or services must be able to support public key implementations to gain the benefits from it.
- Certificate and CA management tools: Management tools provide command-line and GUI-based tools to:
 - Configure CAs
 - Recover archived private keys
 - Import and export keys and certificates
 - Publish CA certificates and CRLs
 - Manage issued certificates

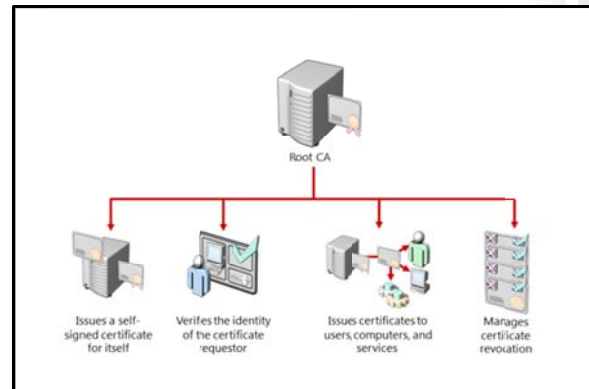


- Authority information access (AIA) and CRL distribution points (CDPs): AIA points determine the location where CA certificates can be found and validated, and CDP locations determine the points where certificate revocation lists can be found during certificate validation process. Because CRLs can become large, (depending on the number of certificates issued and revoked by a CA), you can also publish smaller, interim CRLs called *delta CRLs*. Delta CRLs contain only the certificates revoked since the last regular CRL was published. This allows clients to retrieve the smaller delta CRLs and more quickly build a complete list of revoked certificates. The use of delta CRLs also allows revocation data to be published more frequently, because the size of a delta CRL means that it usually does not require as much time to transfer as a full CRL.
- Hardware security module (HSM): A hardware security module is an optional secure cryptographic hardware device that accelerates cryptographic processing for managing digital keys. It is a high security, specialized storage that is connected to the CA for managing the certificates. An HSM is typically attached to a computer physically. This is an optional add-on in your PKI, and is most widely used in high security environments where there would be a significant impact if a key were compromised.

 **Note:** The most important component of any security infrastructure is physical security. A security infrastructure is not just the PKI implementation. Other elements—such as physical security and adequate security policies—are also important parts of a holistic security infrastructure.

What Are CAs?

A CA is a well-designed and highly trusted service in an enterprise, which provides users and computers with certificates, maintains the CRLs, and optionally responds to OSCP requests. You can install a CA in your environment by deploying the AD CS role on Windows Server 2012. When the first CA is installed, it establishes the PKI in the network, and it provides the highest point in the whole structure. You can have one or more certification authorities in one network, but only one CA can be at the highest point on the CA hierarchy (that CA is called the *root CA*, which will be discussed later in this module).



One of the main purposes of the CA is to issue certificates, revoke certificates, and publish AIA and CRL information. By doing that, the CA ensures that users, services, and computers are issued certificates that can be validated.

A CA performs multiple functions or roles in a PKI. In a large PKI, separation of CA roles among multiple servers is common. A CA provides several management tasks, including:

- Verifying the identity of the certificate requestor.
- Issuing certificates to requesting users, computers, and services.
- Managing certificate revocation.

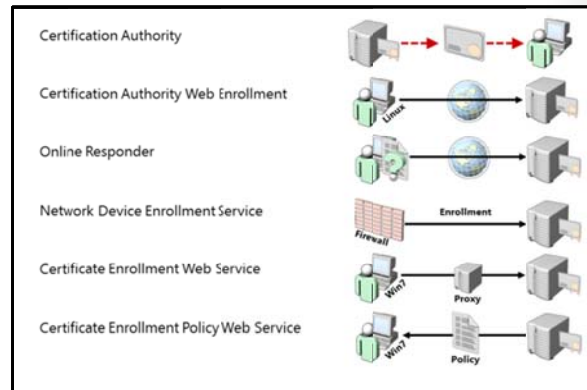
When you deploy a first CA (root CA) in your network, it issues a certificate for itself. After that, other CAs receive certificates from the first CA. You can also choose to issue a certificate for your CA by using one of public CAs.

Overview of the AD CS Server Role in Windows Server 2012

All PKI-related components are deployed as role services of the AD CS server role. This role is made up of several components that are known as role services. Each role service is responsible for a specific portion of the certificate infrastructure, while working together to form a complete solution.

Role services of the AD CS role are:

- CA. This component issues certificates to users, computers, and services. It also manages certificate validity. Multiple CAs can be chained to form a PKI hierarchy.
- CA Web enrollment. This component provides a method to issue and renew certificates for users, computers, and devices that are not joined to the domain, are not connected directly to the network, or are for users of non-Windows® operating systems.
- Online Responder. You can use this component to configure and manage OCSP validation and revocation checking. Online Responder decodes revocation status requests for specific certificates, evaluates the status of those certificates, and returns a signed response containing the requested certificate status information. Unlike in Windows Server 2008 R2, you can install Online Responder on any version of Windows Server 2012. The certificate revocation data can come from a CA on a computer that is running Windows Server 2003, Windows Server 2008, or from a non-Microsoft CA.
- Network Device Enrollment Service. With this component, routers, switches, and other network devices can obtain certificates from AD CS. On Windows Server 2008 R2, this component is only available on the Enterprise and Datacenter editions, but on Windows Server 2012, you can install this role service on any version of Windows Server.
- Certificate Enrollment Web Service. This component works as a proxy between Windows® 7 and Windows 8 client computers and the CA. This component is new to Windows Server 2008 R2 and Windows Server 2012, and requires that the Active Directory forest be at least at the Server 2008 R2 level. It enables users to connect to a CA by means of a web browser to perform the following:
 - Request, renew, and install issued certificates
 - Retrieve CRLs
 - Download a root certificate
 - Enroll over the internet or across forests (new to Windows Server 2008 R2)
- Certificate Authority Policy Web Service. This component is new to Windows Server 2008 R2 and Windows Server 2012. It enables users to obtain certificate enrollment policy information. Combined with the Certificate Enrollment Web Service, it enables policy-based certificate enrollment when the client computer is not a member of a domain, or when a domain member is not connected to the domain.



What Is New in AD CS in Windows Server 2012

Like many other Windows Server roles, AD CS is improved and enhanced in Windows Server 2012. The AD CS role in Windows Server 2012 still has the same six role services, as described in the previous topic. In addition, it now provides multiple new features and capabilities compared to previous versions.

In Windows Server 2008 R2, some of the AD CS role services require a specific Windows Server version. For example, Network Device Enrollment Service (NDES) does not work on the Windows Server 2008 Standard edition, but only on the Windows Server 2008 Enterprise edition). In Windows Server 2012, all role services are available on all Windows Server versions.

The AD CS Server role, in addition to all related role services, can run on Windows Server 2012 with full GUI, Minimal Server Interface, or on a Server Core installation. You can deploy AD CS role services in Windows Server 2012 using Server Manager, or Windows PowerShell® cmdlets, while working locally at the computer or remotely over the network.

From a management perspective, AD CS and its events, and the Best Practices Analyzer tool are now fully integrated into the Server Manager console, which means that you can access all its options directly from Server Manager. AD CS is also fully manageable by using the Windows PowerShell command-line interface.

The Windows Server 2012 version of AD CS also introduces a new certificate template version—version 4—which provides some new capabilities. This will be discussed separately in Lesson 3.

Certificate Enrollment Web Services is also enhanced in Windows Server 2012. This feature, introduced in Windows 7 and Windows Server 2008 R2, allows online certificate requests to come from untrusted Active Directory Domain Services (AD DS) domains or even from computers or devices that are not joined to a domain. AD CS in Windows Server 2012 adds the ability to renew certificates automatically for computers that are part of untrusted AD DS domains, or are not joined to a domain.

From a security perspective, AD CS in Windows Server 2012 provides the ability to require the renewal of a certificate with the same key. Windows Server 2012 also supports generating trusted platform module (TPM)-protected keys using TPM-based key storage providers (KSPs). The benefit of using a TPM-based KSP is true non-exportability of keys that are backed up by the anti-hammering mechanism of TPMs (for example, if a user enters a wrong PIN too many times). To enhance security even further, you can now force encryption of all certificate requests that come to AD CS in Windows Server 2012.

Virtual Smart Cards

Smart cards, as an option for multi-factor authentication, have been used since Windows Server 2000. They provide enhanced security over passwords, as it is much more difficult for an unauthorized user to gain and maintain access to a system. In addition, access to a smart card-protected system requires that a user both have a valid card and know the PIN that provides access to that card. By default, only one copy of the smart card exists, so only one individual can be using their login credentials at a time. In addition, a user will quickly notice if their card has been lost or stolen, especially when their card is combined with access to doors or other functions. This greatly reduces the risk window of credential theft in comparison to passwords.

- All AD CS role services run on all Windows Server versions
- Full integration with Server Manager
- Manageable through Windows PowerShell
- New certificate template version (v4)
- Support for automatic renewal of certificates for non-domain joined computers
- Enforcement of certificate renewal with same key
- Additional security for certificate requests
- Support for Virtual Smart Cards

However, implementation of smart card infrastructure has historically sometimes been too expensive. To implement smart cards, companies had to buy hardware, including smart card readers and smart cards. This cost, in some cases, prevented the deployment of multi-factor authentication.

To address these issues, Windows Server 2012 AD CS introduces a technology that provides the security of smart cards while reducing material and support costs. This is done by providing Virtual Smart Cards. Virtual Smart Cards emulate the functionality of traditional smart cards, but instead of requiring the purchase of additional hardware, they utilize technology that users already own and are more likely to have with them at all times.

Virtual Smart Cards in Windows Server 2012 leverage the capabilities of the TPM chip that is present on most of the computer motherboards produced in the past two years. Because the chip is already in the computer, there is no cost for buying smart cards and smart card readers. However, unlike traditional smart cards, where the user was in a physical possession of the card, in the Virtual Smart Card scenario, a computer (or to be more specific, TPM chip on its motherboard) acts like a smart card. By using this approach, two-factor authentication similar to traditional smart cards is achieved. A user must have his or her computer (which has been set up with the Virtual Smart Card), and also know the PIN necessary to use his or her Virtual Smart Card.

It is important to understand how Virtual Smart Cards protect private keys. Traditional smart cards have their own storage and cryptographic mechanism for protecting the private keys. In the Virtual Smart Card scenario, private keys are protected not by isolation of physical memory, but rather by the cryptographic capabilities of the TPM: all sensitive information that is stored on a smart card is encrypted using the TPM, and then stored on the hard drive in its encrypted form. Although private keys are stored on a hard drive (in encrypted form), all cryptographic operations occur in the secure, isolated environment of the TPM. Private keys never leave this environment in unencrypted form. If the hard drive of the machine is compromised in any way, private keys cannot be accessed, because they are protected and encrypted by TPM. To provide more security, you can also encrypt the drive with Windows BitLocker® Drive Encryption. To deploy Virtual Smart Cards, you need Windows Server 2012 AD CS and a Windows 8 client machine with a TPM chip on motherboard.

Public vs. Private CAs

When you are planning PKI implementation for your organization, one of the first choices you should make is between private and public CAs. It is possible for you to establish PKI by using either of these approaches. If you decide to use a private CA, then you deploy the AD CS server role, and then establish an internal PKI. If you decide to use an external PKI, you do not have to deploy any service internally.

Both approaches have advantages and disadvantages, as specified in the following table.

Internal private CAs:

- Require greater administration than external public CAs
- Cost less than external public CAs, and provide greater control over certificate management
- Are not trusted by external clients by default
- Offer advantages such as customized templates and autoenrollment

External public CAs:

- Are trusted by many external clients (web browsers, operating systems)
- Have slower certificate procurement

CA type	Advantages	Disadvantages
External Public CA	<ul style="list-style-type: none"> • Trusted by many external clients (web browsers, operating systems) • Requires minimal administration 	<ul style="list-style-type: none"> • Higher cost as compared to an internal CA • Cost is based per certificate • Certificate procurement is

CA type	Advantages	Disadvantages
		slower
Internal Private CA	<ul style="list-style-type: none"> • Provides greater control over certificate management • Lower cost as compared to a public CA • Customized templates • Autoenrollment 	<ul style="list-style-type: none"> • By default, not trusted by external clients (web browsers, operating systems) • Requires greater administration

Some organizations have started using a hybrid approach to their PKI architecture. A hybrid approach uses an external public CA for the root CA, and a hierarchy of internal CAs for distribution of certificates. This gives organizations the advantage of having their internally issued certificates trusted by external clients, while still providing the advantages of an internal CA. The only disadvantage is cost. A hybrid approach is typically the most expensive approach, because public certificates for CAs are very expensive.

In addition, you can also choose to deploy internal PKI for internal purposes such as Encrypting File System (EFS), and digital signatures. For external purposes, such as protecting web or mail servers with Secure Socket Layer (SSL), you must buy a public certificate. This approach is not very expensive, and is probably the most cost-effective solution.

What Is a Cross-Certification Hierarchy?

As cross-certification implies, in cross-certification hierarchy the root CA in each CA hierarchy provides a cross-certification certificate to the root CA in the other CA hierarchy. The other hierarchy root CA then installs the supplied certificate. By doing so, the trust flows down to all the subordinate CAs below the level where the cross-certification certificate was installed.

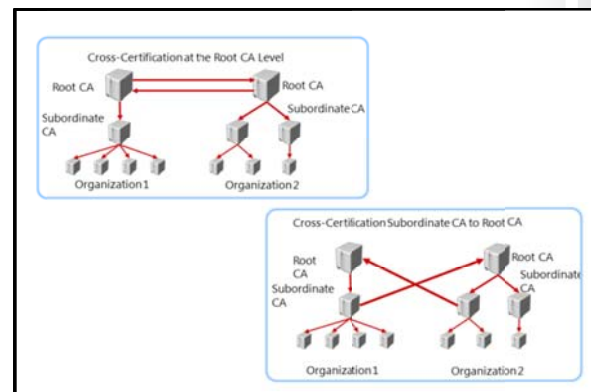
Cross-Certification Benefits

A cross-certification hierarchy provides the following benefits:

- Provides interoperability between businesses and between PKI products
- Joins disparate PKIs
- Assumes complete trust of a foreign CA hierarchy

Companies usually deploy cross-certifications to establish a mutual trust on PKI level, and also to implement some other applications that rely on PKI, such as Active Directory Rights Management Services (AD RMS).

Question: Your company is currently acquiring another company. Both companies run their own PKI. What could you do to minimize disruption and continue to provide PKI services seamlessly?



Lesson 2

Deploying CAs

The first CA that you install will be a root CA. After you install the root CA, you can optionally install a subordinate CA to apply policy restrictions and distribute certificates. You can also use a CAPolicy.inf file to automate additional CA installations and provide additional configuration settings that are not available with the standard GUI-based installation. In this lesson, you will learn about deploying CAs in the Windows Server 2012 environment.

Lesson Objectives

After completing this lesson, you will be able to:

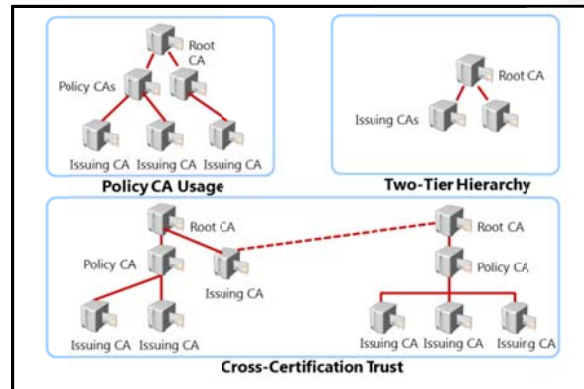
- Describe options for implementing CA hierarchies.
- Explain differences between standalone and enterprise CAs.
- Describe considerations for deploying a root CA.
- Deploy a root CA.
- Describe considerations for deploying a subordinate CA.
- Describe how to use CAPolicy.inf file for installing the CA.

Options for Implementing CA Hierarchies

When you decide to implement PKI in your organization, one of the first decisions you must make is how to design your CA hierarchy. CA hierarchy determines the core design of your internal PKI, and also determines the purpose of each CA in the hierarchy. Each CA hierarchy includes two or more CAs. Usually, the second CA (and all others after that) is deployed with a specific purpose, because only the root CA is mandatory.

The following points describe some scenarios for implementing a CA hierarchy.











- **Policy CA:** Policy CAs are a type of subordinate CA that are located directly below the root CA in a CA hierarchy. You utilize policy CAs to issue CA certificates to subordinate CAs that are located directly below the policy CA in the hierarchy. Use policy CAs when different divisions, sectors, or locations of your organization require different issuance policies and procedures.
- **Cross-certification trust:** In this scenario, two independent CA hierarchies interoperate when a CA in one hierarchy issues a CA certificate to a CA in the other hierarchy. Cross-certification trusts are discussed in more detail later in this module.
- **Two-tier hierarchy:** In a two-tier hierarchy, there is a root CA and at least one subordinate CA. In this scenario, the subordinate CA is responsible for policies, and for issuing certificates to requestors.



Standalone vs. Enterprise CAs

In Windows Server 2012, you can deploy two types of CAs: standalone CA and enterprise CA. These types are not about hierarchy, but about functionality and configuration storage. The most important difference between these two CA types is Active Directory integration and dependency. A standalone CA can work without AD DS, and does not depend on it in any way. An enterprise CA requires AD DS, but it also provides several benefits, such as autoenrollment.

The following table details the most significant differences between standalone and enterprise CAs.

	Standalone CAs		Enterprise CAs
	Must be used if any CA (root/intermediate/policy) is offline, because a standalone CA is not joined to an AD DS domain		Requires the use of AD DS
	Users provide identifying information and specify type of certificate		Can use Group Policy to propagate certificate to trusted root CA certificate store
	Does not require certificate templates		Publishes user certificates and CRLs to AD DS
	All certificate requests are kept pending until administrator approval		Issues certificates based upon a certificate template
			Supports autoenrollment for issuing certificates

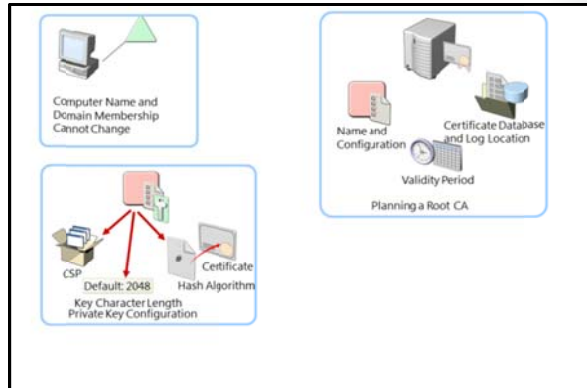
Characteristic	Standalone CA	Enterprise CA
Typical usage	A standalone CA is typically used for offline CAs, but it can be used for a CA that is consistently available on the network.	An enterprise CA is typically used to issue certificates to users, computers, and services, and is not typically used as an offline CA.
Active Directory dependencies	A standalone CA does not depend on AD DS and can be deployed in non-Active Directory environments.	An enterprise CA requires AD DS, which can be used as a configuration and registration database. An enterprise CA also provides a publication point for certificates issued to users and computers.
Certificate request methods	Users can only request certificates from a standalone CA by using a manual procedure or web enrollment.	<ul style="list-style-type: none"> • Users can request certificates from an enterprise CA using the following methods: • Manual Enrollment • Web Enrollment • Autoenrollment • Enrollment agent
Certificate issuance methods	All requests must be manually approved by a certificate administrator.	Requests can be automatically issued or denied, based on the template's discretionary access control list (DACL).

Most commonly, the root CA (which is the first CA deployed) is deployed as standalone CA, and it is taken offline after it issues a certificate for itself and for a subordinate CA. Alternatively, a subordinate CA is usually deployed as an enterprise CA, and is configured in one of scenarios described in the previous topic.

Considerations for Deploying a Root CA

Before you deploy a root CA, there are several decisions that you should make. First, you should decide if you will be deploying an offline root CA or not. Based on that decision, you will also decide if you will be deploying a standalone root CA or an enterprise root CA.

Usually, if you are deploying a single-layer CA hierarchy—which means that you deploy only a single CA—it is most common to choose enterprise root CA. However, if you are deploying a two-layer hierarchy, the most common scenario is to deploy a standalone root CA and an enterprise subordinate CA.



The next factor to consider is the operating system installation type. AD CS is supported in both the full installation and the Server Core installation scenarios. Server Core provides a smaller attack surface and less administrative overhead, and therefore should be strongly considered for AD CS in an enterprise environment. In Windows Server 2012, you can also use Windows PowerShell to deploy and manage the AD CS role.

You should also be aware that you cannot change computer names or computer domain memberships after you deploy on that computer a CA of any type, nor can you change the domain name. Therefore, it is important to determine these attributes before installing a CA.

The following table details additional considerations.

Consideration	Description
A cryptographic service provider (CSP) that is used to generate a new key	<ul style="list-style-type: none"> The default CSP is the Microsoft® Strong Cryptographic Provider. Any provider whose name starts with a number sign (#) is a cryptography Next Generation (CNG) provider.
The key character length	The default key length for the Microsoft Strong Cryptographic Provider is 2,048 characters. This is the minimum recommended value for a root CA.
The hash algorithm that is used to sign certificates issued by a CA	The default value of the hash algorithm is SHA-1.
The validity period for certificates issued by a CA	The default value for certificates is five years.
The status of the root server (online or offline)	The root server should be deployed as an offline CA. This enhances security and safeguards the root certificate (because it is not available to attack over the network).

Specifically, if you decide to deploy an offline standalone root CA, there are some specific considerations that you should have in mind:

- Before you issue a subordinate certificate from root CA, make sure that you provide at least one CDP and AIA location that will be available to all clients. This is because, by default, a standalone root CA

has the CDP and AIA located on itself. Therefore, when you take the Root CA off the network, revocation check will fail, as the CDP and AIA locations will be inaccessible. When you define these locations, you should manually copy CRL and AIA information to that location.

- Set a validity period for CRLs that root CA publishes to a long period of time (for example, one year). This means that you will have to turn on root CA once per year to publish a new CRL, and then copy it to a location that is available to clients. If you fail to do so, after the CRL on the root CA expires, revocation check for all certificates will also fail.
- Publish root CA certificate to a trusted root certification authority store on all server and client machines, by using Group Policy. You must do this manually, because a standalone CA cannot do it automatically, unlike an enterprise CA. You can also publish the root CA certificate to AD DS by using the **certutil** command-line tool.

Demonstration: Deploying a Root CA

In this demonstration, you will see how to deploy an Enterprise root CA.

Demonstration Steps

Deploy a Root CA

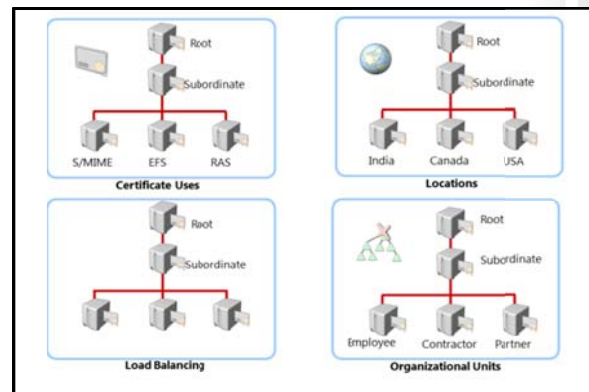
1. In Server Manager, add the **Active Directory Certificate Services** role.
2. Select the **Certification Authority** role service.
3. After the installation completes successfully, click the text **Configure Active Directory Certificate Services** on the destination server.
4. Select to install **Enterprise Root CA**.
5. Set the Key length to **4096**.
6. Name the CA **AdatumRootCA**.

Considerations for Deploying a Subordinate CA

You can use a subordinate CA to implement policy restrictions for PKI, and to distribute certificates to clients. After installing a root CA for the organization, you can install one or more subordinate CAs.

When you are using a subordinate CA to distribute certificates to users or computers that have an account in an AD DS environment, you can install the subordinate CA as an enterprise CA. Then, you can use the data from the client accounts in AD DS to distribute and manage certificates, and to publish certificates to AD DS. However, to complete this procedure, you must be a member of the local Administrators group, or have equivalent permissions. If the subordinate CA will be an enterprise CA, you also need to be a member of the Domain Admins group or have equivalent permissions.

From a security perspective, a recommended scenario would be to have an offline root standalone CA and an Enterprise subordinate CA.



A subordinate CA is usually deployed to achieve some of the following functionalities:

- Usage: You may issue certificates for a number of purposes, such as secure email and network authentication. The issuing policy for these uses may be distinct, and separation provides a basis for administering these policies.
- Organizational divisions: You may have different policies for issuing certificates, depending upon an entity's role in the organization. You can create subordinate CAs to separate and administer these policies.
- Geographic divisions: Organizations often have entities at multiple physical sites. Limited network connectivity between these sites may require individual subordinate CAs for many or all sites.
- Load balancing: If you will be using your PKI to issue and manage a large number of certificates, having only one CA can result in considerable network load for that single CA. Using multiple subordinate CAs to issue the same kind of certificates divides the network load between CAs.
- Backup and fault tolerance: Multiple CAs increase the possibility that your network will always have operational CAs available to respond to user requests.

How to Use the CAPolicy.inf File for Installation

If you want to deploy root or subordinate CA, and you want to both predefine some values for use during installation and define some additional parameters, you can use the CAPolicy.inf file to complete these steps. The CAPolicy.inf file is a plain text file that contains various settings that are used when installing the AD CS role, or when renewing the CA certificate. The CAPolicy.inf file is not required to install AD CS, but without it, the default settings will be applied, and in many cases, the default settings are insufficient. You can use the CAPolicy.inf file to configure CAs in more complicated deployments.

The CAPolicy.inf file is stored in the %Windir% folder of the root or subordinate CA. This file defines the following:

- CPS
- Object Identifier
- CRL publication intervals
- CA renewal settings
- Key size
- Certificate validity period
- CDP and AIA paths

Each CAPolicy.inf file is divided into sections, and has a simple structure, which can be described as follows:

- A section is an area in the .inf file that contains a logical group of keys. A section always appears in brackets in the .inf file.
- A key is the parameter that is to the left of the equal (=) sign.
- A value is the parameter that is to the right of the equal sign.

For example, if you want to specify Authority Information Access point in the CAPolicy.inf file, you will use following syntax:

```
[AuthorityInformationAccess]
URL=http://pki.adatum.com/CertData/adatumCA.crt
```

In this example, AuthorityInformationAccess is a section, URL is the key, and **http://pki.adatum.com/CertData/adatumCA.crt** is the value.

You can also specify some CA server settings in the CAPolicy.inf file. One example of the section that specifies these settings is :

```
[certsrv_server]
RenewalKeyLength=2048
RenewalValidityPeriod=Years
RenewalValidityPeriodUnits=5
CRLPeriod=Days
CRLPeriodUnits=2
CRLDeltaPeriod=Hours
CRLDeltaPeriodUnits=4
ClockSkewMinutes=20
LoadDefaultTemplates=True
AlternateSignatureAlgorithm=0
ForceUTF8=0
EnableKeyCounting=0
```



Note: All parameters from the previous examples are optional.

You can also use the CAPolicy.inf file when installing AD CS to define the following:

- Certification practice statement (CPS): Describes the practices that the CA uses to issue certificates. This includes the types of certificates issued, information for issuing, renewing, and recovering certificates, and other details about the CA's configuration.
- Object identifier (also known as OID): Identifies a specific object or attribute.
- CRL publication intervals: Defines the interval between publications for the base CRL.
- CA renewal settings: Defines renewal settings as follows:
 - Key size: Defines the length of the key pair used during the root CA renewal.
 - Certificate validity period: Defines the validity period for a root CA certificate.
 - CDP and AIA paths: Provides the path used for root CA installations and renewals.

Once you have created your CAPolicy.inf file, you must copy it into the %systemroot% folder of your server (for example, C:\Windows) before you install the AD CS role, or before you renew the CA certificate.



Note: The CAPolicy.inf file is processed for both the root and subordinate CA installations and renewals.

Lesson 3

Deploying and Managing Certificate Templates

Certificate templates define how a certificate can be requested and for what it can be used. Templates are configured on the CA, and they are stored in the Active Directory database. There are different versions of templates: the Microsoft Windows 2000 Server Enterprise CA supports version 1 certificate templates, the Windows Server 2003 Enterprise Edition supports versions 1 and 2 templates, and Windows Server 2008 Enterprise supports versions 1, 2, and 3 certificate templates. Windows Server 2012 introduces version 4 templates, yet still also supports all three previous template versions.

Two types of certificate template categories are users and computers, and each can be used for multiple purposes. You can assign Full Control, Read, Write, Enroll, and Autoenroll permissions to certificate templates. You can update certificate templates by modifying the original certificate template, copying a template, or superseding existing certificate templates. In this lesson, you will learn how to manage and deploy certificate templates.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe certificate templates.
- Describe certificate template versions in Windows Server 2012.
- Configure certificate template permissions.
- Configure certificate template settings.
- Describe options for updating a certificate template.
- Modify and enable a certificate template.

What Are Certificate Templates?

Certificate templates allow administrators to customize the distribution method of certificates, define certificate purposes, and mandate the type of usage allowed by a certificate. Administrators can easily create templates, and can then quickly deploy them to the enterprise by using the built-in GUI or command-line management utilities.

Associated with each certificate template is its DACL, which defines what security principals have permissions to read and configure the template, and to enroll or autoenroll for certificates based on the template. The certificate templates and their permissions are defined in AD DS and are valid within the forest. If more than one enterprise CA is running in the Active Directory forest, permission changes will affect all CAs.

When you define a certificate template, the definition of the certificate template must be available to all CAs in the forest. This is accomplished by storing the certificate template information in the Configuration naming context, where CN=Configuration, and DC=ForestRootName. The replication of this information depends on the Active Directory replication schedule, and the certificate template may not be available to all CAs until replication completes. Storage and replication are accomplished automatically.

Certificate template defines:

- Format and contents of a certificate
- Process for creating and submitting a valid certificate request
- Security principles that are allowed to read, enroll, or use Autoenroll for a certificate that will be based on the template
- Permissions required to modify a certificate template



Note: Prior to Windows Server 2008 R2, only the Enterprise version of Windows Server supported management of certificate templates. In Windows Server 2008 R2 and Windows Server 2012, you can also manage certificate templates in the Standard editions.

Certificate Template Versions in Windows Server 2012

Windows Server 2012 Certification Authority supports four versions of certificate templates. Certificate templates versions 1, 2 and 3 are legacy from previous versions of Windows Server, while version 4 is new to Windows Server 2012.

Certificate template versions correspond to the Windows Server operating system version. Windows 2000 Server, Windows Server 2003, Windows Server 2008, and Windows Server 2012 correspond to version 1, version 2, version 3, and version 4 respectively.

Version 1:	<ul style="list-style-type: none"> • Introduced in Windows Server 2000, provided for backward compatibility in later versions • Created by default when a CA is installed • Cannot be modified (except for permissions) or removed, but can be duplicated to become version 2 or 3 templates (which can then be modified)
Version 2:	<ul style="list-style-type: none"> • Default template introduced with Windows Server 2003 • Allows customization of most settings in the template • Several preconfigured templates are provided when a CA is installed
Version 3:	<ul style="list-style-type: none"> • Supports advanced Suite B cryptographic settings • Includes advanced options for encryption, digital signatures, key exchange, and hashing • Only supports Windows Server 2008 and Windows Server 2008 R2 servers • Only supports Windows Vista and Windows 7 client computers
Version 4:	<ul style="list-style-type: none"> • Available only for Windows Server 2012 and Windows 8 clients • Supports both CSPs and KSPs • Supports renewal with the same key

Aside from corresponding with Windows Server operating system versions, certificate template versions also have some functional differences as follows:

- Windows 2000 Advanced Server operating system provides support for version 1 certificate templates. The only modification allowed to version 1 templates is changing permissions to either allow or disallow enrollment of the certificate template. When you install an enterprise CA, version 1 certificate templates are created by default. As of July 13, 2010, Windows 2000 Server is no longer supported by Microsoft.
- Windows Server 2003 Enterprise Edition operating systems provide support for version 1 and version 2 templates. You can customize several settings in the version 2 templates. The default installation provides several preconfigured version 2 templates. You can add version 2 templates based on the requirements of your organization. Alternatively, you can duplicate a version 1 certificate template to create a new version 2 of the template. You can then modify and secure the newly created version 2 certificate template. When new templates are added to a Windows Server 2003 enterprise CA, they are version 2 by default.
- Windows Server 2008 Enterprise operating systems bring support for new, version 3 certificate templates. Additionally, support for version 1 and version 2 is provided. Version 3 certificate templates support several features of a Windows Server 2008 enterprise CA, such as CNG. CNG provides support for Suite B cryptographic algorithms such as elliptic curve cryptography (ECC). In Windows Server 2008 Enterprise, you can duplicate default version 1 and version 2 templates to bring them up to version 3. Windows Server 2008 provides two new certificate templates by default: Kerberos Authentication and OCSP Response Signing. In Windows Server 2008 R2, the Standard version was also able to support certificate templates. When you use version 3 certificate templates, you can use CNG encryption and hash algorithms for the certificate requests, issued certificates, and protection of private keys for key exchange and key archival scenarios.
- Windows Server 2012 operating systems provide support for version 4 certificate templates, and for all other versions from earlier editions of Windows Server. These certificate templates are available only to Windows Server 2012 and Windows 8. To help administrators separate what features are supported by which operating system version, the **Compatibility** tab was added to the certificate template properties tab. It marks options as unavailable in the certificate template properties, depending upon the selected operating system versions of certificate client and CA. Version 4

certificate templates also support both CSPs and KSPs. They can also be configured to require renewal with a same key.






Upgrading certificate templates is a process that applies only in situations where the CA has been upgraded from Windows Server 2008 or 2008 R2 to Windows Server 2012. After the upgrade, you can upgrade the certificate templates by launching the CA Manager console and accepting the upgrade prompt by clicking **Yes**.

Configuring Certificate Template Permissions

To configure certificate template permissions, you need to define the DACL for each certificate template in the **Security** tab. The permissions that are assigned to a certificate template will define which users or groups can read, modify, enroll, or autoenroll for that certificate template.

You can assign the following permissions to certificate templates:

- **Full Control:** The **Full Control** permission allows a security principal to modify all attributes of a certificate template, which includes permissions for the certificate template itself. It also includes permission to modify the security descriptor of the certificate template.
- **Read:** The **Read** permission allows a user or computer to view the certificate template when enrolling for certificates. The **Read** permission is also required by the certificate server to find the certificate templates in AD DS.
- **Write:** The **Write** permission allows a user or computer to modify the attributes of a certificate template, which includes permissions assigned to the certificate template itself.
- **Enroll:** The **Enroll** permission allows a user or computer to enroll for a certificate based on the certificate template. However, to enroll for a certificate, you must also have **Read** permissions for the certificate template.
- **Autoenroll:** The **Autoenroll** permission allows a user or computer to receive a certificate through the autoenrollment process. However, the **Autoenroll** permission requires the user or computer to also have both **Read** and **Enroll** permissions for a certificate template.

Permission	Description
 Full Control	Allows a designated user, group, or computer to modify all attributes including ownership and permissions.
 Read	Allows a designated user, group, or computer to read the certificate in AD DS when enrolling
 Write	Allows a designated user, group, or computer to modify all the attributes except permissions
 Enroll	Allows a designated user, group, or computer to enroll for the certificate template
 Autoenroll	Allows a designated user, group, or computer to receive a certificate through the Autoenrollment process



As a best practice, you should assign certificate template permissions to global or universal groups only. This is because the certificate template objects are stored in the configuration naming context in AD DS. You cannot assign permissions by using domain local groups that are found within an Active Directory domain. You should never assign certificate template permissions to individual user or computer accounts.

As a best practice, keep the **Read** permission allocated to the Authenticated Users group. This permission allocation allows all users and computers to view the certificate templates in AD DS. This permission assignment also allows the CA that is running under the System context of a computer account to view the certificate templates when assigning certificates.

Configuring Certificate Template Settings

Besides configuring security settings for certificate templates, you can also configure several other settings for each template. Be aware however, that the number of configurable options depends on the certificate template version. For example, version 1 certificate templates do not allow modification of any settings, except for security, while certificate templates from higher versions allow you to configure most of the available options. Windows Server 2012 provides several default certificate templates for purposes that include code signing (for digitally signing software), EFS (for encrypting data), and the ability for users to log on with a smart card. To customize a template for your company, duplicate the template and then modify the certificate configuration.

For each certificate template, you can customize several settings, such as: validity time, purpose, CSP, private key exportability, and issuance requirements

Category	Single Purpose Example	Multiple Purposes Example
 Users	<ul style="list-style-type: none"> • Basic EFS • Authenticated session • Smart card logon 	<ul style="list-style-type: none"> • Administrator • User • Smart card user
 Computers	<ul style="list-style-type: none"> • Web server • IPsec 	<ul style="list-style-type: none"> • Computer • Domain controller

For example, you can configure the following:

- Format and content of a certificate based on the certificate's intended use



Note: . The intended use of a certificate may relate to users or to computers, based on the types of security implementations that are required to use the PKI.

- Process of creating and submitting a valid certificate request
- CSP supported
- Key length
- Validity period
- Enrollment process or enrollment requirements

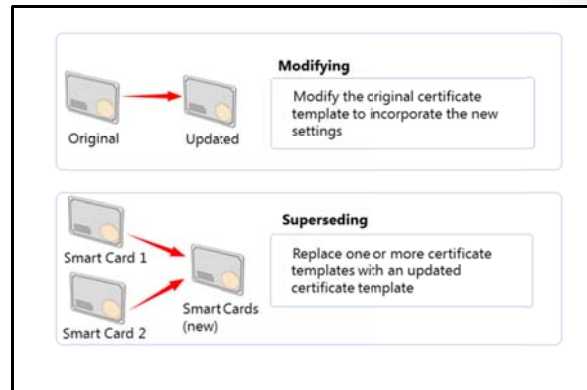
You can also define certificate purpose in certificate settings. Certificate templates can have the following purposes:

- **Single Purpose:** A single purpose certificate serves a single purpose, such as allowing users to log on with a smart card. Organizations utilize single purpose certificates in cases where the certificate configuration differs from other certificates that are being deployed. For example, if all users will receive a certificate for smart card logon but only a couple of groups will receive a certificate for EFS, organizations will generally keep these certificates and templates separate to ensure that users only receive the required certificates.
- **Multiple Purposes:** A multi-purpose certificate serves more than one purpose (often unrelated) at the same time. While some templates (such as the User template) serve multiple purposes by default, organizations will often modify templates to serve additional purposes. For example, if a company intends on issuing certificates for three purposes, those purposes can be combined into a single certificate template to ease the administrative effort and maintenance.

Options for Updating a Certificate Template

The CA hierarchy in most organizations has one certificate template for each job function. For example, there may be a certificate template for file encryption and another for code signing. Additionally, there may be a few templates that cover functions for most of the common groups of subjects.

As an IT administrator, you may need to modify an existing certificate template because of incorrect settings or other issues in the original certificate template. You may also need to merge multiple existing certificate templates into a single template.



You can update a certificate template by either modifying the template, or superseding the existing template:

- **Modify the original certificate template:** To modify a certificate template of version 2, 3, or 4, you need to make changes and then apply them to that template. After this, any certificate issued by a CA based on that certificate template will include the modifications that you made.
- **Supersede existing certificate templates:** The CA hierarchy of an organization may have multiple certificate templates that provide the same or similar functionality. In such a scenario, you can supersede or replace the multiple certificate templates by using a single certificate template. You can make this replacement in the Certificate Templates console by designating that a new certificate template supersedes, or replaces, the existing certificate templates.

Demonstration: Modifying and Enabling a Certificate Template

In this demonstration, you will see how to modify and enable a certificate template.

Demonstration Steps

Modify and enable a certificate template

1. On LON-SVR1, open the Certificate Templates console.
2. Review the list of available templates.
3. Open Properties of IPsec certificate template and review available settings.
4. Duplicate the **Exchange User** certificate template. Name it **Exchange User Test1**, and then configure it to supersede the **Exchange User** template.
5. Allow Authenticated Users to enroll for the Exchange User Test1 template.
6. Publish the template on LON-SVR1.

Lesson 4

Implementing Certificate Distribution and Revocation

One step in deploying PKI in your organization will be to define methods for certificate distribution and enrollment. In addition, during the certificate management process, there will be times that you may need to revoke certificates. There may be a number of reasons for revoking certificates, such as if a key becomes compromised, or if someone leaves the organization. You need to ensure that network clients can determine which certificates are revoked before accepting authentication requests. To ensure scalability and high availability, you can deploy the AD CS Online Responder, which can be used to provide certificate revocation status. In this lesson, you will learn about methods for certificate distribution and certificate revocation.

Lesson Objectives

After completing this lesson, you will be able to:





- Describe options for certificate enrollment.
- Describe how autoenrollment works.
- Describe the Restricted Enrollment Agent.
- Explain how to configure the Restricted Enrollment Agent.
- Describe the Network Device Enrollment Service.
- Explain how certificate revocation works.
- Describe considerations for publishing AIAs and CDPs.
- Describe an Online Responder.
- Configure Online Responder.

Options for Certificate Enrollment

In Windows Server 2012, several methods can be used to enroll for a user or computer certificate. The use of these methods depends on specific scenarios. For example, autoenrollment will probably be used to mass-deploy certificates to a large number of users or computers, while manual enrollment will be used for certificates dedicated just to specific security principals.

The following list describes the different enrollment methods, and when to use them:

- **Autoenrollment:** Using this method, the administrator defines the permissions and the configuration of a certificate template. These definitions help the requestor to automatically request, retrieve, and renew certificates without end-user interaction. This method is used for AD DS domain computers. The certificate must be configured for autoenrollment through Group Policy.
- **CA Web enrollment:** Using this method, you can enable a website CA so that users can obtain certificates. To use CA Web enrollment, you must install Internet Information Server (IIS) and the web enrollment role on the CA of AD CS. To obtain a certificate, the requestor logs on to the website, selects the appropriate certificate template, and then submits a request. The certificate is issued

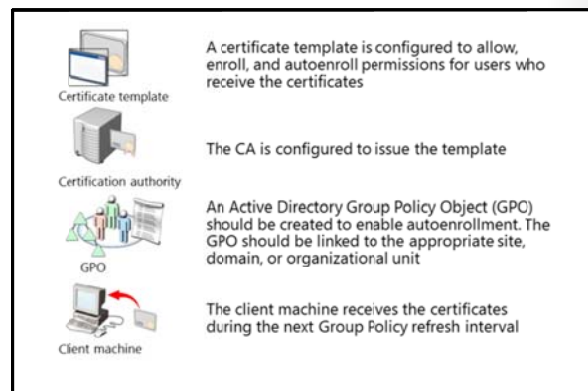
Method	Use
 Autoenrollment	<ul style="list-style-type: none"> • To automate the request, retrieval, and storage of certificates for domain-based computers
 Manual enrollment	<ul style="list-style-type: none"> • To request certificates by using the Certificates console or Certreq.exe when the requestor cannot communicate directly with the CA
 Web enrollment	<ul style="list-style-type: none"> • To request certificates from a website that is located on a CA • To issue certificates when autoenrollment is not available
 Enroll on behalf	<ul style="list-style-type: none"> • To provide IT staff with the right to request certificates on behalf of another user (Enrollment Agent)

automatically if the user has the appropriate permissions to enroll for the certificate. The CA Web enrollment method should be used to issue certificates when autoenrollment cannot be used. This can happen in the case of an Advanced Certificate request. However, there can also be cases where autoenrollment can be used for certain certificates, but not for all certificates.

- **Manual enrollment:** Using this method, the private key and a certificate request are generated on a device, such as a web service or a computer. The certificate request is then transported to the CA to generate the certificate that is requested. The certificate is then transported back to the device for installation. Use this method when the requestor cannot communicate directly with the CA, or if the device does not support autoenrollment.
- **Enrollment on behalf (Enrollment Agent):** Using this method, a CA administrator creates an Enrollment Agent account for a user. The user with Enrollment Agent rights can then enroll for certificates on behalf of other users. For example, use this method if you need to allow a manager to preload logon certificates of new employees on smart cards

How Does Autoenrollment Work?

One of the most common methods for deploying certificates in an Active Directory environment is to use autoenrollment. This method provides an automated way to deploy certificates to both users and computers within the PKI. You can use autoenrollment in environments that meet specific requirements, such as the use of certificate templates and Group Policy in AD DS. It is important to note, however, that you cannot use autoenrollment with a standalone CA. You must have an enterprise CA available to make use of autoenrollment.



You can use autoenrollment to deploy public key–based certificates automatically to users and computers in an organization. The Certificate Services administrator duplicates a certificate template, and then configures the permissions to allow Enroll and Autoenroll permissions for the users who will receive the certificates. Domain-based Group Policies, such as computer-based and user-based policies, can activate and manage autoenrollment.

By default, Group Policy is applied when you restart computers, or at logon for users. Also by default, Group Policy is refreshed every 90 minutes on domain members. This Group Policy setting is named Certificate Services Client - Auto-Enrollment.

An internal timer triggers autoenrollment every eight hours after the last autoenrollment activation. The certificate template might specify user interaction for each request. For such a request, a pop-up window appears approximately 60 seconds after the user logs on.

Many certificates can be distributed without the client even being aware that enrollment is taking place. These include most types of certificates that are issued to computers and services, as well as many certificates issued to users.

To enroll clients automatically for certificates in a domain environment, you must:

- Have membership in Domain Admins or Enterprise Admins, (or equivalent), which is the minimum required to complete this procedure.
- Configure a certificate template with Autoenroll permissions.

- Configure an autoenrollment policy for the domain.

What Is Credential Roaming?

Credential Roaming allows organizations to store certificates and private keys in AD DS, separately from application state or configuration information.

Credential Roaming uses existing logon and autoenrollment mechanisms to download certificates and keys to a local computer whenever a user logs on and, if desired, remove them when the user logs off. In addition, the integrity of these credentials is maintained under any conditions, such as when certificates are updated, or when users log on to more than one computer at a time. This avoids the scenario where a user is autoenrolled for a certificate on each new machine to which he or she logs on.

Credential Roaming is triggered any time a private key or certificate in the user's local certificate store changes, whenever the user locks or unlocks the computer, and whenever Group Policy is refreshed.

All certificate-related communication between components on the local computer and between the local computer and AD DS is signed and encrypted. Credential Roaming is supported in Windows 7 and newer Windows operating systems.

What Is the Restricted Enrollment Agent?

In earlier versions of Windows Server CA, such as Windows Server 2003, it is not possible to permit an Enrollment Agent to enroll only a certain group of users. As a result, every user with an Enrollment Agent certificate is able to enroll on behalf of any user in an organization.

The Restricted Enrollment Agent allows you to limit the permissions for users who are designated as Enrollment Agents, to enroll for smart card certificates on behalf of other users. The Restricted Enrollment Agent is a functionality that was introduced in the Windows Server 2008 Enterprise operating system.

The Restricted Enrollment Agent:

- Limits permissions of the enrollment agent
- Requires Windows Server 2008 Enterprise or Windows Server 2012 CA
- Uses version 3 or version 4 certificate templates

Typically, one or more authorized individuals within an organization are designated as Enrollment Agents. The Enrollment Agent needs to be issued an Enrollment Agent certificate, which enables the agent to enroll for smart card certificates on behalf of users. Enrollment agents are typically members of corporate security, IT security, or help desk teams, because these individuals have already been entrusted with safeguarding valuable resources. In some organizations, such as banks that have many branches, help desk and security workers might not be conveniently located to perform this task. In this case, designating a branch manager or other trusted employee to act as an Enrollment Agent is required to enable smart card credentials to be issued from multiple locations.

On a Windows Server 2012 CA, the restricted Enrollment Agent features allow an Enrollment Agent to be used for one or many certificate templates. For each certificate template, you can choose on behalf of which users or security groups the Enrollment Agent can enroll. You cannot constrain an Enrollment Agent based on a certain Active Directory organizational unit (OU) or container. Instead, you must use security groups.



Note: Using restricted Enrollment Agents will affect the performance of the CA. To optimize performance, you should minimize the number of accounts that are listed as Enrollment Agents. You minimize the number of accounts in the Enrollment Agent's permissions list. As a best practice, use group accounts in both lists instead of individual user accounts.

Demonstration: Configuring the Restricted Enrollment Agent

In this demonstration, you will see how to configure the Restricted Enrollment Agent.

Demonstration Steps

Configure the Restricted Enrollment Agent

1. On LON-SVR1, open the Certificate Templates console.
2. Configure **Allie Bellew** permissions to enroll for an Enrollment Agent certificate.
3. Publish the Enrollment Agent certificate template.
4. Log on to LON-CL1 as **Adatum\Allie** with the password **Pa\$\$w0rd**.
5. Open a MMC console and add the Certificates snap-in.
6. Request the Enrollment Agent certificate.
7. Switch to LON-SVR1, and open the properties of AdatumRootCA.
8. Configure the Restricted Enrollment Agent so that Allie can only issue certificates based on the **User** template, and only for the **Marketing** security group.

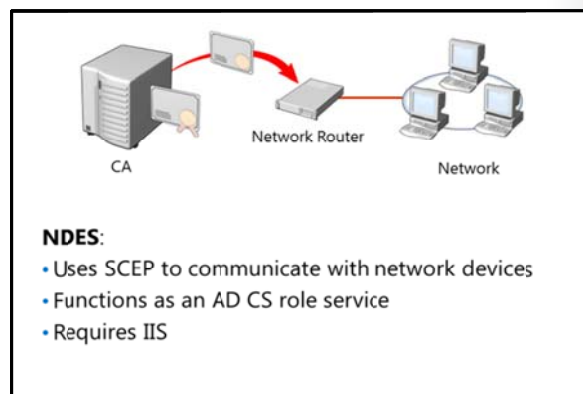
What Is Network Device Enrollment Service?

The Network Device Enrollment Service (NDES) is the Microsoft implementation of Simple Certificate Enrollment Protocol (SCEP). SCEP is a communication protocol that makes it possible for software that is running on network devices such as routers and switches—which cannot otherwise be authenticated on the network—to enroll for X.509 certificates from a CA.

You can use NDES as an Internet Server API (ISAPI) filter on IIS to perform the following functions:

- Create and provide one-time enrollment passwords to administrators.
- Retrieve awaiting requests from the CA.
- Collect and process SCEP enrollment requests for the software that runs on network devices.

This feature applies to organizations that have PKIs with one or more Windows Server 2012–based CAs, and that want to enhance the security of their network devices. Port security, based on 802.1x, requires certificates be installed on switches and access points. Secure Shell (SSH), instead of Telnet, requires a certificate on the router, switch, or access point. NDES is the service that allows administrators to install certificates on devices using SCEP.



Adding support for NDES can enhance the flexibility and scalability of an organization's PKI. Therefore, this feature should interest PKI architects, planners, and administrators.

Before installing NDES, you must decide:

- Whether to set up a dedicated user account for the service, or use the Network Service account.
- The name of the NDES registration authority and what country/region to use. This information is included in any SCEP certificates that are issued.
- The CSP to use for the signature key that is used to encrypt communication between the CA and the registration authority.
- The CSP to use for the encryption key that is used to encrypt communication between the registration authority and the network device.
- The key length for each of these keys.

In addition, you need to create and configure the certificate templates for the certificates that are used in conjunction with NDES.

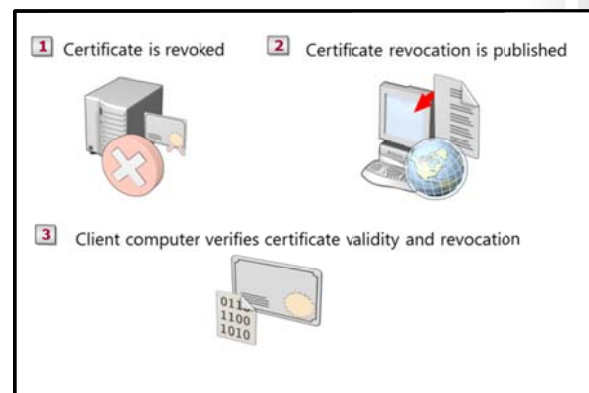
Installing NDES on a computer creates a new registration authority and deletes any preexisting registration authority certificates on the computer. Therefore, if you plan to install NDES on a computer where another registration authority has already been configured, any pending certificate requests should be processed and any unclaimed certificates should be claimed before you install NDES.

How Does Certificate Revocation Work?

Revocation is the process in which you disable validity of one or more certificates. By initiating the revoke process, you actually publish a certificate thumbprint in the corresponding CRL.

An overview of the certificate revocation life cycle is outlined as follows:

- A certificate is revoked from the CA MMC snap-in. During revocation, a reason code and a date and time are specified. This is optional, but recommended to fill.
- The CRL is published using the CA MMC snap-in (or the scheduled revocation list is published automatically based on the configured value). CRLs can be published in AD DS, some shared folder location, or on a website.
- When Windows client computers are presented with a certificate, they use a process to verify revocation status by querying the issuing CA. This process determines whether the certificate is revoked, and then presents the information to the application requesting the verification. The Windows client computer uses one of the CRL locations specified in certificate to check its validity.



The Windows operating systems include a CryptoAPI, which is responsible for the certificate revocation and status checking processes. The CryptoAPI utilizes the following phases in the certificate checking process:

- **Certificate Discovery:** Certificate discovery collects CA certificates, AIA information in issued certificates, and details of the certificate enrollment process.
- **Path validation:** Path validation is the process of verifying the certificate through the CA chain (or *path*) until the root CA certificate is reached.

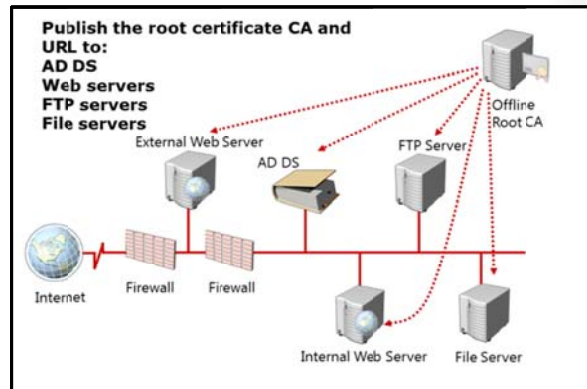
- Revocation checking: Each certificate in the certificate chain is verified to ensure that none of the certificates are revoked.
- Network retrieval and caching: Network retrieval is performed by using OCSP. CryptoAPI is responsible for checking the local cache first for revocation information and if there is no match, making a call using OCSP, which is based on the URL provided by the issued certificate.

Considerations for Publishing AIAs and CDPs

When you are managing and issuing certificates, it is very important to properly configure certificate extensions that are used to verify the certificate of the CA, and the certificate that is being used by the user. These extensions, called AIA and CDP, are part of each certificate. They must point to proper locations, or PKI may not function correctly.

What Is AIA?

AIA addresses are the URLs—addresses that uniquely identify each location on the Internet or intranet—in the certificates that a CA issues. These addresses tell the verifier of a certificate where to retrieve the CA's certificate. AIA access URLs can be HTTP, File Transfer Protocol (FTP), Lightweight Directory Address Protocol (LDAP), or FILE addresses.



What Is CDP?

CDP is a certificate extension that indicates from where the certificate revocation list for a CA can be retrieved. It can contain none, one, or many HTTP, FILE, or LDAP URLs.

AIA and CDP Publishing

If you use only an online CA, these values are configured by default locally on the CA. However, if you want to deploy an offline root CA or if you want to publish AIA and CDP to an internet facing location, you must reconfigure these values so that they apply to all certificates issued by the root CA. The AIA and CDP extensions define where client applications can locate AIA and CDP information for the root CA. The formatting and publishing of AIA and CDP extension URLs are generally the same for root CAs and subordinate CAs. You can publish the root CA certificate and the CRL to the following locations:

- Active Directory
- Web servers
- File Transfer Protocol (FTP) servers
- File servers

Publication Points

To ensure accessibility to all computers in the forest, publish the offline root CA certificate and the offline root CAs CRL to Active Directory by using the **Certutil** command. This places the root CA certificate and CRL in the Configuration naming context, which Active Directory replicates to all domain controllers in the forest.

For computers that are not members of Active Directory, place the CA certificate and CRL on web servers by using the HTTP protocol. Locate the web servers on the internal network, and also on the external

network if external client computers (or internal clients from external networks) require access. This is very important if you are using internally issued certificates outside your company.

You can also publish certificates and CRLs to ftp:// and FILE:// URLs, but it is recommended that you use only LDAP and HTTP URLs, because they are the most widely supported URL formats for interoperability purposes. The order in which you list the CDP and AIA extensions is important because the certificate chaining engine searches the URLs sequentially. Place the LDAP URL first in the list if your certificates are mostly used internally.

What Is an Online Responder?

By using OCSP, an Online Responder provides clients with an efficient way to determine the revocation status of a certificate. OCSP submits certificate status requests using HTTP.

Clients access CRLs to determine the revocation status of a certificate. CRLs might be large, and clients might utilize a large amount of time to search through these CRLs. An Online Responder can dynamically search these CRLs for the clients and respond only to the requested certificate.

You can use a single Online Responder to determine revocation status information for certificates that are issued by a single CA, or by multiple CAs. However, you can use more than one Online Responder to distribute CA revocation information.

You can install an Online Responder on any computer that runs Windows Server 2008 Enterprise or Windows Server 2012. You should install an Online Responder and a CA on different computers. The following operating systems can use Online Responder for validation of certificate status:

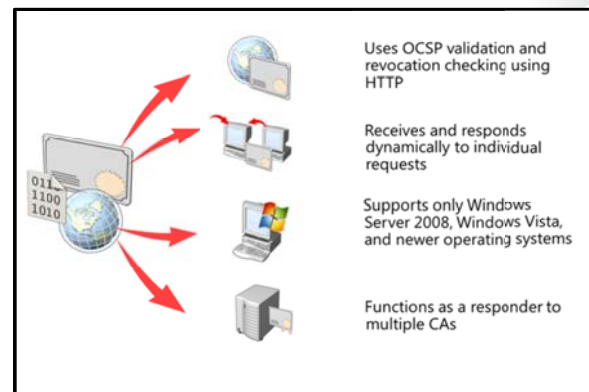
- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012
- Windows Vista
- Windows 7
- Windows 8

For scalability and high availability, you can deploy the Online Responder in a load-balanced array using Network Load Balancing (NLB), which processes certificate status requests. You can monitor and manage each member of the array independently. To configure the Online Responder, you must use the Online Responder management console.

You must configure the CAs to include the URL of the Online Responder in the AIA extension of issued certificates. The OCSP client uses this URL to validate the certificate status. You must also issue the OCSP Response Signing certificate template, so that the Online Responder can also enroll that certificate.

How to Install and Configure Online Responder

You can install Online Responders on computers that are running Windows Server 2008 R2 or Windows Server 2012. You should install Online Responders after the CAs, but before issuing any client certificates.



The certificate revocation data is derived from a published CRL that can come from a CA on a computer that is running Windows Server 2008 or newer, or Windows Server 2003, or from a non-Microsoft CA.

Before configuring a CA to support the Online Responder service, the following must be present:

- IIS must be installed on the computer during the Online Responder installation. The correct configuration of IIS for the Online Responder is installed automatically when you install an Online Responder.
- An OCSP Response Signing certificate template must be configured on the CA, and autoenrollment used to issue an OCSP Response Signing certificate to the computer on which the Online Responder will be installed.
- The URL for the Online Responder must be included in the AIA extension of certificates issued by the CA. This URL is used by the Online Responder client to validate certificate status.

After an Online Responder has been installed, you need to create a revocation configuration for each CA and CA certificate that is served by an Online Responder. A revocation configuration includes all of the settings that are needed to respond to status requests regarding certificates that have been issued using a specific CA key. These configuration settings include:

- CA certificate. This certificate can be located on a domain controller, in the local certificate store, or imported from a file.
- Signing certificate for the Online Responder. This certificate can be selected automatically for you, selected manually (which involves a separate import step after you add the revocation configuration), or you can use the selected CA certificate.
- Revocation provider that will provide the revocation data used by this configuration. This information is entered as one or more URLs where the valid base and delta CRLs can be obtained.

Demonstration: Configuring an Online Responder

In this demonstration, you will see how to configure an Online Responder.

Demonstration Steps

Configure an Online Responder

1. On LON-SVR1, use Server Manager to add an Online Responder role service to the existing AD CS role.
2. Configure a new AIA distribution location on **AdatumRootCA** to be **http://lon-svr1/ocsp**.
3. On **AdatumRootCA**, publish the OCSP Response signing certificate template, and allow Authenticated users to enroll.
4. Open the Online Responder Management console.
5. Add revocation configuration for **AdatumRootCA**.
6. Enroll for OCSP Response signing certificate.
7. Ensure that the revocation configuration status displays as working.

Lesson 5

Managing Certificate Recovery

Certificate or key recovery is one of the most important management tasks during the certificate life cycle. You use a key archival and recovery agent for data recovery if you lose your public and private keys. You can also use automatic or manual key archival and key recovery methods to ensure that you can gain access to data in the event that your keys are lost. In this lesson, you will learn how to manage key archival and recovery in AD CS in Windows Server 2012 AD CS.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the process of key archival and recovery.
- Configure Automatic Key Archival.
- Configure CA for Key Archival.
- Explain key recovery.
- Recover a lost key.

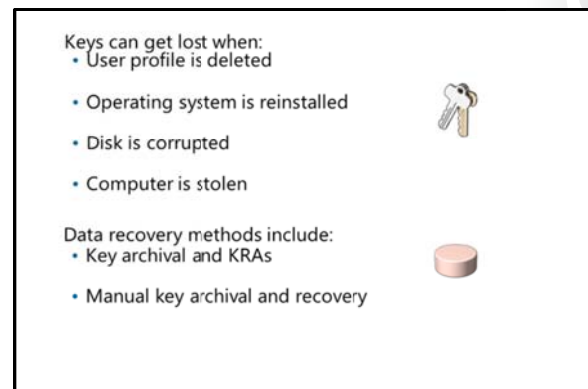
Overview of Key Archival and Recovery

If you lose your public and private keys, you will not be able to access any data that is encrypted by using the certificate's public key. This data can include Encrypting File System (EFS) and Secure/Multipurpose Internet Mail Extensions (S/MIME). Therefore, archival and recovery of public and private keys are important.

Conditions for Losing Keys

You may lose key pairs due to the following conditions:

- User profile is deleted or corrupted. A CSP encrypts a private key and stores the encrypted private key in the local file system and registry in the user profile folder. Deletion or corruption of the profile results in the loss of the private key material.
- Operating system is reinstalled. When you reinstall the operating system, the previous installations of the user profiles are lost, including the private key material.
- Disk is corrupted. If the hard disk becomes corrupted and the user profile is unavailable, the private key material is lost automatically.
- Computer is stolen. If a user's computer is stolen, the user profile with the private key material is unavailable.



Key Archival and Recovery Agents

You use key archival and Key Recovery Agents (KRA) for data recovery. You can ensure that CA administrators can recover private keys by archiving them. KRAs are designated users who are able to retrieve the original certificate, private key, and public key that were used to encrypt the data, from the CA database. A specific certificate template is applied to a KRA. When you enable key archival in a version 2 certificate template, the CA encrypts and stores that private key in its database. In situations where the

CA has stored the subject's private key in the CA database, you can use key recovery to recover a corrupted or lost key.

During the key recovery process, the certificate manager retrieves the encrypted file that contains the certificate and private key from the CA database. Next, a KRA decrypts the private key from the encrypted file and returns the certificate and private key to the user.

Security for Key Archival

When you have a configured CA to issue a KRA certificate, any user with **Read** and **Enroll** permission on the KRA certificate template can enroll and become a KRA. As a result, Domain Admins and Enterprise Admins receive permission by default. However, you must ensure the following:

- Only trusted users are allowed to enroll for this certificate.
- The KRA's recovery key is stored in a secure manner.
- The server where the keys are archived is in a separate physical secure location.

Understanding Key Archival and Recovery

Key recovery implies that the private key portion of a public-private key pair may be archived and recovered. Private key recovery does not recover any data or messages. It merely enables a user to retrieve lost or damaged keys, or for an administrator to assume the role of a user for data access or data recovery purposes. In many applications, data recovery cannot occur without first performing key recovery.

The key recovery procedure is as follows:

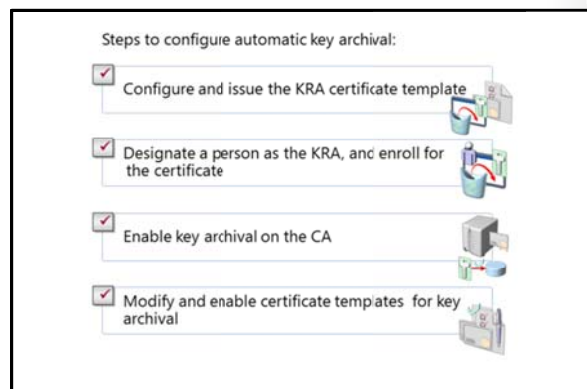
1. The user requests a certificate from a CA and provides a copy of the private key as part of the request. The CA, which is processing the request, archives the encrypted private key in the CA database and issues a certificate to the requesting user.
2. The issued certificate can be used by an application such as EFS to encrypt sensitive files.
3. If, at some point, the private key is lost or damaged, the user can contact the company's Certificate Manager to recover the private key. The Certificate Manager, with the help of the KRA, recovers the private key, stores it in a protected file format, and sends it back to the user.
4. After the user stores the recovered private key in the user's local keys store, it once again can be used by an application such as EFS to decrypt previously encrypted files or to encrypt new ones.

Configuring Automatic Key Archival

Before you can use key archival, you must perform several configuration steps. The key archival feature is not enabled by default, and you should configure both CA and certificate templates for key archival and key recovery.

The following steps describe the automatic key archival process:

1. Configure the KRA certificate template. Only Enterprise Administrators or Domain Administrators are allowed to request a KRA certificate. If you want to enroll some other user with a KRA certificate, you must specify it on the template DACL.



2. Configure Certificate Managers:
 - a. CA enforces a person to be a Certificate Manager, if defined. The Certificate Manager usually holds a private key for valid KRA certificates. By default, the CA Administrator is a Certificate Manager for all users, except for cases with another explicit definition. However, as a best practice, you should separate these two roles if possible.
 - b. A CA Officer is defined as a Certificate Manager. This user has the security permission to issue and manage certificates. The security permissions are configured on a CA in the Certification Authority MMC snap-in, in the **CA Properties** dialog box, from the **Security** tab.
 - c. A KRA is not necessarily a CA Officer or a Certificate Manager. These roles may be segmented as separate roles. A KRA is a person who holds a private key for a valid KRA certificate.
3. Enable KRA:
 - a. Log on as Administrator of the server, or as CA Administrator if role separation is enabled.
 - b. In the CA console, right-click the CA name, and then click **Properties**. To enable key archival, on the **Recovery Agents** tab, click **Archive the key**.
 - c. By default, the CA uses one KRA. However, you must first select the KRA certificate for the CA to begin archival by clicking **Add**.
 - d. The system finds valid KRA certificates, and then displays available KRA certificates. These are generally published to AD DS by an enterprise CA during enrollment. KRA certificates are stored under the KRA container in the Public Key Services branch of the configuration partition in AD DS. Because CA issues multiple KRA certificates, each KRA certificate will be added to the multivalued user attribute of the CA object.
 - e. Select one certificate, and then click **OK**. Ensure that you have selected the intended certificate.
 - f. After you have added one or more KRA certificates, click **OK**. KRA certificates are only processed at service start.
4. Configure user templates:
 - a. In the Certificate Templates MMC, right-click the key archival template, and then click **Properties**.
 - b. To always enforce key archival for the CA, in the **Properties** dialog box, on the **Request Handling** tab, select the **Archive subject's encryption private key** check box. In Windows Server 2008 or later CAs, select the **Use advanced symmetric algorithm to send the key to the CA** option.

Demonstration: Configuring CA for Key Archival

In this demonstration, you will see how to configure automatic key archival.

Demonstration Steps

Configure automatic key archival

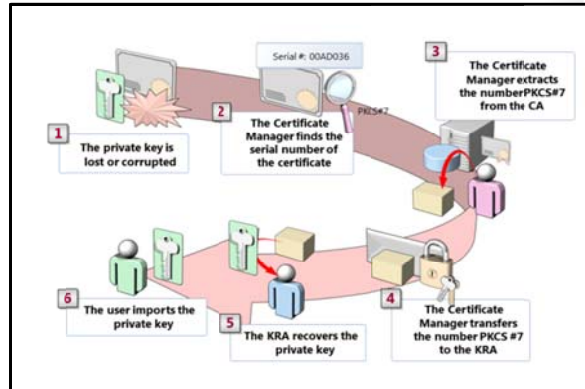
1. Configure **adatumRootCA** to issue Key Recovery Agent certificates without approval.
2. Enroll Administrator for Key Recovery Agent certificate.
3. Configure **adatumRootCA** to use certificate enrolled in step 2 as Key Recovery Agent.
4. Configure **Exchange User Test 1** certificate template to allow key archival.

5. Configure **adatumRootCA** to allow key archival.

Recovering a Lost Key

Key recovery consists of several steps, and you must strictly follow the procedure to recover archived keys. The procedure for key recovery is as follows:

1. Find recovery candidates. You will require two pieces of information to perform key recovery. First, the Certificate Manager or the CA Administrator locates the correct certificate entry in the CA database. Then, the Certificate Manager or the CA Administrator obtains the serial number of the correct certificate entry and the KRA certificate required for key recovery.
2. Retrieve PKCS #7 BLOB from the database. This is the first half of the key recovery step. A Certificate Manager or a CA Administrator retrieves the correct BLOB from the CA database. The certificate and the encrypted private key to be recovered are present in PKCS #7 BLOB. The private key is encrypted alongside the public key of one or more KRAs.
3. Recover key material and save to PKCS #12 (.pfx). This is the second half of the key recovery step. The holder of one of the KRA private keys decrypts the private key to be recovered. In addition, the holder generates a password-protected .pfx file that contains the certificate and private key.



Import recovered keys. The password-protected .pfx file is delivered to the end user. This user imports the .pfx file into the local user certificate store. Alternatively, the KRA or an administrator can perform this part of the procedure on behalf of the user.

Demonstration: Recovering a Lost Private Key (optional)

In this demonstration, you will see how to recover a lost private key.

Demonstration Steps

Recover a lost private key

1. Enroll Administrator for Exchange User Test1 certificate.
2. Delete the certificate from Administrator personal store to simulate key loss.
3. On LON-SVR1 in CA console, retrieve the serial number of lost certificate.
4. Use command **Certutil -getkey <serialnumber> outputblob** to generate blob file.
5. Use command **Certutil -recoverkey outputblob recover.pfx**, to recover the private key.
6. Import the private key back to administrator personal store.

Lab: Implementing Active Directory Certificate Services

Scenario

As A. Datum Corporation has expanded, its security requirements have also increased, and the security department is particularly interested in enabling secure access to critical web sites, and in providing additional security for features such as EFS, smart cards, and the Windows 7 and Windows 8 DirectAccess feature. To address these and other security requirements, A. Datum Corporation has decided to implement a PKI using the AD CS role in Windows Server 2012.

As one of the senior network administrators at A. Datum Corporation, you are responsible for implementing the AD CS deployment. You will be deploying the CA hierarchy, developing the procedures and process for managing certificate templates, and deploying and revoking certificates.

Objectives

- Deploy a standalone root CA, and an enterprise subordinate CA.
- Configure certificate templates.
- Configure certificate enrollment.
- Configure certificate revocation.
- Configure and perform private key archival and recovery.

Lab Setup

Estimated Time: 120 minutes

- 20412A-LON-DC1
- 20412A-LON-SVR1
- 20412A-LON-SVR2
- 20412A-LON-CA1
- 20412A-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 **20412A-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on using the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
5. Repeat steps 2 and 3 for **20412A-LON-SVR1**, **20412A-LON-SVR2**, **20412A-LON-CA1** and **20412A-LON-CL1**. Do not log on until instructed to do so.

Exercise 1: Deploying a standalone root CA

Scenario

A. Datum Corporation wants to start using certificates for various purposes, and you now need to install the appropriate CA infrastructure. Because they are using AD DS with Windows Server 2012 AD DS, you decided to implement the AD CS role. When you were reviewing available designs, you decided to

implement a standalone root CA. This CA will be taken offline after it issues a certificate for a subordinate CA.

The main tasks for this exercise are as follows:

1. Install the Active Directory® Certificate Services (AD CS) server role on non-domain joined server.
2. Configure a new certificate revocation location.

► **Task 1: Install the Active Directory® Certificate Services (AD CS) server role on non-domain joined server**

1. Log on to **LON-CA1** as **Administrator** using the password **Pa\$\$w0rd**.
2. Use the Add Roles and Features Wizard to install the **Active Directory Certificate Services** role.
3. After installation completes successfully, click the text **Configure Active Directory Certificate Services** on the destination server.
4. Configure the AD CS role as a standalone root CA. Name it **AdatumRootCA**.
5. Set the key length to **4096**, accept all other values as default.

► **Task 2: Configure a new certificate revocation location**

1. On LON-CA1, open the Certification Authority console.
2. Open the Properties window for **AdatumRootCA**.
3. Configure new locations for CDP to be on **http://lon-svr1.adatum.com/CertData/<CaName><CRLNameSuffix><DeltaCRLAllowed>.crl**
4. Select options: **Include in the CDP extensions of issued certificates** and **Include in CRLs. Clients use this to find Delta CRL locations.**
5. Configure new locations for AIA to be on **http://lon-svr1.adatum.com/CertData /<ServerDNSName>_<CaName><CertificateName>.crt**
6. Select the **Include in the AIA extension of issued certificates** check box.
7. Publish the certificate revocation list on LON-CA1.
8. Export the root CA certificate, and copy the .cer file to **\\lon-svr1\C\$**.
9. Copy the content of folder **C:\Windows\System32\CertSrv\CertEnroll** to **\\lon-svr1\C\$**.

Results: After completing this exercise, you will have installed and configured a standalone root CA.

Exercise 2: Deploying an Enterprise Subordinate CA

Scenario

After deploying the standalone root CA, the next step is to deploy an enterprise subordinate CA. A. Datum Corporation wants to use an enterprise subordinate CA to utilize AD DS integration. In addition, because root CA is standalone, you want to publish its certificate to all clients.

The main tasks for this exercise are as follows:

1. Install and configure AD CS role on LON-SVR1.
2. Install a subordinate Certification Authority (CA) certificate.
3. Publish the RootCA certificate through Group Policy.

► **Task 1: Install and configure AD CS role on LON-SVR1**

1. Log on to **LON-SVR1** as **Adatum\Administrator** with the password of **Pa\$\$w0rd**.
2. Install the **Active Directory Certificate Services** role on LON-SVR1. Include the Certification Authority and Certification Authority Web Enrollment role services.
3. After installation is successful, click **Configure Active Directory Certificate Services** on the destination server.
4. Select the **Certification Authority** and **Certification Authority Web Enrollment** role services.
5. Configure LON-SVR1 to be an **Enterprise CA**.
6. Configure the CA Type to be a **Subordinate CA**.
7. For the CA Name type **Adatum-IssuingCA**.
8. Save the request file to the local drive.

► **Task 2: Install a subordinate Certification Authority (CA) certificate**

1. On LON-SVR1, install the RootCA.cer certificate in the Trusted Root Certification Authority store.
2. Navigate to Local Disk (C:) and copy the **AdatumRootCA.crl** and **LON-CA1_AdatumRootCA.crt** files to **C:\inetpub\wwwroot\CertData**.
3. Copy the **LON-SVR1.Adatum.com_Adatum-IssuingCA.req** request file to **\\lon-ca1\C\$**.
4. Switch to LON-CA1.
5. From the Certification Authority console on LON-CA1, submit a new certificate request, by using .req file that you copied in step 3.
6. Issue the certificate and export it to p7b format with complete chain. Save the file to **\\lon-svr1\C\$\SubCA.p7b**.
7. Switch to LON-SVR1.
8. Install the SubCA certificate on LON-SVR1 using the Certification Authority console.
9. Start the service.

► **Task 3: Publish the RootCA certificate through Group Policy**

1. On LON-DC1, from Server Manager open the Group Policy Management Console.
2. Edit the Default Domain Policy.
3. Publish the **RootCA.cer** file from **\\lon-svr1\C\$** to Trusted Root Certification Authorities store in Computer Configuration\Policies\Windows Settings\Security Settings\Public Key Policies.

Results: After completing this exercise, you will have deployed and configured an enterprise subordinate CA

Exercise 3: Configuring Certificate Templates

Scenario

After deploying the CA infrastructure, the next step is to deploy the certificate templates that are required in the organization. For the beginning, A. Datum Corporation wants to implement a new Web server certificate and implement smart card certificates for users. They also want to implement new certificates on the LON-SVR2 web server.

The main tasks for this exercise are as follows:

1. Create a new template based on the Web server template.
2. Create a new template for users that includes smart card logon.
3. Configure the templates so they can be issued.
4. Update the Web server certificate on the LON-SVR2 Web Server.

► **Task 1: Create a new template based on the Web server template**

1. On LON-SVR1, from the Certification Authority console, open the Certificate Templates Console.
2. Duplicate the **Web Server** template.
3. Create a new template and name it **Adatum Web Server**.
4. Configure validity for **3 years**.
5. Configure the private key as exportable.

► **Task 2: Create a new template for users that includes smart card logon**

1. In the Certificate Templates Console, duplicate the **User** certificate template.
2. Name the new template **Adatum Smart Card User**.
3. On the **Subject Name** tab, clear both the **Include e-mail name in subject name** and the **E-mail name** check boxes.
4. Add **Smart Card Logon** to Application Policies of the new certificate template.
5. Configure this new template to supersede the **User** template.
6. Allow Authenticated Users to **Read**, **Enroll**, and **Autoenroll** for this certificate.
7. Close the Certificate Templates Console.

► **Task 3: Configure the templates so they can be issued**

- Configure LON-SVR1 to issue certificates based on the **Adatum Smart Card User** and **Adatum Web Server** templates.

► **Task 4: Update the Web server certificate on the LON-SVR2 Web Server**

1. Log on to **LON-SVR2** as **Adatum\Administrator** with the password of **Pa\$\$w0rd**.
2. Refresh the Group Policy and restart server if needed.
3. From Server Manager, open the **Internet Information Services (IIS) Manager**.
4. Enroll for a domain certificate using the following parameters:
 - Common name: **lon-svr2.adatum.com**
 - Organization: **Adatum**
 - Organizational Unit: **IT**
 - City/locality: **Seattle**
 - State/province: **WA**
 - Country/region: **US**
5. Create HTTPS binding for Default Web Site, and associate it with new certificate.

Results: After completing this exercise, you will have created and published new certificate templates.

Exercise 4: Configuring Certificate Enrollment

Scenario

The next step in implementing the PKI at A. Datum Corporation is configuring certificate enrollment. A. Datum wants to enable different options for distributing the certificates. Users should be able to enroll automatically, and smart card users should get their smart cards from Enrollment Agents. Adatum has delegated enrollment agent rights for the Marketing department group to Allie Bellew.

The main tasks for this exercise are as follows:

1. Configure autoenrollment for users.
2. Verify autoenrollment.
3. Configure the Enrollment Agent for smart card certificates.

► Task 1: Configure autoenrollment for users

1. On LON-DC1, open Group Policy Management.
2. Edit the Default Domain Policy.
3. Navigate to **User Configuration**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, and then click to highlight **Public Key Policies**.
4. Enable the **Certificate Services Client – Auto-Enrollment** option, and enable **Renew expired certificates, update pending certificates, and remove revoked certificates** and **Update certificates that use certificate templates**.
5. Enable the Certificate Services Client – Certificate Enrollment Policy.
6. Close Group Policy Management Editor and Group Policy Management console.

► Task 2: Verify autoenrollment

1. On LON-SVR1, open Windows PowerShell and use `gpupdate /force` to refresh Group Policy.
2. Open an mmc.exe console and add the Certificates snap-in focused on the user account.
3. Verify that you have been issued a certificate based on the **Adatum Smart Card User** template.

► Task 3: Configure the Enrollment Agent for smart card certificates

1. On LON-SVR1, from the Certification Authority console, open the Certificate Templates console.
2. Allow **Allie Bellew** to enroll for an Enrollment Agent certificate.
3. Publish the Enrollment Agent certificate template.
4. Log on to **LON-CL1** as **Allie**, and enroll for an Enrollment Agent certificate.
5. On LON-SVR1, open properties of **Adatum-IssuingCA**, and configure Restricted Enrollment Agent so that Allie can only issue certificates based on **Adatum Smart Card User**, for security group **Marketing**.

Results: After completing this exercise, you will have configured and verified autoenrollment for users, and configured an enrollment agent for smart cards.

Exercise 5: Configuring Certificate Revocation

Scenario

As part of configuring the certificate infrastructure, A. Datum Corporation wants to configure revocation components on newly established CAs. You will configure CRL and Online Responder components.

The main tasks for this exercise are as follows:

1. Configure Certified Revocation List (CRL) distribution.
2. Install and configure an Online Responder.

► Task 1: Configure Certified Revocation List (CRL) distribution

1. On LON-SVR1, in the Certification Authority console, right-click **Revoked Certificates**, and then click **Properties**.
2. Set the CRL publication interval to **1 Day**, and set the Delta CRL publication interval to **1 hour**.
3. Review CDP locations on **Adatum-IssuingCA**.

► Task 2: Install and configure an Online Responder

1. On LON-SVR1, use Server Manager to add an Online Responder role service to the existing AD CS role.
2. When the message displays that installation succeeded, click **Configure Active Directory Certificate Services on the destination server**. Configure the online responder.
3. On **LON-SVR1**, open the Certification Authority console.
4. Configure the new AIA distribution location on Adatum-IssuingCA to be <http://lon-svr1/ocsp>.
5. On Adatum-IssuingCA, publish the OCSP Response signing certificate template, and allow Authenticated users to enroll.
6. Open the Online Responder Management console.
7. Add revocation configuration for Adatum-IssuingCA.
8. Enroll for an OCSP Response signing certificate.
9. Ensure that revocation configuration is working.

Results: After completing this exercise, you will have configured certificate revocation settings.

Exercise 6: Configuring Key Recovery

Scenario

As a part of establishing a PKI, you want to configure and test procedures for recovery of private keys. You want to assign a KRA certificate for an administrator, and configure CA and specific certificate templates to allow key archiving. In addition, you want to test a procedure for key recovery.

The main tasks for this exercise are as follows:

1. Configure the CA to issue Key Recovery Agent (KRA) certificates.
2. Acquire the KRA certificate.
3. Configure the CA to allow key recovery.
4. Configure a custom template for key archival.
5. Verify key archival functionality.

► **Task 1: Configure the CA to issue Key Recovery Agent (KRA) certificates**

1. On LON-SVR1, in the Certification Authority console, right-click the **Certificates Templates** folder, and then click **Manage**.
2. In the Certificates Templates console, open the **Key Recovery Agent certificate properties** dialog box.
3. On the **Issuance Requirements** tab, clear the **CA certificate manager approval** check box.
4. On the **Security** tab, notice that only Domain Admins and Enterprise Admins groups have the **Enroll** permission.
5. Right-click the **Certificates Templates** folder, and enable the **Key Recovery Agent** template.

► **Task 2: Acquire the KRA certificate**

1. Create an MMC console window that includes having the Certificates snap-in for the current user loaded.
2. Use the Certificate Enrollment Wizard to request a new certificate, and enroll the KRA certificate.
3. Refresh the console window, and view the KRA in the personal store.

► **Task 3: Configure the CA to allow key recovery**

1. On LON-SVR1, in the Certification Authority console window, open the **Adatum-IssuingCA Properties** dialog box.
2. On the **Recovery Agents** tab, click **Archive the key**, and then add the certificate by using the **Key Recovery Agent Selection** dialog box.
3. Restart Certificate Services when prompted.


► **Task 4: Configure a custom template for key archival**

1. On LON-SVR1, open the Certificates Templates console.
2. Duplicate the User template, and name it **Archive User**.
3. On the **Request Handling** tab, set the option for the Archive subject's encryption private key. Using the archive key option, the KRA can obtain the private key from the certificate store.
4. Click the **Subject Name** tab, clear the **E-mail name** and **Include e-mail name in subject name** check boxes.
5. Add the **Archive User template** as a new certificate template to issue.

► **Task 5: Verify key archival functionality**

1. Log on to **LON-CL1** as **Adatum\Aidan**, using the password **Pa\$\$wOrd**.
2. Create an MMC console window that includes the Certificates snap-in.
3. Request and enroll a new certificate based on the **Archive User** template.
4. From the personal store, locate the Archive User certificate.
5. Delete the certificate for Alan Brewer to simulate a lost key.
6. Switch to LON-SVR1.
7. Open the Certification Authority console, expand **Adatum-IssuingCA**, and then click **Issued Certificates** store.
8. In the Certificate Authority Console, note the serial number of the certificate that has been issued for Alan Brewer.

9. On LON-SVR1, open a command prompt, and type:
certutil -getkey <serial number> outputblob

 **Note:** Replace serial number with the serial number that you wrote down.

10. Verify that the **Outputblob** file has appeared in the **C:\Users\Administrator** folder.
11. To convert the Outputblob file into an importable .pfx file, at the command prompt, type **Certutil-recoverkey outputblob aidan.pfx**.
12. Enter the password **Pa\$\$w0rd** for the certificate.
13. Verify the creation of the recovered key in the C:\Users\Administrator folder.
14. Cut and paste the **aidan.pfx** file to the root of drive C on LON-CL1.
15. Switch to LON-CL1, and import the **aidan.pfx** certificate.
16. Verify that the certificate appears in the Personal store.

Results: After completing this exercise, you will have implemented key archival, and tested private key recovery.

► To prepare for the next module

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

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20412A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20412A-LON-CL1**, **20412A-LON-SVR1**, **20412A-LON-CA1** and **20412A-LON-SVR2**.

Lab Review

Question: Why is it not recommended to install just an Enterprise root CA?

Question: What is the main benefit of OCSP over CRL?

Question: What must you do to recover private keys?

Module Review and Takeaways

Question: What are some reasons that an organization would utilize PKI?

Question: What are some reasons that an organization would use an enterprise root CA?

Question: List the requirements to use autoenrollment for certificates.

Question: What are the steps to configure an Online Responder?

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
The location of the CA certificate that is specified in the authority information access extension is not configured to include the certificate name suffix. Clients may not be able to locate the correct version of the issuing CA's certificate to build a certificate chain, and certificate validation may fail.	
CA is not configured to include CRL distribution point locations in the extensions of issued certificates. Clients may not be able to locate a CRL to check the revocation status of a certificate, and certificate validation may fail.	
CA was installed as an enterprise CA, but Group Policy settings for user autoenrollment have not been enabled. An enterprise CA can use autoenrollment to simplify certificate issuance and renewal. If autoenrollment is not enabled, certificate issuance and renewal may not occur as expected.	

Real-world Issues and Scenarios

Contoso, Ltd wants to deploy PKI for supporting and securing several services. They have decided to use Windows Server 2012 Certificate Services as a platform for PKI. Certificates will be primarily used for EFS, digital signing, and for Web servers. Because documents that will be encrypted are important, it is crucial to have a disaster recovery strategy in case of key loss. In addition, clients that will access secure parts of the company website must not receive any warning in their browsers.

1. What kind of deployment should Contoso, Ltd choose?
2. What kind of certificates should Contoso use for EFS and digital signing?
3. What kind of certificates should Contoso use for a website?
4. How will Contoso ensure that EFS–encrypted data is not lost if a user loses a certificate?

Best Practice

- When deploying CA infrastructure, deploy a standalone (non-domain joined) root CA, and an enterprise subordinate CA (issuing CA). After the enterprise subordinate CA receives a certificate from RootCA, take RootCA offline.
- Issue a certificate for RootCA for a long period of time such as 15 or 20 years.

- Use autoenrollment for certificates that are widely used.
- Use a Restricted Enrollment Agent whenever possible.
- Use Virtual Smart Cards for improving logon security.

Tools

- Certificate Authority console
- Certificate Templates console
- Certificates console
- Certutil.exe

Module 11

Implementing Active Directory Rights Management Services

Contents:

Module Overview	11-1
Lesson 1: AD RMS Overview	11-2
Lesson 2: Deploying and Managing an AD RMS Infrastructure	11-7
Lesson 3: Configuring AD RMS Content Protection	11-13
Lesson 4: Configuring External Access to AD RMS	11-19
Lab: Implementing AD RMS	11-24
Module Review and Takeaways	11-31

Module Overview

Active Directory® Rights Management Services (AD RMS) provides a method for protecting content that goes beyond simply encrypting storage devices using Windows® BitLocker® Drive Encryption, or individual files using Encrypting File System (EFS). AD RMS provides a method to protect data in transit and at rest, and ensures that it is accessible only to authorized users for a specific duration.

This module introduces you to AD RMS, and describes how to deploy it, how to configure content protection, and how to make AD RMS-protected documents available to external users.

Objectives

After completing this module, you will be able to:

- Provide an overview of AD RMS.
- Deploy and manage an AD RMS infrastructure.
- Configure AD RMS content protection.
- Configure external access to AD RMS.

Lesson 1

AD RMS Overview

Prior to deploying AD RMS, you need to know how AD RMS works, what components are included in an AD RMS deployment, and how you should deploy AD RMS. You must also understand the concepts behind various AD RMS certificates and licenses.

This lesson provides an overview of AD RMS, and the scenarios in which you can use it to protect an organization's confidential data.

Lesson Objectives

After completing this lesson you will be able to:

- Describe AD RMS.
- Explain the scenarios in which you can use AD RMS.
- List the AD RMS components.
- List the different AD RMS certificates and licenses.
- Explain how AD RMS works.

What Is AD RMS?

AD RMS is an information protection technology that is designed to minimize the possibility of data leakage. *Data leakage* is the unauthorized transmission of information, either to people within the organization or people outside the organization, who should not be able to access that information. AD RMS integrates with existing Microsoft products and operating systems including Exchange, SharePoint, and the Microsoft Office Suite.

AD RMS can protect data in transit and at rest. For example, AD RMS can protect documents sent as email messages, ensuring that a message cannot be opened even if it is accidentally addressed to the wrong recipient. You can also use AD RMS to protect data stored on devices such as removable USB drives. A drawback of file and folder permissions is that once the file is copied to another location, the original permissions no longer apply. A file that is copied to a USB drive will inherit the permissions on the destination device. Once copied, a file that was read-only can be made editable by altering the file and folder permissions. With AD RMS, the file can be protected in any location, irrespective of file and folder permissions that grant access. With AD RMS, only the users who are authorized to open the file will be able to view the contents of that file.

- Information protection technology
- Designed to reduce information leakage
- Integrated with Windows operating systems, Microsoft Office, Exchange, and SharePoint
- Based on Symmetric and Public Key Cryptography
- Protects data at rest, in transit, and in use

Usage Scenarios for AD RMS

The primary use for AD RMS is to control the distribution of sensitive information. You can use AD RMS in combination with encryption techniques to secure data when it is in storage or in transit. There can be many reasons to control the distribution of sensitive information, such as needing to ensure that only authorized staff members have access to a file, ensuring that sensitive email messages cannot be forwarded, or ensuring that details of an unreleased project are not made public. Consider the following scenarios.

- Prevent the transmission of sensitive information
- Comply with privacy regulations
- Can be used with encryption to protect data in transit and at rest

Scenario 1

The CEO copies a spreadsheet file containing the compensation packages of an organization's executives from a protected folder on a file server to the CEO's personal USB drive. During the commute home, the CEO leaves the USB drive on the train, where someone with no connection to the organization finds it. Without AD RMS, whoever finds the USB drive can open the file. With AD RMS, it is possible to ensure that the file cannot be opened by unauthorized users.

Scenario 2

An internal document should be viewable by a group of authorized people within the organization. These people should not be able to edit or print the document. While it is possible to use the native functionality of Microsoft® Office Word to restrict these features, doing so requires each person to have a Windows Live® account. With AD RMS, you can configure these permissions based on existing accounts in AD DS.

People within the organization should not be able to forward sensitive e-mail messages that have been assigned a particular classification. With AD RMS, you can allow a sender to assign a particular classification to a new e-mail message, and that classification will ensure that the recipient cannot forward the message.

Overview of the AD RMS Components

The AD RMS root certification cluster is the first AD RMS server that you deploy in a forest. The AD RMS root certification cluster manages all licensing and certification traffic for the domain in which it is installed. AD RMS stores configuration information either in a Microsoft SQL Server® database or in the Windows Internal Database. In large environments, the SQL Server database is hosted on a server that is separate from the server that hosts the AD RMS role.

AD RMS licensing-only clusters are used in distributed environments. Licensing-only clusters do not provide certification, but do allow the distribution of licenses that are used for content consumption and publishing. Licensing-only clusters are often deployed to large branch offices in organizations that use AD RMS.

- AD RMS server
 - Licenses AD RMS protected content
 - Certifies identity of trusted users and devices
- AD RMS client
 - Built into Windows Vista, Windows 7, and Windows 8 operating systems.
 - Interacts with AD RMS enabled applications
- AD RMS enabled applications
 - Allows publication and consumption of AD RMS protected content
 - Includes Microsoft Office, Microsoft Exchange, and SharePoint
 - Can be created using AD RMS SDKs.

AD RMS Server

AD RMS servers must be members of an AD DS domain. When you install AD RMS, information about the location of the cluster is published to AD DS to a location known as the service connection point. Computers that are members of the domain query the service connection point to determine the location of AD RMS services.

AD RMS Client

AD RMS client is built into the Windows Vista, Windows 7, and Windows 8 operating systems. The AD RMS client allows AD RMS-enabled applications to enforce the functionality dictated by the AD RMS template. Without the AD RMS client, AD RMS-enabled applications would be unable to interact with AD RMS-protected content.

AD RMS Enabled Applications

AD RMS-enabled applications allow users to create and consume AD RMS-protected content. For example, Microsoft Outlook allows users to view and create protected email messages. Microsoft Word allows users to view and create protected word processing documents.

AD RMS Certificates and Licenses

To understand how AD RMS works, you need to be familiar with its different certificates and license types. Each of these certificates and licenses functions in a different way. Some certificates, such as the server licensor certificate (SLC), are critically important and you must back them up on a regular basis.

- Server licensor certificate
- AD RMS machine certificate
- Rights Account Certificate
- Client licensor certificate
- Publishing license
- End use license

SLC

The SLC is generated when you create the AD RMS cluster. It has a validity of 250 years. The SLC allows the AD RMS cluster to issue:

- SLCs to other servers in the cluster.
- Rights Account Certificates to clients.
- Client licensor certificates.
- Publishing licenses.
- Use licenses.
- Rights policy template.

The SLC public key encrypts the content key in a publishing license. This allows the AD RMS server to extract the content key and issue end use licenses (EULs) against the publishing key.

AD RMS Machine Certificate

The AD RMS machine certificate is used to identify a trusted computer or device. The certificate identifies the client computer's lockbox. The machine certificate public key encrypts the Rights Account Certificate private key. The machine certificate private key decrypts the Rights Account Certificates.

Rights Account Certificate

The Rights Account Certificate (RAC) identifies a specific user. The default validity time for a RAC is 365 days. RACs can only be issued to users in AD DS whose user accounts have email addresses that are

associated with them. A RAC is issued the first time a user attempts to access AD RMS-protected content. You can adjust the default validity time using the Rights Account Certificate Policies node of the Active Directory Rights Management Services console.

A temporary RAC has a validity time of 15 minutes. Temporary RACs are issued when a user is accessing AD RMS-protected content from a computer that is not a member of the same or trusted forest as the AD RMS cluster. You can adjust the default validity time using the Rights Account Certificate Policies node of the Active Directory Rights Management Services console.

AD RMS supports the following additional RACs:

- Active Directory Federation Services (AD FS) RACs are issued to federated users. They have a validity of seven days.
- Two types of Windows Live® ID RAC are supported. Windows Live ID RACs used on private computers have a validity of six months. Windows Live ID RACs used on public computers are valid until the user logs off.

Client Licensor Certificate

A client licensor certificate allows a user to publish AD RMS-protected content when the client computer is not connected to the same network as the AD RMS cluster. The client licensor certificate public key encrypts the symmetric content key and includes it in the publishing license that it issues. The client licensor certificate private key signs any publishing licenses that are issued when the client is not connected to the AD RMS cluster.

Client licensor certificates are tied to a specific user's RAC. If another user who has not been issued a RAC attempts to publish AD RMS-protected content from the same client, they will be unable to until the client is connected to the AD RMS cluster and can issue that user with a RAC.

Publishing License

A publishing license (PL) determines the rights that apply to AD RMS-protected content. For example, the publishing license determines if the user can edit, print, or save a document. The publishing license contains the content key, which is encrypted using the public key of the licensing service. It also contains the URL and the digital signature of the AD RMS server.

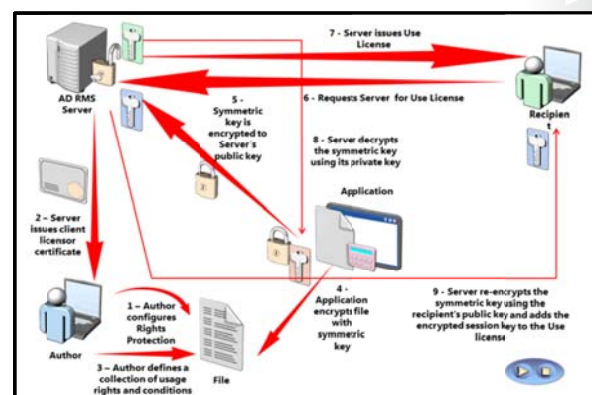
End Use License

An EUL is required to consume AD RMS-protected content. The AD RMS server issues one EUL per user per document. EULs are cached by default.

How AD RMS Works

AD RMS works in the following manner:

1. An author receives a client licensor certificate from the AD RMS server the first time he or she configures rights protection for information.
2. The author is able to define a collection of usage rights and conditions for the file. When the author does this, the application encrypts the file with a symmetric key.



3. This symmetric key is encrypted to the public key of the AD RMS server that is used by the author.
4. The recipient of the file opens it using an AD RMS application or browser. It is not possible to open AD RMS-protected content unless the application or browser supports AD RMS. If the recipient does not have an account certificate on the current device, one will be issued to the user at this point. The application or browser transmits a request to the author's AD RMS server for a Use License.
5. The AD RMS server determines if the recipient is authorized. If the recipient is authorized, the AD RMS server issues a Use License.
6. The AD RMS server decrypts the symmetric key that was encrypted in step 3, using its private key.
7. The AD RMS server re-encrypts the symmetric key using the recipient's public key and adds the encrypted session key to the Use License.

Lesson 2

Deploying and Managing an AD RMS Infrastructure

Before deploying AD RMS, it is important to have a deployment plan that is appropriate for your organization's environment. AD RMS deployment in a single-domain forest is different from AD RMS deployment in scenarios where you need to support the publication and consumption of content across multiple forests, to trusted partner organizations, or across the public Internet. Before deploying AD RMS, you also need to have an understanding of the client requirements, and an appropriate strategy for backing up and recovering AD RMS.

This lesson provides an overview of deploying AD RMS, and the steps you need to take to back up, recover, and decommission an AD RMS infrastructure.

Lesson Objectives

After completing this lesson you will be able to:

- Describe AD RMS deployment scenarios.
- Configure the AD RMS cluster.
- Explain how to install the first server of an AD RMS cluster.
- Describe AD RMS client requirements.
- Explain how to implement an AD RMS backup and recovery strategy.
- Explain how to decommission and remove AD RMS.

AD RMS Deployment Scenarios

An AD RMS deployment consists of one or more servers known as a *cluster*. An AD RMS cluster is not a high-availability failover cluster. When you are deploying AD RMS, you should host the server so that it is highly available. AD RMS is commonly deployed as a highly available virtual machine.

When you deploy AD RMS in a single forest, you have a single AD RMS cluster. This is the most common form of AD RMS deployment. You add servers to the AD RMS cluster as needed, to provide additional capacity.

When you deploy AD RMS across multiple forests, each forest must have its own AD RMS root cluster. It is necessary to configure AD RMS Trusted Publishing Domains to ensure that AD RMS content can be protected and consumed across the multiple forests.

You can also deploy AD RMS to extranet locations. In this deployment, the AD RMS licensing server is accessible to hosts on the internet. You use this type of deployment to support collaboration with external users.

You can deploy AD RMS with Active Directory Federation Services (AD FS) or the Microsoft Federation Gateway. In this scenario, users leverage federated identity to publish and consume rights-protected content.

As a best practice, you should not deploy AD RMS on a domain controller. You can only deploy AD RMS on a domain controller if the service account is a member of the Domain Admins group.

- AD RMS in a single forest
- AD RMS in multiple forests
- AD RMS used on an extranet
- AD RMS integrated with AD FS

Configuring the AD RMS Cluster

Once you have deployed the AD RMS server role, you need to configure the AD RMS cluster before it is possible to use AD RMS. Configuring the AD RMS cluster involves performing the following steps:

1. AD RMS cluster: Choose whether to create a new AD RMS root cluster, or join an existing cluster.
2. Configuration database: Select whether to use an existing SQL Server instance in which to store the AD RMS configuration database, or to configure and install the Windows Internal Database locally. You can use SQL Server 2008, SQL Server 2008 R2, or SQL Server 2012 to support an AD RMS deployment in Windows Server® 2012. As a best practice, use a SQL Server database that is hosted on a separate server.
3. Service account: Microsoft recommends using a standard domain user account with additional permissions. You can use a managed service account as the AD RMS service account.
4. Cryptographic mode: Choose the strength of the cryptography used with AD RMS.
 - Cryptographic Mode 2 uses RSA 2048-bit keys and SHA-256 hashes.
 - Cryptographic Mode 1 uses RSA 1045-bit keys and SHA-1 hashes.
5. Cluster key storage: Choose where the cluster key is stored. You can either have it stored within AD RMS, or use a special cryptographic service provider (CSP). If you choose to use a CSP, you need to manually distribute the key if you want to add additional servers.
6. Cluster key password: This password encrypts the cluster key, and is required if you want to join other AD RMS servers to the cluster, or if you want to restore the cluster from backup.
7. Cluster website: Choose which website on the local server will host the AD RMS cluster website.
8. Cluster address: Specify the fully qualified domain name (FQDN) used with the cluster. You have the option of choosing between an Secure Sockets Layer (SSL)–encrypted and non-SSL-encrypted website. If you choose non-SSL-encrypted, you will be unable to add support for Identity Federation. Once you set the cluster address and port, you cannot change them without completely removing AD RMS.
9. Licensor certificate: Choose is the friendly name used by the SLC. It should represent the function of the certificate.
10. Service connection point registration: Choose whether the service connection point is registered in AD DS when the AD RMS cluster is created. The service connection point allows computers that are members of the domain to locate the AD RMS cluster automatically. Only users that are members of the Enterprise Admins group are able to register the service connection point. You can perform this step after the AD RMS cluster is created—you do not have to perform it during the configuration process.

Involves choosing the following:

- New or join existing cluster
- Configuration database location
- Service account
- Cryptographic mode
- Cluster key storage
- Cluster key password
- Cluster website
- Cluster address
- Server certificate
- Licensor certificate
- SCP registration

Demonstration: Installing the First Server of an AD RMS Cluster

In this demonstration, you will deploy AD RMS on a computer that is running Windows Server 2012.

Demonstration Steps

Configure Service Account

1. Log on to **LON-DC1** with the **Adatum\Administrator** account and the password **Pa\$\$w0rd**.
2. Use the Active Directory Administrative Center to create an Organizational Unit (OU) named **Service Accounts** in the **adatum.com** domain.
3. Create a new user account in the Service Accounts OU with the following properties:
 - o First name: **ADRMSSVC**
 - o User UPN logon: **ADRMSSVC**
 - o Password: **Pa\$\$w0rd**
 - o Password never expires: **Enabled**
 - o User cannot change password: **Enabled**

Prepare DNS

- Use the DNS Manager console to create a host (A) resource record in the **adatum.com** zone with the following properties:
 - o Name: **adrms**
 - o IP Address: **172.16.0.21**

Install the AD RMS role

1. Log on to **LON-SVR1** with the **Adatum\Administrator** account using the password **Pa\$\$w0rd**.
2. Use the Add Roles and Features Wizard to add the AD RMS role to LON-SVR1 using the following option:
 - o Role services: **Active Directory Rights Management Server**

Configure AD RMS

1. In Server Manager, from the **AD RMS** node, click **More** to start post deployment configuration of AD RMS.
2. In the AD RMS Configuration Wizard, provide the following information:
 - o **Create a new AD RMS root cluster**
 - o **Use Windows Internal Database on this server**
 - o **Use Adatum\ADRMSSVC as the service account**
 - o Cryptographic Mode: **Cryptographic Mode 2**
 - o Cluster Key Storage: **Use AD RMS centrally managed key storage**
 - o Cluster Key Password: **Pa\$\$w0rd**
 - o Cluster Web Site: **Default Web Site**
 - o Connection Type: **Use an unencrypted connection**
 - o Fully Qualified Domain Name: **http://adrms.adatum.com**
 - o Port: **80**

- Licensor Certificate: **Adatum AD RMS**
 - Register AD RMS Service Connection Point: **Register the SCP Now**
3. Log off from LON-SVR1.



Note: You must sign out before you can manage AD RMS

AD RMS Client Requirements

AD RMS content can only be published and consumed by computers that are running the AD RMS client. All versions of Windows Vista®, Windows 7, and Windows 8 client operating systems include AD RMS client software. Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012 operating systems also include the AD RMS client. These operating systems do not require additional configuration to consume and publish AD RMS-protected content.

AD RMS client software is available for download to computers that are running the Windows XP operating system, and Mac OS X. This client software must be installed before it is possible for users of these operating systems to be able to consume and publish AD RMS-protected content.

AD RMS requires compatible applications. Server applications that support AD RMS include the following:

- Microsoft Exchange Server 2007
- Exchange Server 2010
- Exchange Server 2013
- Microsoft Office SharePoint Server 2007
- SharePoint Server 2010
- SharePoint Server 2013

Client applications, such as those included in Microsoft Office 2007, Office 2010, and Office 2013 can publish and consume AD RMS-protected content. You can use the AD RMS Software Development Kit (SDK) to create applications that can publish and consume AD RMS-protected content. Microsoft XPS viewer and Windows Internet Explorer® are also able to view AD RMS-protected content.

- Client included in Windows Vista, Windows 7, Windows 8 operating systems
- Client included in Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012 operating systems
- Client available for download for previous versions of Windows operating systems, and Mac OS X
- AD RMS-enabled applications include Office 2007, Office 2010, and Office 2013
- Microsoft Exchange Server 2007, Exchange Server 2010 and Exchange Server 2013 support AD RMS

Implementing an AD RMS Backup and Recovery Strategy

To prevent data loss, you must ensure that the AD RMS server is backed up in such a way that it can be recovered in the event of file corruption or server failure. If the AD RMS server becomes inaccessible, all AD RMS-protected content also becomes inaccessible.

A simple strategy for implementing AD RMS backup and recovery is to run AD RMS server as a virtual machine, and then use an enterprise backup product such as Microsoft System Center 2012 - Data Protection Manager to perform regular virtual machine backups. Some of the

important components that require backups are the private key, certificates, the AD RMS database, and templates. You can also perform a full server backup, by running AD RMS server on a virtual machine.

As a best practice, you need to back up the AD RMS private key and all certificates used by AD RMS. The simplest method of doing this is to export the certificates to a safe location. You must also back up the AD RMS database on a regular basis. The method you use to do this depends on whether AD RMS uses SQL Server or the Windows Internal Database. To back up templates, configure the templates to be exported to a shared folder, and then back up these templates.

When you are performing recovery of the AD RMS role, it may be necessary to delete the **ServiceConnectionPoint** object from AD DS. You need to do this if you are recovering an AD RMS root configuration server, and the server attempts to provision itself as a licensing-only server.

- Back up private key and certificates
- Ensure that AD RMS database is backed up regularly
- Export templates to back them up
- Run AD RMS server in virtual machine, and perform full server backup

Decommissioning and Removing AD RMS

Prior to removing an AD RMS server, you should decommission that server. Decommissioning AD RMS puts the cluster into a state where consumers of AD RMS-protected content are able to obtain special keys that decrypts that content, irrespective of the existing restrictions that were placed on the use of that content. If you do not have a decommissioning period, and if you simply remove the AD RMS server, then the AD RMS-protected content will become inaccessible.

To decommission AD RMS, perform the following steps:

1. Log on to the server that is hosting AD RMS, and that you wish to decommission.
2. Modify the access control list (ACL) of the file **decommissioning.asmx**. Grant the Everyone group **Read & Execute** permission on the file. This file is stored in the %systemdrive%\inetpub\wwwroot_wmcs\decommission folder.
3. In the Active Directory Rights Management Services console, expand the **Security Policies** node, and then click the **Decommissioning** node.
4. In the Actions pane, select **Enable Decommissioning**.
5. Click **Decommission**.

- Decommission an AD RMS cluster prior to removing it
 - Decommissioning provides key that decrypts previously published AD RMS content
 - Leave server in decommissioned state until all AD RMS-protected content is migrated
- Export the server licensor certificate prior to uninstalling the AD RMS role

6. When prompted to confirm that you want to decommission the server, click **Yes**.

After the AD RMS decommissioning process is complete, you should export the server licenser certificate prior to uninstalling the AD RMS role.

Lesson 3

Configuring AD RMS Content Protection

AD RMS uses rights policy templates to enforce a consistent set of policies when protecting content. When configuring AD RMS, you also need to develop strategies to ensure that users can still access protected content from a computer that is not connected to the AD RMS cluster. You also need to develop strategies for excluding some users from being able to access AD RMS-protected content, and strategies to ensure that protected content can be recovered in the event that it has expired, the template has been deleted, or if the author of the content is no longer available.

Lesson Objectives

After completing this lesson you will be able to:

- Describe the function of rights policy templates.
- Explain how to create a rights policy template.
- Implement strategies to ensure rights policy templates are available for offline use.
- Describe exclusion policies.
- Explain how to create an exclusion policy to exclude an application.
- Implement an AD RMS super users group.

What Are Rights Policy Templates?

Rights policy templates allow you to configure standard methods of implementing AD RMS policies across the organization. For example, you can configure standard templates that grant view-only rights, block the ability to edit, save, and print, or if used with Microsoft Exchange Server, block the ability to forward, reply, and reply all to messages.

Rights policy templates are created using the Active Directory Rights Management Services console. They are stored in the AD RMS database, and can also be stored in XML format. When content is consumed, the client checks with AD RMS to verify that it has the most recent version of the template.

A document author can choose to protect content by applying an existing template. This is done using an AD RMS-aware application. For example, in Office Word, you apply a template by using the **Protect Document** function. When you do this, Office Word queries AD DS to determine the location of the AD RMS server. Once the location of the AD RMS server is acquired, templates that are available to the content author can be used.

AD RMS templates support the following rights:

- **Full Control.** Gives a user full control over an AD RMS-protected document.
- **View.** Gives a user the ability to view an AD RMS-protected document.
- **Edit.** Allows a user to modify an AD RMS-protected document.

- Allow authors to apply standard forms of protection across the organization
- Different applications allow different forms of rights
- Can configure rights related to viewing, editing and printing documents
- Can configure content expiration rights
- Can configure content revocation

- **Save.** Allows a user to use the **Save** function with an AD RMS–protected document.
- **Export (Save as).** Allows a user to use the **Save As** function with an AD RMS–protected document.
- **Print.** Allows an AD RMS–protected document to be printed.
- **Forward.** Used with Exchange Server. Allows the recipient of an AD RMS–protected message to forward that message.
- **Reply.** Used with Exchange Server. Allows the recipient of an AD RMS–protected message to reply to that message.
- **Reply All.** Used with Exchange Server. Allows the recipient of an AD RMS–protected message to use the Reply All function to reply to that message.
- **Extract.** Allows the user to copy data from the file. If this right is not granted, the user cannot copy data from the file.
- **Allow Macros.** Allows the user to utilize macros.
- **View Rights.** Allows the user to view assigned rights.
- **Edit Rights.** Allows the user to modify the assigned rights.

Rights can only be granted, and cannot be explicitly denied. For example, to ensure that a user cannot print a document, the template associated with the document must not include the **Print** right.

Administrators are also able to create custom rights that can be used with custom AD RMS–aware applications.

AD RMS templates can also be used to configure documents with the following properties:

- **Content Expiration.** Determines when the content expires. The options are:
 - **Never.** The content never expires.
 - **Expires on a particular date.** Content expires at a particular date and time.
 - **Expires after.** The content expires a particular number of days after it is created.
- **Use license expiration.** Determines the time interval in which the use license will expire, and a new one will need to be acquired.
- **Enable users to view protected content using a browser add-on.** Allows content to be viewed using a browser add-on. Does not require the user have an AD RMS–aware application.
- **Require a new use license each time content is consumed.** When you enable this option, client-side caching is disabled. This means that the document cannot be consumed when the computer is offline.
- **Revocation policies.** Allows the use of a revocation list. This allows an author to revoke permission to consume content. You can specify how often the revocation list is checked, with the default being once every 24 hours.

Once an AD RMS policy template is applied to a document, any updates to that template will also be applied to that document. For example, if you have a template without a content expiration policy that is used to protect documents, and you modify that template to include a content expiration policy, those protected documents will now have an expiration policy. Template changes are reflected when the EUL is acquired. If EULs are configured not to expire and the user who is accessing the document already has a license, then they may not receive the updated template.

You should avoid deleting templates, because documents that use those templates will become inaccessible to everyone except for members of the super users group. As a best practice, archive templates instead of deleting them.

You can view the rights associated with a template by selecting the template within the Active Directory Rights Management Services console, and then in the **Actions** menu, clicking **View Rights Summary**.

Demonstration: Creating a Rights Policy Template

In this demonstration, you will create a rights policy template that allows users to view a document, but not to perform other actions.

Demonstration Steps

- In the Active Directory Rights Management Services console, use the Rights Policy Template node to create a Distributed Rights Policy Template with the following properties:
 - Language: **English (United States)**
 - Name: **ReadOnly**
 - Description: **Read-only access. No copy or print.**
 - Users and rights: **executives@adatum.com**
 - Rights for Anyone: **View**
 - **Grant owner (author) full control right with no expiration**
 - Content Expiration: Expires after **7** days
 - Use license expiration: Expires after **7** days
 - Require a new use license every time content is consumed (disable client-side caching): **Enabled**

Providing Rights Policy Templates for Offline Use

If users are going to publish AD RMS–connected templates when they are not connected to the network, you need to ensure that they have access to a local copy of the available rights policy templates.

You can configure computers to acquire and store published rights policy templates automatically, so that they are available offline. To enable this feature, computers must be running the following Windows operating systems:

- Windows Vista SP1 or newer
- Windows 7
- Windows 8
- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012

- Ensure that templates are published to a shared folder
- Enable the AD RMS Rights Policy Template Management (Automated) scheduled task
- Edit the registry key and specify the shared folder location

To enable this functionality, in the Task Scheduler, enable the **AD RMS Rights Policy Template Management (Automated) Scheduled Task**, and then edit the following registry key:

HKEY_CURRENT_USER\Software\Microsoft\Office\12.0\Common\DRM

Provide the following location for templates to be stored:

%LocalAppData%\Microsoft\DRM\Templates

When computers that are running these operating systems are connected to the domain, the AD RMS client polls the AD RMS cluster for new templates, or updates to existing templates.

You can configure a shared folder for templates by performing the following steps:

1. In the Active Directory Rights Management Services console, right-click the **Rights Policy Templates** node, and then click **Properties**.
2. On the **Rights Policy Templates Properties** dialog box, specify the location of the shared folder to which templates will be published.

What Are Exclusion Policies?

Exclusion policies allow you to prevent specific user accounts, client software, or applications from using AD RMS.

User Exclusion

The User Exclusion policy allows you to configure AD RMS so that specific user accounts—which are identified based on email addresses—are unable to obtain Use Licenses. You do this by adding each user's RAC to the exclusion list. User Exclusion is disabled by default. Once you have enabled User Exclusion, you can exclude specific RACs.

Allows you to:

- Block specific users from accessing AD RMS-protected content by blocking their Rights Account Certificate
- Block specific applications from creating or consuming AD RMS protected content
- Block specific versions of the AD RMS client

You can use user exclusion in the event that you needed to lock a specific user out of AD RMS-protected content. For example, when users leave the organization, you might exclude their RACs to ensure that they are unable to access protected content. You can block the RACs that are assigned to both internal users and external users.

Application Exclusion


Application Exclusion allows you to block specific applications—such as Office PowerPoint—from creating or consuming AD RMS-protected content. You specify applications based on executable names. You also specify a minimum and a maximum version of the application. Application Exclusion is disabled by default.



Note: It is possible to circumvent Application Exclusion by renaming an executable file.

Lockbox Exclusion

Lockbox exclusion allows you to exclude AD RMS clients, such as those used with specific operating systems such as Windows XP and Windows Vista. Lockbox version exclusion is disabled by default. Once you have enabled Lockbox version exclusion, you must specify the minimum lockbox version that can be used with the AD RMS cluster.

 **Additional Reading:** To find out more about enabling exclusion policies, refer to the following TechNet webpage: <http://technet.microsoft.com/en-us/library/cc730687.aspx>

Demonstration: Creating an Exclusion Policy to Exclude an Application

In this demonstration, you will see how to exclude a specific application from AD RMS.

Demonstration Steps

1. In the Active Directory Rights Management Services console, enable Application exclusion.
2. In the **Exclude Application** dialog box, enter the following information:
 - Application File name: **Powerpnt.exe**
 - Minimum version: **14.0.0.0**
 - Maximum version: **16.0.0.0**

AD RMS Super Users Group

The AD RMS super users group provides a data recovery mechanism for AD RMS-protected content. This mechanism is useful in the event that AD RMS-protected data needs to be recovered, such as when content has expired, when a template has been deleted, or when you do not have access.

Members of the super users group are assigned Owner Use Licenses for all content that is protected by the AD RMS cluster on which that particular super users group is enabled. Members of the super users group are able to reset the AD RMS server's private key password.

As members of the super users group can access any AD RMS-protected content, you must be especially careful when you are managing the membership of this group. If you choose to use the AD RMS super users group, you should consider implementing restricted groups policy and auditing to limit group membership, and audit any changes that are made. Super User activity is written to the Application event log.

The super users group is disabled by default.

You enable the super users group by performing the following steps:

1. In the Active Directory Rights Management Services console, expand the server node, and then click **Security Policies**.
2. In the **Security Policies** area, under **Super Users**, click **Change Super User Settings**.
3. In the Actions pane, click **Enable Super Users**.

To set a particular group as the super users group:

1. In the **Security Policies\Super Users Super Users** area, click **Change super user group**.
2. Provide the e-mail address associated with the super users group.

- Not configured by default
- Data recovery mechanism for AD RMS-protected content
 - Can recover content that has expired
 - Can recover content if template deleted
 - Can recover content without requiring author credentials
- Must be an Active Directory group with an assigned email address.

Lesson 4

Configuring External Access to AD RMS

It is often necessary to allow users that are not a part of the organization access to AD RMS–protected content. This could be a situation where an external user is a contractor who requires access to sensitive materials, or a partner organization where your users will require access to protected content published by their AD RMS server. AD RMS provides a number of different options for granting external users access to protected content.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the options available for making AD RMS–protected content accessible to external users.
- Explain how to implement Trusted User Domains.
- List the steps necessary to deploy Trusted Publishing Domains.
- Describe the steps necessary to configure AD RMS to share protected content to users with Windows Live IDs.
- Determine the appropriate solution for sharing AD RMS–protected content with external users.

Options for Enabling External Users to Access AD RMS

Trust policies allow users who are external to the organization the ability to consume AD RMS–protected content. For example, a trust policy can allow users in Bring Your Own Device (BYOD) environments to consume AD RMS–protected content, even though those computers are not members of the organization's AD DS domain. AD RMS trusts are disabled by default, and need to be enabled before being used. AD RMS supports the following trust policies.

- **Trusted User Domains**
 - Exchange protected content between two organizations
- **Trusted Publishing Domains**
 - Consolidate AD RMS architecture
- **Federation Trust**
 - One AD RMS infrastructure accessible to AD FS partners
- **Windows Live ID**
 - Allow standalone users access to AD RMS content
- **Microsoft Federation Gateway**
 - Allow an AD RMS cluster to work with Microsoft Federation Gateway without requiring a direct Federation Trust

Trusted User Domains

Trusted User Domains (TUD) allows an AD RMS cluster to process requests for client licensor certificates, or use licenses from people who have RACs issued by a different AD RMS cluster. For example, A. Datum Corporation and Trey Research are separate organizations that have each deployed AD RMS. TUD allows each organization to publish and consume AD RMS–protected content to and from the partner organization without having to implement AD DS trusts or AD FS.

Trusted Publishing Domains

Trusted Publishing Domains (TPD) allows one AD RMS cluster to issue EULs to content that uses publishing licenses that are issued by a different AD RMS cluster. TPD consolidates existing AD RMS infrastructure.

Federation Trust

Federation Trust provides Single Sign On (SSO) for partner technologies. Federated partners can consume AD RMS–protected content without deploying their own AD RMS infrastructure. Federation Trust requires deployment of AD FS.

Windows Live ID Trust

You can use Windows Live ID to allow standalone users that have Windows Live IDs to consume AD RMS–protected content generated by users in your organization. However, Windows Live ID users are unable to create content that is protected by the AD RMS cluster.

Microsoft Federation Gateway

Microsoft Federation Gateway allows an AD RMS cluster to process requests to publish and consume AD RMS–protected content from external organizations, by accepting claims-based authentication tokens from the Microsoft Federation Gateway. Rather than configuring a Federation Trust, each organization has a relationship with the Microsoft Federation Gateway. The Microsoft Federation Gateway acts as a trusted broker.



Additional Reading: You can learn more about AD RMS Trust Policies at the following link: <http://technet.microsoft.com/en-us/library/cc755156.aspx>

Implementing TUD

TUD allows AD RMS to service requests from users who have RACs issued by different AD RMS deployments. You can use exclusions with each TUD to block access to specific users and groups.

To configure AD RMS to support service requests from users who have RACs issued by different AD RMS deployments, you add the organization to the list of TUDs. TUDs can be one-way, where organization A is a TUD of organization B, or bi-directional, where organization A and organization B are TUDs of each other. In one-way deployments, it is possible for the users of the TUD to consume the content of the local AD RMS deployment, but they cannot publish AD RMS–protected content by using the local AD RMS cluster.

You need to enable anonymous access to the AD RMS licensing service in Internet Information Services (IIS) when using TUD, as by default, accessing the service requires authenticating using Integrated Windows Authentication.

To add a TUD, perform the following steps:

1. The TUD of the AD RMS deployment that you want to trust must have already been exported, and the file must be available. (TUD files use the .bin extension.)
2. In the AD RMS console, expand **Trust Policies**, and then click **Trusted User Domains**.
3. In the Actions pane, click **Import Trusted User Domain**.
4. In the **Trusted User Domain** dialog box, enter the path to the exported TUD file with the .bin extension.
5. Provide a name to identify this TUD. If you have configured federation, you can also choose to extend the trust to federated users of the imported server.

You can use the **Import-RmsTUD PowerShell** cmdlet, which is part of the ADRMSADMIN PowerShell module, to add a TUD.

- Allows AD RMS to service requests to users with RACs from different AD RMS clusters.
- TUDs:
 - Support exclusions to individual users and groups
 - Can be one way or bi-directional
- Must export TUD from partner before importing TUD locally

To export a TUD, perform the following steps:

1. In the Active Directory Rights Management Services console, expand **Trust Policies**, and then click **Trusted User Domains**.
2. In the Actions pane, click **Export Trusted User Domain**.
3. Save the TUD file with a descriptive name.

You can also use the **Export-RmsTUD** cmdlet to export an AD RMS server TUD.

Implementing TPD

You can use TPD to set up a trust relationship between two AD RMS deployments. An AD RMS TPD, which is a local AD RMS deployment, can grant EULs for content published using the Trusted Publishing domain's AD RMS deployment. For example, Contoso, Ltd and A. Datum Corporation are set up as TPD partners. TPD allows users of the Contoso AD RMS deployment to consume content published using the A. Datum AD RMS deployment, by using EULs that are granted by the Contoso AD RMS deployment.

- Allows a local AD RMS deployment to issue EULs to content protected by partner AD RMS cluster
- Involves importing SLC of partner AD RMS cluster
- No limit to number of supported TPDs

You can remove a TPD at any time. When you do this, clients of the remote AD RMS deployment will not be able to issue EULs to access content protected by your AD RMS cluster.

When you are configuring a TPD, you import the SLC of another AD RMS cluster. TPDs are stored in XML format, and are protected by passwords.

To export a TPD, perform the following steps:


1. In the AD RMS console, expand **Trust Policies**, and then click **Trusted Publishing Domains**.
2. In the Results pane, choose the certificate for the AD RMS domain that you want to export, and then in the Actions pane, click **Export Trusted Publishing Domain**.
3. Choose a strong password and a filename for the TPD.

When you are exporting a TPD, it is possible to save it as a V1-compatible trusted publishing domain file. This allows the TPD to be imported into organizations that are using AD RMS clusters on earlier versions of the Windows Server operating system, such as the version available in Windows Server 2003. You can use the **Export-RmsTPD** cmdlet to export a TPD.

To import a TPD, perform the following steps:

1. In the Active Directory Rights Management Services console, expand **Trust Policies**, and then click **Trusted Publishing Domains**.
2. In the Actions pane, click **Import Trusted Publishing Domain**.
3. Specify the path of the Trusted Publishing Domain file that you want to import.
4. Enter the password to open the Trusted Publishing Domain file, and enter a display name that identifies the TPD.

You can also use the **Import-RmsTPD** cmdlet to import a TPD.

 **Additional Reading:** You can learn more about importing TPDs on the following Microsoft TechNet reference: <http://technet.microsoft.com/en-us/library/cc771460.aspx>

Sharing AD RMS–Protected Documents by Using Windows Live ID

You can use Windows Live ID as a method of providing RACs to users who are not part of your organization.

To trust Windows Live ID–based RACs, perform the following steps:


1. In the Active Directory Rights Management Services console, expand **Trust Policies**, and then click **Trusted User Domains**.
2. In the Actions pane, click **Trust Windows Live ID**.

- Provide RACs to users who are not part of an organization
- Users with Windows Live accounts can consume AD RMS protected content
- Users with Windows Live accounts cannot publish AD RMS-protected content

To exclude specific Windows Live ID email domains, right-click the Windows Live ID certificate, click **Properties**, and then click the **Excluded Windows Live IDs** tab. You can then enter the Windows Live IDs that you want to exclude from being able to procure RACs.

To allow users with Windows Live IDs to obtain RACs from your AD RMS cluster, you need to configure IIS to support anonymous access. To do this, perform the following steps:

1. Open the IIS Manager console on the AD RMS server.
2. Navigate to the **Sites\Default Web Site_wmcs** node, right-click the **Licensing** virtual directory, and then click **Switch to Content View**.
3. Right-click **license.asmx**, and then click **Switch to Content View**.
4. Double-click **Authentication**, and then enable **Anonymous Authentication**.
5. Repeat this step for the file **ServiceLocator.asmx**.

 **Additional Reading:** You can learn more about using Windows Live ID to establish RACs for users at the following link: <http://technet.microsoft.com/en-us/library/cc753056.aspx>

Considerations for Implementing External User Access to AD RMS

The type of external access that you configure depends on the types of external users that need access to your organization's content.

When you are determining which method to use, consider the following questions:

- Does the external user belong to an organization that has an existing AD RMS deployment?
- Does the external user's organization have an existing federated trust with the internal organization?
- Has the external user's organization established a relationship with the Microsoft Federation Gateway?
- Does the external user need to publish AD RMS-protected content that is accessible to internal RAC holders?

- Use Windows Live ID to issue RACs to users who are not part of organizations, and who need to consume content
- Use TUD for RACs issued by a different AD RMS cluster
- Use TPD to allow local RACs to access remotely published AD RMS content
- Use Federation Trust between organizations that have a federated relationship
- Use Microsoft Federation Gateway when no direct federated relationship exists

It is possible that organizations may use one solution before settling on another. For example, during initial stages, only a small number of external users may require access to AD RMS-protected content, in which case, using Windows Live IDs for RACs may be appropriate. When larger numbers of external users from a single organization require access, a different solution may be appropriate. The financial benefit a solution brings to an organization must exceed the cost of implementing that solution.

Lab: Implementing AD RMS

Scenario

Because of the highly confidential nature of the research that is performed at A. Datum Corporation, the security team at A. Datum wants to implement additional security for some of the documents that the Research department creates. The security team is concerned that anyone with **Read** access to the documents can modify and distribute the documents in any way that they choose. The security team would like to provide an extra level of protection that stays with the document even if it is moved around the network or outside the network.

As one of the senior network administrators at A. Datum, you need to plan and implement an AD RMS solution that will provide the level of protection requested by the security team. The AD RMS solution must provide many different options that can be adapted for a wide variety of business and security requirements.

Objectives

- Install and configure AD RMS.
- Configure AD RMS Templates.
- Implement AD RMS Trust Policies.
- Verify AD RMS Deployment.

Lab Setup

Estimated Time: 60 minutes

- 20412A-LON-DC1
- 20412A-LON-SVR1
- 20412A-LON-CL1
- 20412A-MUN-DC1
- 20412A-MUN-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 **20412A-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on using the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
5. Repeat step 2 for **20412A-LON-SVR1**, **20412A-MUN-DC1**, **20412A-LON-CL1**, and **20412A-MUN-CL1**. Do not log on until directed to do so.

Exercise 1: Installing and Configuring AD RMS

Scenario

The first step in deploying AD RMS at A. Datum Corporation is to deploy a single server in an AD RMS cluster. You will begin by configuring the appropriate DNS records and the AD RMS service account, and then will continue with installing and configuring the first AD RMS server. You will also enable the AD RMS super users group.

The main tasks for this exercise are as follows:

1. Configure Domain Name System (DNS) and the Active Directory® Rights Management Services (AD RMS) service account.
2. Install and configure the AD RMS server role.
3. Configure the AD RMS Super Users group.


► Task 1: Configure Domain Name System (DNS) and the Active Directory® Rights Management Services (AD RMS) service account

1. Log on to **LON-DC1** with the **Adatum\Administrator** account and the password **Pa\$\$w0rd**.
2. Use Active Directory Administrative Center to create an OU named **Service Accounts** in the **adatum.com** domain.
3. Create a new user account in the Service Accounts OU with the following properties:
 - First name: **ADMSSVC**
 - User UPN logon: **ADMSSVC**
 - Password: **Pa\$\$w0rd**
 - Password never expires: **Enabled**
 - User cannot change password: **Enabled**
4. Create a new Global security group in the Users container named **ADRMS_SuperUsers**. Set the email address of this group as **ADRMS_SuperUsers@adatum.com**.
5. Create a new global security group in the Users container named **Executives**. Set the email address of this group as **executives@adatum.com**.
6. Add the user accounts **Aidan Delaney** and **Bill Malone** to the Executives group.
7. Use the DNS Manager console to create a host (**A**) resource record in the **adatum.com** zone with the following properties:
 - Name: **adrms**
 - IP Address: **172.16.0.21**

► Task 2: Install and configure the AD RMS server role

1. Log on to **LON-SVR1** with the **Adatum\Administrator** account and the password **Pa\$\$word**.
2. Use the Add Roles and Features Wizard to add the Active Directory Rights Management Services role to LON-SVR1 using the following option:
 - Role services: **Active Directory Rights Management Services**
3. From the AD RMS node in Server Manager, click **More** to start post deployment configuration of AD RMS.
4. On the AD RMS Configuration Wizard, provide the following information:

- **Create a new AD RMS root cluster**
 - **Use Windows Internal Database on this server**
 - Service account: **Adatum\ADRMSSVC**
 - Cryptographic Mode: **Cryptographic Mode 2**
 - Cluster Key Storage: **Use AD RMS centrally managed key storage**
 - Cluster Key Password: **Pa\$\$w0rd**
 - Cluster Web Site: **Default Web Site**
 - Connection Type: **Use an unencrypted connection**
 - Fully Qualified Domain Name: **http://adrms.adatum.com**
 - Port: **80**
 - Licensor Certificate: **Adatum AD RMS**
 - Register AD RMS Service Connection Point: **Register the SCP Now**
5. Log off LON-SVR1.

 **Note:** You must sign out before you can manage AD RMS. This lab uses port 80 for convenience. In production environments, you would protect AD RMS using an encrypted connection.

► Task 3: Configure the AD RMS Super Users group

1. Log on to **LON-SVR1** with the **Adatum\Administrator** account and the password **Pa\$\$w0rd**.
2. Open the Active Directory Rights Management Services console.
3. From the Active Directory Rights Management Services console, enable Super Users.
4. Set the **AD RMS_SuperUsers** group as the Super Users group.

Results: After completing this exercise, you should have installed and configured AD RMS.

Exercise 2: Configuring AD RMS Templates

Scenario

After deploying the AD RMS server, the next step is to configure the rights policy templates and exclusion policies for the organization. You will deploy both components.

The main tasks for this exercise are as follows:

1. Configure a new rights policy template.
2. Configure the rights policy template distribution.
3. Configure an exclusion policy.

► Task 1: Configure a new rights policy template

- On LON-SVR1, use the Rights Policy Template node of the Active Directory Rights Management Services console to create a Distributed Rights Policy Template with the following properties:
 - Language: **English (United States)**
 - Name: **ReadOnly**
 - Description: **Read only access. No copy or print**

- Users and rights: **executives@adatum.com**
- Rights for Anyone: **View**
- **Grant owner (author) full control right with no expiration**
- Content Expiration: **7** days
- Use license expiration: **7** days
- Require a new use license **every time content is consumed (disable client-side caching)**

► **Task 2: Configure the rights policy template distribution**

1. On LON-SVR1, open a Windows PowerShell prompt and issue the following commands each followed by Enter:

```
Cmd.exe
mkdir c:\rmstemplates
net share RMTEMPLATES=C:\rmstemplates /GRANT:ADATUM\ADRMSSVC,FULL
mkdir c:\docshare
net share docshare=c:\docshare /GRANT:Everyone,FULL
```

2. In the Active Directory Rights Management Services console, set the Rights Policy Templates file location to **\\LON-SVR1\RMTEMPLATES**.
3. In Windows Explorer, view the c:\rmstemplates folder. Verify that the **ReadOnly.xml** template is present.

► **Task 3: Configure an exclusion policy**

1. In the Active Directory Rights Management Services console, enable Application exclusion.
2. In the **Exclude Application** dialog box, enter the following information:
 - Application File name: **Powerpnt.exe**
 - Minimum version: **14.0.0.0**
 - Maximum version: **16.0.0.0**

Results: After completing this exercise, you should have configured AD RMS templates.

Exercise 3: Implementing the AD RMS Trust Policies

Scenario

As part of the deployment, you need to ensure that AD RMS functionality is extended to the Trey Research AD RMS deployment. You will configure the required trust policies, and then validate that you can share protected content between the two organizations.

The main tasks for this exercise are as follows:

1. Export the Trusted User Domains policy.
2. Export the Trusted Publishing Domains policy.
3. Import the Trusted User Domain policy from the partner domain.
4. Import the Trusted Publishing Domains policy from the partner domain.
5. Configure anonymous access to the AD RMS licensing server.

► Task 1: Export the Trusted User Domains policy

1. On LON-SVR1, open a Windows PowerShell prompt and issue the following commands:

```
Cmd.exe
mkdir c:\export
net share export=c:\export /GRANT:Everyone,FULL
```

2. Use the Active Directory Rights Management Services console to export the Trusted User Domains policy to the **\\LON-SVR1\export** share as **ADATUM-TUD.bin**.
3. Log on to **MUN-DC1** with the **TREYRESEARCH\Administrator** account and the password **Pa\$\$w0rd**.
4. On MUN-DC1, open the Active Directory Rights Management Services console.
5. Export the **Trusted User domains** policy to the **\\LON-SVR1\export** share as **TREYRESEARCH-TUD.bin**.

► Task 2: Export the Trusted Publishing Domains policy

1. Switch to LON-SVR1.
2. Use the Active Directory Rights Management Services console to export the Trusted Publishing Domains policy to the **\\LON-SVR1\export** share as **ADATUM-TPD.xml**. Protect this file using the password **Pa\$\$w0rd**.
3. Switch to MUN-DC1.
4. Use the Active Directory Rights Management Services console to export the Trusted Publishing Domains policy to the **\\LON-SVR1\export** share as **TREYRESEARCH-TPD.xml**. Protect this file using the password **Pa\$\$w0rd**.

► Task 3: Import the Trusted User Domain policy from the partner domain

1. Switch to LON-SVR1.
2. Import the Trusted User Domain policy for Treyresearch by importing the file **\\LON-SVR1\export\treyresearch-tud.bin**. Use the display name **TreyResearch**.
3. Switch to MUN-DC1.
4. Import the Trusted User Domain policy for Trey Research by importing the file **\\LON-SVR1\export\adatum-tud.bin**. Use the display name **Adatum**.

► Task 4: Import the Trusted Publishing Domains policy from the partner domain

1. Switch to LON-SVR1.
2. Import the Trey Research Trusted Publishing Domain by importing the file **\\LON-SVR1\export\treyresearch-tpd.xml** using the password **Pa\$\$w0rd** and the display name **Trey Research**.
3. Switch to MUN-SVR1.
4. Import the Adatum Trusted Publishing Domain by importing the file **\\LON-SVR1\export\adatum-tpd.xml** using the password **Pa\$\$w0rd** and the display name **Adatum**.

► Task 5: Configure anonymous access to the AD RMS licensing server

- On LON-SVR1, use Internet Information Services (IIS) to enable anonymous authentication on the following two files under Default Web Site_wmcs\Licensing

- **license.asmx**
- **ServiceLocator.asmx**

Results: After completing this exercise, you should have implemented the AD RMS trust policies.

Exercise 4: Verifying the AD RMS Deployment

Scenario

As a final step in the deployment, you will validate that the configuration is working correctly.

The main tasks for this exercise are as follows:

1. Create a rights-protected document.
2. Verify internal access to protected content.
3. Open the rights-protected document as an unauthorized user.
4. Open and edit the rights-protected document as an authorized user at Trey Research.

► Task 1: Create a rights-protected document

1. Log on to **LON-CL1** with the **Adatum\Aidan** account and the password **Pa\$\$w0rd**.
2. Open Microsoft Word 2010.
3. Create a document named **Executives Only**.
4. In the document, type the following text:
This document is for executives only, it should not be modified.
5. From the **Permissions** item, choose to restricted access. Grant **bill@adatum.com** permission to read the document.
6. Save the document in the share **\\lon-svr1\docshare**.
7. Log off from LON-CL1.

► Task 2: Verify internal access to protected content

1. Log on to **LON-CL1** with the **Adatum\Bill** account using the password **Pa\$\$w0rd**.
2. In the **\\lon-svr1\docshare** folder, open the **Executives Only** document.
3. When prompted, provide the credentials, **Adatum\Bill** with the password of **Pa\$\$w0rd**.
4. Verify that you are unable to modify or save the document.
5. Select a line of text in the document.
6. Right-click the line of text. Verify that you cannot modify this text.
7. View the document permissions.
8. Log off from LON-CL1.

► Task 3: Open the rights-protected document as an unauthorized user

1. Log on to **LON-CL1** as **Adatum\Carol** using the password **Pa\$\$w0rd**.
2. In the **\\lon-svr1\docshare** folder, attempt to open the **Executives Only** document.
3. Verify that Carol does not have permission to open the document.

4. Log off from LON-CL1.

► **Task 4: Open and edit the rights-protected document as an authorized user at Trey Research.**

1. Log on to **LON-CL1** with the **Adatum\Aidan** account using the password **Pa\$\$w0rd**.
2. Open Microsoft Word 2010.
3. Create a new document named **\\LON-SVR1\docshare\TreyResearch-Confidential.docx**.
4. In the document, type the following text:

This document is for Trey Research only, it should not be modified.
5. Restrict the permission so that **april@treymresearch.net** is able to open the document.
6. Log on to **MUN-CL1** as **TREYRESEARCH\April**.
7. Use Windows Explorer to navigate to **\\LON-SVR1\docshare**. Use the credentials **Adatum\Administrator** and **Pa\$\$w0rd** to connect.
8. Copy the **TreyReserch-Confidential.docx** document to the desktop.
9. Attempt to open the document. When prompted enter the following credentials, select **Remember my credentials**, and then click **OK**:
 - Username: **April**
 - Password: **Pa\$\$w0rd**
10. Verify that you can open the document, but that you cannot make modifications to this document.
11. View the permissions that the **april@treymresearch.com** account has for the document.

Results: After completing this exercise, you should have verified that the AD RMS deployment is successful.

► **To prepare for the next module**

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

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

Lab Review

Question: What steps can you take to ensure that Information Rights Management can be used with the AD RMS role?

Module Review and Takeaways

Question: What are the benefits of having an SSL certificate installed on the AD RMS server when you are performing AD RMS configuration?

Question: You need to provide access to AD RMS–protected content to five users who are unaffiliated contractors, and are not members of your organization. Which method should you use to provide this access?

Question: You want to block users from protecting Office PowerPoint content using AD RMS templates. What steps should you take to accomplish this goal?

Best Practice

Prior to deploying AD RMS, you must analyze your organization’s business requirements and create the necessary templates. You should meet with users to inform them of AD RMS functionality and also ask for feedback on the types of templates that they would like to have available.

Strictly control membership of the Super Users group. Users in this group can access all protected content. Granting a user membership of this group gives them complete access to all AD RMS protected content.

Module 12

Implementing Active Directory Federation Services

Contents:

Module Overview	12-1
Lesson 1: Overview of AD FS	12-2
Lesson 2: Deploying AD FS	12-11
Lesson 3: Implementing AD FS for a Single Organization	12-17
Lesson 4: Deploying AD FS in a B2B Federation Scenario	12-23
Lab: Implementing AD FS	12-28
Module Review and Takeaways	12-36

Module Overview

Active Directory® Federation Services (AD FS) in Windows Server 2012 provides flexibility for organizations who want to enable their users to log on to applications that may be located on the local network, at a partner company, or in an online service. With AD FS, an organization can manage its own user accounts, and users only have to remember one set of credentials. However those credentials can be used to provide access to a variety of applications, which can be located in a variety of places.

This module provides an overview of AD FS, and then details how to configure AD FS in both a single organization scenario and in a partner organization scenario.

Objectives

After completing this module, you will be able to:

- Describe AD FS.
- Explain how to configure the AD FS prerequisites, and deploy the AD FS services.
- Describe how to implement AD FS for a single organization.
- Deploy AD FS in a business-to-business federation scenario.

Lesson 1

Overview of AD FS

AD FS is the Microsoft implementation of an identity federation framework that enables organizations to establish federation trusts and share resources across organizational and Active Directory Domain Services (AD DS) boundaries. AD FS is compliant with common web services standards, so as to enable interoperability with identity federation solutions provided by other vendors.

AD FS is designed to address a variety of business scenarios, where the typical authentication mechanisms used in a single organization do not work. This lesson provides an overview of the concepts and standards that are implemented in AD FS, and also the business scenarios that can be addressed with AD FS.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe Identity Federation.
- Describe claims-based identity.
- Describe web services.
- Describe AD FS.
- Explain how AD FS enables single sign-on (SSO) within a single organization.
- Explain how AD FS enables SSO between business partners.
- Explain how AD FS enables SSO between on-premises and cloud-based services.

What Is Identity Federation?

Identity federation enables the distribution of identification, authentication, and authorization across organizational and platform boundaries. You can implement identity federation within a single organization to enable access to diverse web applications, or between two organizations that have a relationship of trust between them.

To establish an identity federation partnership, both partners agree to create a federated trust relationship. This federated trust is based on an ongoing business relationship, and enables the organizations to implement business processes identified in the business relationship.

Identity Federation:

- Enables distributed identification, authentication, and authorization across organizational and platform boundaries.
- Requires a federated trust relationship between two organizations or entities.
- Enables organizations to retain control over who can access resources.
- Enables organizations to retain control of their user and group accounts.



Note: A federated trust is not the same as a forest trust that organizations can configure between Active Directory Domain Services (AD DS) forests. In a federated trust, the AD FS servers in two organizations never have to communicate directly with each other. In addition, all communication in a federation deployment occurs over HTTPS, so you do not need to open multiple ports on any firewalls to enable federation.

As a part of the federated trust, each partner defines what resources are accessible to the other organization, and how access to the resources is enabled. For example, to update a sales forecast, a sales

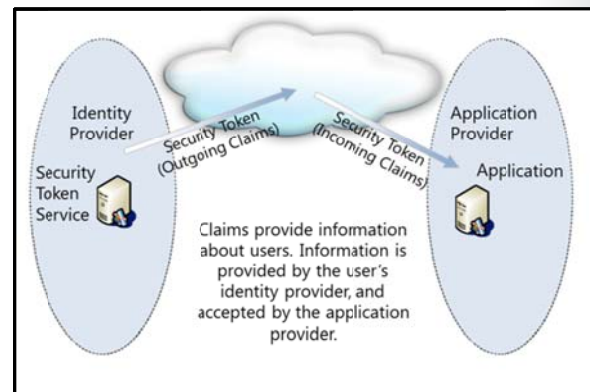
representative may need to collect information from a supplier's database that is hosted on the supplier's network. The administrator of the domain for the sales representative is responsible for ensuring that the appropriate sales representatives are members of the group requiring access to the supplier's database. The administrator of the organization where the database is located is responsible to ensure that the partner's employees only have access to the data they require.

In an identity federation solution, user identities and their associated credentials are stored, owned, and managed by the organization where the user is located. As part of the identity federation trust, each organization also defines how the user identities are shared securely to restrict access to resources. Each partner must define the services that it makes available to trusted partners and customers, and which other organizations and users it trusts. Each partner must also define what types of credentials and requests it accepts, and its privacy policies to ensure that private information is not accessible across the trust.

Identity federation can also be used within a single organization. For example, an organization may plan to deploy several web-based applications that require authentication. By using AD FS, the organization can implement one authentication solution for all of the applications, making it easy for users in multiple internal domains or forests to access the application. The solution can also be extended to external partners in the future, without changing the application.

What Is Claims-Based Identity?

Claims-based authentication is designed to address issues by extending typical authentication and authorization mechanisms outside the boundaries that are associated with that mechanism. For example, in most organizations, when users log on to the network, they are authenticated by an AD DS domain controller. A user who provides the right credentials to the domain controller is granted a security token. Applications that are running on servers in the same AD DS environment trust the security tokens that are provided by the AD DS domain controllers, because the servers can communicate with the same domain controllers where the users are authenticated.



The problem with this type of authentication is that it does not easily extend outside the boundaries of the AD DS forest. Although it is possible to implement Kerberos or NTLM-based trusts between two AD DS forests, client computers and domain controllers on both sides of the trust must communicate with domain controllers in the other forest to make decisions about authentication and authorization. This communication requires network traffic that is sent on multiple ports, so these ports must be open on all firewalls between the domain controllers and other computers. The problem becomes even more complicated when users have to access resources that are hosted in cloud-based systems, such as Windows® Azure™ or Microsoft® Office 365.

Claims-based authentication provides a mechanism for separating user authentication and authorization from individual applications. With claims-based authentication, users can authenticate to a directory service that is located within their organization, and be granted a claim based on that authentication. The claim can then be presented to an application that is running in a different organization. The application is designed to enable user access to the information or features, based on the claims presented. All communication also occurs over HTTPS.

The claim that is used in claims-based authentication is a statement about a user that is defined in one organization or technology, and trusted in another organization or technology. The claim could include a variety of information. For example, the claim could define the user's e-mail address, User Principal Name (UPN), and information about specific groups to which the user belongs. This information is collected from the authentication mechanism when the user successfully authenticates.

The organization that manages the application defines what types of claims will be accepted by the application. For example, the application may require the email address of the user to verify the user identity, and it may then use the group membership that is presented inside the claim to determine what level of access the user should have within the application.

Web Services Overview

For claims-based authentication to work, organizations have to agree on the format for exchanging claims. Rather than have each business define this format, a set of specifications broadly identified as *web services* has been developed. Any organization interested in implementing a federated identity solution can use this set of specifications.

Web services are a set of specifications that are used for building connected applications and services, whose functionality and interfaces are exposed to potential users through web technology standards such as Extensible Markup Language (XML), SOAP, Web Services Description Language (WSDL), and HTTP(S). The goal for creating web applications using web services is to simplify interoperability for applications across multiple development platforms, technologies, and networks.

To enhance interoperability, web services are defined by a set of industry standards. Web services are based on the following standards:

- Most web services use XML to transmit data through HTTP(S). With XML, developers can create their own customized tags, thereby facilitating the definition, transmission, validation, and interpretation of data between applications and between organizations.
- Web services expose useful functionality to web users through a standard web protocol. In most cases, the protocol used is SOAP, which is the communications protocol for XML web services. SOAP is a specification that defines the XML format for messages, and essentially describes what a valid XML document looks like.
- Web services provide a way to describe their interfaces in enough detail to enable a user to build a client application to communicate with the service. This description is usually provided in an XML document called a WSDL document. In other words, a WSDL file is an XML document that describes a set of SOAP messages, and how the messages are exchanged.
- Web services are registered so that potential users can find them easily. This is done with Universal Discovery Description and Integration (UDDI). A UDDI directory entry is an XML file that describes a business and the services it offers.

WS-* Security Specifications

There are many components included in web services specifications (also known as "WS-*" specifications). However, the most relevant specifications for an AD FS environment are the WS-Security specifications. The specifications that are part of the WS-Security specifications include the following:

Web services use a set of open specifications to develop applications that can interoperate across boundaries

Web services:

- Are developed using industry standards such as XML, SOAP, WSDL, and UDDI
- Define the security specifications used by identity federation systems
- Define the SAML standard for exchanging claims between federation partners

- **WS-Security - SOAP Message Security and X.509 Certificate Token Profile:** WS-Security describes enhancements to SOAP messaging that provide quality of protection through message integrity, message confidentiality, and single message authentication. WS-Security also provides a general-purpose—yet extensible—mechanism for associating security tokens with messages and a mechanism to encode binary security tokens—specifically X.509 certificates and Kerberos tickets—in SOAP messages.
- **WS-Trust:** WS-Trust defines extensions that build on WS-Security to request and issue security tokens and to manage trust relationships.
- **WS-Federation:** WS-Federation defines mechanisms that WS-Security can use to enable attribute-based identity, authentication, and authorization federation across different trust realms.
- **WS-Federation Passive Requestor Profile:** This WS-Security extension describes how passive clients such as web browsers can be authenticated and authorized, and how the clients can submit claims in a federation scenario. Passive requestors of this profile are limited to the HTTP or HTTPS protocol.
- **WS-Federation Active Requestor Profile:** This WS-Security extension describes how active clients, such as SOAP-based mobile device applications, can be authenticated and authorized, and how the clients can submit claims in a federation scenario.

Security Assertion Markup Language

The Security Assertion Markup Language (SAML) is an XML-based standard for exchanging claims between an identity provider and a service or application provider. SAML assumes that a user has been authenticated by an identity provider, and that the identity provider has populated the appropriate claim information in the security token. When the user is authenticated, the Identity Provider passes a SAML assertion to the service provider. On the basis of this assertion, the service provider can make authorization and personalization decisions within an application. The communication between federation servers is based around an XML document that stores the X.509 certificate for token-signing, and the SAML 1.1 or 2.0 token.

What Is AD FS?

(AD FS is the Microsoft implementation of an identity federation solution that uses claims-based authentication. AD FS provides the mechanisms to implement both the identity provider and the service provider components in an identity federation deployment.

AD FS provides the following features:

- **Enterprise claims provider for claims-based applications:** You can configure an AD FS server as a claims provider, which means that it can issue claims about authenticated users. This enables an organization to provide its users with access to claims-aware applications in another organization by using SSO.
- **Federation Service for identity federation across domains:** This service offers federated web SSO across domains, thereby enhancing security and reduces overhead for IT administrators.

AD FS is the Microsoft identity federation solution that can use claims-based authentication

AD FS includes the following features:

- Web SSO
- Web services interoperability
- Passive and smart client support
- Extensible architecture
- Enhanced security

The Windows Server 2012 version of AD FS includes:

- Integration with Dynamic Access Control
- Integration with the Windows Server 2012 operating system
- New Windows PowerShell cmdlets




Note: The Windows Server® 2012 version of AD FS is built on AD FS version 2.0, which is the second generation of AD FS released by Microsoft. The first version, AD FS 1.0, required

AD FS web agents to be installed on all web servers that were using AD FS, and provided both claims-aware and NT token-based authentication. AD FS 1.0 did not support active clients or SAML.

AD FS Features

The following are some of the key features of AD FS:

- **Web SSO:** Many organizations have deployed AD DS. After authenticating to AD DS through Integrated Windows authentication, users can access all other resources that they have permission to access within the AD DS forest boundaries. AD FS extends this capability to intranet or Internet-facing applications, enabling customers, partners, and suppliers to have a similar, streamlined user experience when they access an organization's web-based applications.
- **Web services interoperability:** AD FS is compatible with the web services specifications. AD FS employs the federation specification of WS-*, called WS-Federation. WS-Federation makes it possible for environments that do not use the Windows identity model to federate with Windows environments.
- **Passive and smart client support:** Because AD FS is based on the WS-* architecture, it supports federated communications between any WS-enabled endpoints, including communications between servers and passive clients, such as browsers. AD FS on Windows Server 2012 also enables access for SOAP-based smart clients, such as servers, mobile phones, personal digital assistants (PDAs), and desktop applications. AD FS implements the WS-Federation Passive Requestor Profile and WS-Federation Active Requestor Profile standards for client support.
- **Extensible architecture:** AD FS provides an extensible architecture that supports various security token types, including SAML and Kerberos authentication, and the ability to perform custom claims transformations. For example, AD FS can convert from one token type to another, or add custom business logic as a variable in an access request. Organizations can use this extensibility to modify AD FS to coexist with their current security infrastructure and business policies.
- **Enhanced security:** AD FS also increases the security of federated solutions by delegating responsibility of account management to the organization closest to the user. Each individual organization in a federation continues to manage its own identities, and is capable of securely sharing and accepting identities and credentials from other members' sources.

 **Additional Reading:** For information on the different identity federation products that can interoperate with AD FS, and for step-by-step guides on how to configure the products, see the AD FS 2.0 Step-by-Step and How To Guides, located at <http://technet.microsoft.com/en-us/library/adfs2-step-by-step-guides%28v=ws.10%29.aspx>.

New Features in Windows Server 2012 AD FS

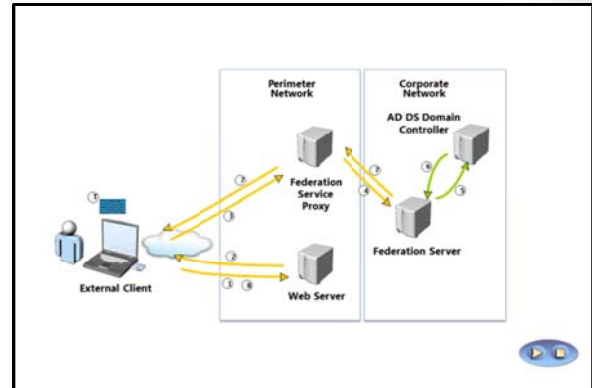
The version of AD FS that is shipping with Windows Server 2012 includes several new features:

- **Integration with the Windows Server 2012 operating system.** In Windows Server 2012, AD FS is included as a server role that you can install using Server Manager. When you install the server role, all required operating system components are installed automatically.
- **Integration with Dynamic Access Control (DAC).** When you deploy DAC, you can configure user and device claims that are issued by AD DS domain controllers. AD FS can consume the AD DS claims that the domain controllers issue. This means that AD FS can make authorization decisions based on both user accounts and computer accounts.
- **Windows PowerShell® cmdlets for administering AD FS.** Windows Server 2012 provides several new cmdlets that you can use to install and configure the AD FS server role.

How AD FS Enables SSO in a Single Organization

For many organizations, configuring access to applications and services may not require an AD FS deployment. If all users are members of the same AD DS forest, and if all applications are running on servers that are members of the same forest, you can usually just use AD DS authentication to provide access to the application. However, there are several scenarios where you can use AD FS to optimize the user experience by enabling SSO:

- The applications may not be running on Windows servers or on any servers that support AD DS authentication, or on Windows Server servers that are not domain-joined. The applications may require SAML or web services for authentication and authorization.
- Large organizations frequently have multiple domains and forests that may be the results of mergers and acquisitions, or due to security requirements. Users in multiple forests might require access to the same applications.
- Users from outside the office might require access to applications that are running on internal servers. The external users may be logging on to the applications from computers that are not part of the internal domain.



Note: Implementing AD FS does not necessarily mean that users are not prompted for authentication when they access applications. Depending on the scenario, users may be prompted for their credentials. However, users always authenticate using their internal credentials in the trusted account domain, and they never need to remember alternate credentials for the application. In addition, the internal credentials are never presented to the application or to the partner AD FS server.

Organizations can use AD FS to enable SSO in these scenarios. Because all users and the application are in the same AD DS forest, the organization only has to deploy a single federation server. This server can operate as the claims provider so that it authenticates user requests and issues the claims. The same server is also the relying party, or the consumer of the claims to provide authorization for application access.

Note: The slide and the following description use the terms federation server and federation service proxy to describe AD FS server roles. The *federation server* is responsible for issuing claims, and in this scenario, is also responsible for consuming the claims. The *Federation Service Proxy* is a proxy component that is recommended for deployments where users outside the network need access to the AD FS environment. These components are covered in more detail in the next lesson.

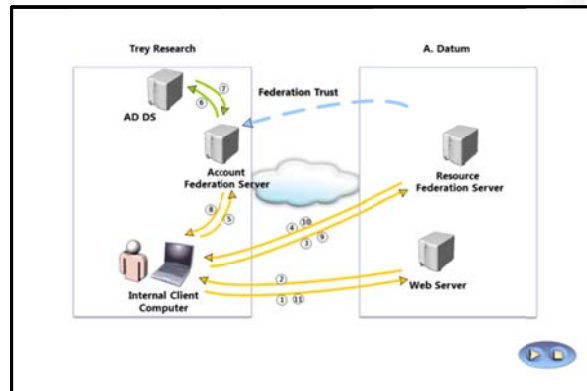
The following steps describe the communication flow in this scenario.

1. The client computer, which is located outside the network, must access a web-based application on the web server. The client computer sends an HTTPS request to the web server.
2. The web server receives the request, and identifies that the client computer does not have a claim. The web server redirects the client computer to the Federation Service Proxy.

3. The client computer sends an HTTPS request to the Federation Service Proxy. Depending on the scenario, the Federation Service Proxy may prompt the user for authentication, or use Integrated Windows authentication to collect the user credentials.
4. The Federation Service Proxy passes on the request and the credentials to the federation server.
5. The federation server uses AD DS to authenticate the user.
6. If authentication is successful, the federation server collects AD DS information about the user, which is then used to generate the user's claims.
7. If the authentication is successful, the authentication information and other information is collected in a security token and passed back to the client computer, through the Federation Service Proxy.
8. The client then presents the token to the web server. The web resource receives the request, validates the signed tokens, and uses the claims in the user's token to provide access to the application.

How AD FS Enables SSO in a Business-to-Business Federation

One of the most common scenarios for deploying AD FS is to provide SSO in a business-to-business (B2B) federation. In the scenario, the organization that requires access to another organization's application or service can manage their own user accounts and define their own authentication mechanisms. The other organization can define what applications and services are exposed to users outside the organization, and what claims it accepts to provide access to the application. To enable application or service sharing in this scenario, the organizations have to establish a federation trust, and then define the rules for exchange claims between the two organizations.



The slide for this topic demonstrates the flow of traffic in a federated B2B scenario using a claims-aware web application. In this scenario, users at Trey Research have to access a web-based application at A. Datum Corporation. The AD FS authentication process for this scenario is as follows:

1. A user at Trey Research uses a web browser to establish an HTTPS connection to the web server at A. Datum Corporation.
2. The web application receives the request and verifies that the user does not have a valid token stored in a cookie by the web browser. Because the user is not authenticated, the web application redirects the client to the federation server at A. Datum (by using an HTTP 302 redirect message).
3. The client computer sends an HTTPS request to the A. Datum Corporation's federation server. The federation server determines the home realm for the user. In this case, the home realm is Trey Research.
4. The client computer is redirected again to the federation server in the user's home realm, Trey Research.
5. The client computer sends an HTTPS request to the Trey Research federation server.
6. If the user is already logged on to the domain, the federation server can take the user's Kerberos ticket and request authentication from AD DS on the user's behalf, using Integrated Windows authentication. If the user is not logged onto their domain, the user is prompted for credentials.

7. The AD DS domain controller authenticates the user, and sends the success message back to the federation server, along with other information about the user that can be used to generate the user's claims.
8. The federation server creates the claim for the user based on the rules defined for the federation partner. The claims data is placed in a digitally-signed security token, and then sent to the client computer, which posts it back to A. Datum Corporation's federation server.
9. A. Datum Corporation's federation server validates that the security token came from a trusted federation partner.
10. A. Datum Corporation's federation server creates and signs a new token, which it sends to the client computer, which then sends the token back to the original URL requested.
11. The application on the web server receives the request and validates the signed tokens. The web server issues the client a session cookie indicating that it has been successfully authenticated, and a file-based persistent cookie is issued by the federation server (good for 30 days by default) to eliminate the home realm discovery step during the cookie lifetime. The server then provides access to the application, based on the claims provided by the user.

How AD FS Enables SSO with Online Services

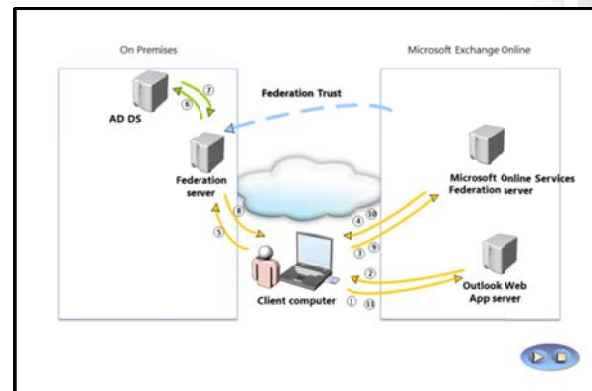
As organizations move services and applications to cloud-based services, it is increasingly important that these organizations have some way to simplify the authentication and authorization experience for their users as they consume the cloud-based services. Cloud-based services add another level of complexity to the IT environment, as they are located outside the direct administrative control of the IT administrators, and may be running on many different platforms.

You can use AD FS to provide an SSO experience to users across the various cloud-based platforms available. For example, once users are authenticated with AD DS credentials, they could then access Microsoft Online Services, such as hosted Microsoft Exchange Online or SharePoint® Online, by using those domain credentials.

AD FS can also provide SSO to non-Microsoft cloud providers. Because AD FS is based on open standards, it can interoperate with any compliant claims-based system.

The process for accessing a cloud-based application is quite similar to the B2B scenario. One example of a cloud-based service that uses AD FS for authentication is a hybrid Exchange Online deployment. In this type of deployment, an organization deploys some or all of their mailboxes in an Office 365 and Exchange Online environment. However, the organization manages all of their user accounts in their on-premises AD DS environment. The deployment uses the Microsoft Online Services Directory Synchronization Tool to synchronize user account information from the on-premises deployment to the Exchange Online deployment.

When users try to log on to their Exchange Online mailbox, the user must be authenticated using their internal AD DS credentials. If users try to log on directly to the Exchange Online environment, they are redirected back to the internal AD FS deployment to authenticate before they are given access.



The following steps describe what happens when a user tries to access their online mailbox using a web browser:

1. The user opens a web browser and sends an HTTPS request to the Exchange Online Microsoft Outlook® Web App server.
2. The Outlook Web App server receives the request and verifies whether the user is part of a hybrid Exchange Server deployment. If this is the case, the server redirects the client computer to the Microsoft Online Services federation server.
3. The client computer sends an HTTPS request to the Microsoft Online Services federation server.
4. The client computer is redirected again to the on-premises federation server.
5. The client computer sends an HTTPS request to the on-premises federation server.
6. If the client computer is already logged on to the domain, the on-premises federation server can take the user's Kerberos ticket and request authentication from AD DS on the user's behalf, using Integrated Windows authentication. If the user is logging on from outside the network or from a computer that is not a member of the internal domain, the user is prompted for credentials.
7. The AD DS domain controller authenticates the user, and sends the success message back to the federation server, along with other information about the user that the federation server can use to generate the user's claims.
8. The federation server creates the claim for the user based on the rules defined during the AD FS server setup. The claims data is placed in a digitally-signed security token, and then sent to the client computer, which posts it back to the Microsoft Online Services federation server.
9. The Microsoft Online Services federation server validates that the security token came from a trusted federation partner. This trust is configured when you configure the hybrid Exchange Server environment.
10. The Microsoft Online Services federation server creates and signs a new token, which it sends to the client computer, which then sends the token back to the Outlook Web App server.
11. The Outlook Web App server receives the request and validates the signed tokens. The server issues the client a session cookie indicating that it has authenticated successfully. The user is then granted access to their Exchange Server mailbox.

Lesson 2

Deploying AD FS

After you understand how AD FS works, the next step is deploying the service. Before deploying AD FS, you must understand the components that you will need to deploy, and the prerequisites that you must meet, particularly in regard to certificates. This lesson provides an overview of deploying the AD FS server role in Windows Server 2012.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe the components that you can include in an AD FS deployment.
- List the prerequisites for an AD FS deployment.
- Describe the Public Key Infrastructure (PKI) and certificate requirements for the AD FS deployment.
- Describe the AD FS federation server roles.
- Install the AD FS server role.

AD FS Components

AD FS is installed as a server role in Windows Server 2012. However, there are many different components that you install and configure in an AD FS deployment.

The following table lists the AD FS components.

AD FS Components	
• Federation Server	• Relying Parties
• Federation Server Proxy	• Claims Provider Trust
• Claims	• Relying Party Trust
• Claim Rules	• Certificates
• Attribute Store	• Endpoints
• Claims Providers	

Component	What does it do?
Federation server	The federation server issues, manages, and validates requests involving identity claims. All implementations of AD FS require at least one Federation Service for each participating forest.
Federation server proxy	The federation server proxy is an optional component that you usually deploy in a perimeter network. It does not add any functionality to the AD FS deployment, but is deployed just to provide a layer of security for connections from the Internet to the federation server.
Claims	A claim is a statement that is made by a trusted entity about an object such as a user. The claim could include the user's name, job title, or any other factor that might be used in an authentication scenario. With Windows Server 2012, the object can also be a device used in a DAC deployment.
Claim rules	Claim rules determine how claims are processed by the federation servers. For example, a claim rule may state that an email address is accepted as a valid claim, or that a group name from one organization is translated into an application-specific role in the other organization. The rules are usually

Component	What does it do?
	processed in real time as claims are made.
Attribute store	AD FS uses an attribute store to look up claim values. AD DS is a common attribute store and is available by default if AD FS is installed on a domain-joined server.
Claims providers	The claims provider is the server that issues claims and authenticates users. A claims provider enables one side of the AD FS authentication and authorization process. The claims provider manages the user authentication, and then issues the claims that the user presents to a relying party.
Relying parties	The relying party is where the application is located, and it enables the second side of the AD FS authentication and authorization process. The relying party is a web service that consumes claims from the claims provider. The relying party server must have the Microsoft Windows Identity Foundation installed, or use the AD FS 1.0 claims-aware agent.
Claims provider trust	Configuration data that defines rules under which a client may request claims from a claims provider and subsequently submit them to a relying party. The trust consists of various identifiers such as names, groups and various rules.
Relying party trust	The AD FS configuration data that is used to provide claims about a user or client to a relying party. It consists of various identifiers, such as names, groups, and various rules.
Certificates	AD FS uses digital certificates when communicating over SSL or as part of the token issuing process, the token receiving process, and the metadata publishing process. Digital certificates are also used for token signing.
Endpoints	Endpoints are mechanisms that enable access to the AD FS technologies including token issuance and metadata publishing. AD FS comes with built-in endpoints that are responsible for a specific functionality.



Note: Many of these components are described in more detail throughout the remainder of this module.

AD FS Prerequisites

Before deploying AD FS, you must ensure that your internal network meets some basic prerequisites. The configuration of the following network services is critical for a successful AD FS deployment:

- Network connectivity: The following network connectivity is required:
 - The client computer must be able to communicate with the web application, the resource federation server or federation server proxy, and the account federation server or federation proxy using HTTPS.

Infrastructure critical to a successful AD FS deployment includes:

- TCP/IP network connectivity
- AD DS
- Attribute stores
- DNS
- Compatible operating systems



- The federation server proxies must be able to communicate with the federation servers in the same organization using HTTPS
- Federation servers and internal client computers must be able to communicate with domain controllers for authentication.
- AD DS: AD DS is a critical piece of AD FS. Domain controllers should be running Windows Server 2003 Service Pack 1 (SP1) as a minimum. Federation servers must be joined to an AD DS domain. The Federation Service proxy does not have to be domain-joined. Although you can install AD FS on a domain controller, it is not recommended due to security implications.
- Attribute stores: AD FS uses an attribute store to build claim information. The attribute store contains information about users, which is extracted from the store by the AD FS server after the user has been authenticated. AD FS supports the following attribute stores:
 - Active Directory Application Mode (ADAM) in Windows Server 2003
 - Active Directory Lightweight Directory Services (AD LDS) in Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012
 - Microsoft SQL Server® 2005 (all editions)
 - Microsoft SQL Server 2008 (all editions)
 - A custom attribute store



Note: AD DS can be used both as the authentication provider and as an attribute store. AD FS can also use AD LDS as an attribute store. In AD FS 1.x, AD LDS can be used as an authentication store, but in the current version of AD FS, AD LDS can only be used as an attribute store.

- DNS: Name resolution allows clients to find federation servers. The client computers must resolve the DNS names for all federation servers to which they connect, and the web applications that the client computer is trying to use. If the client computer is external to the network, the client computer must resolve the DNS name for the Federation Service Proxy, not the internal federation server. The Federation Service Proxy must resolve the name of the internal federation server. If internal users have to access the internal federation server directly, and external users have to connect through the federation server proxy, you will need to configure different DNS records in the internal and external DNS zones.
- Operating system prerequisites: You can only deploy the Windows Server 2012 version of AD FS as a server role on a Windows Server 2012 server.

PKI and Certificate Requirements

AD FS is designed to enable computers to communicate securely, even though they may be located in different locations. In this scenario, most of the communications between computers passes through the Internet. To provide security for the network traffic, all communications are protected using Secure Sockets Layer (SSL). This factor means that it is important to correctly choose and assign SSL certificates to the AD FS servers. To provide SSL security, AD FS servers use certificates as service communication certificates, token-signing certificates, and token-decrypting certificates.

AD FS federation services require:

- Service communication certificates
- Token-signing certificates
- Token-decrypting certificates

When choosing certificates, ensure that the service communication certificate and the token-signing certificate are trusted by all federation partners and clients

Service Communication Certificates

You use a service communication certificate to secure SSL communications to the websites running on the AD FS server. This certificate is bound to the default website on the AD FS server. You can choose which certificate to use when you configure the AD FS server role on the server, and can change the assigned certificate after deployment by using the AD FS console. This certificate is also called a server authentication certificate.

Token-Signing Certificates

The token-signing certificate is used to sign every token that a federation server issues. This certificate is critical in an AD FS deployment because the token signature indicates which federation server issued the token. This certificate is used by the claims provider to identify itself, and it is used by the relying party to verify that the token is coming from a trusted federation partner.

The relying party also requires a token-signing certificate to sign the tokens that it prepares for other AD FS components, such as web applications and clients. These tokens must be signed by the relying party's token-signing certificate to be validated by the destination applications.

When you configure a federation server, the server assigns a self-signed certificate as the token-signing certificate. Because no other parties trust the self-signed certificate, you might choose to replace the self-signed certificate with a trusted certificate. As an alternative, you can configure all federation servers in partner organizations to trust the self-signed certificate. You can have multiple token-signing certificates configured on the federation server, but only the primary certificate is used to sign tokens.

Token-Decrypting Certificates

Token-decrypting certificates are used to encrypt the entire user token before transmitting the token across the network. To provide this functionality, the public key from the relying party federation server certificate must be provided to the claims provider federation server. The certificate is sent without the private key. The claims provider server uses the public key from the certificate to encrypt the user token. When the token is returned to the relying party federation server, it uses the private key from the certificate to decrypt the token. This provides an extra layer of security when transmitting the certificates across the Internet.

When you configure a federation server, the server assigns a self-signed certificate as the token-decrypting certificate. Because no other parties have to trust this certificate, it is possible to continue to use this certificate without replacing it with a trusted certificate.



Note: Federation server proxies only require a service communication certificate. The certificate is used to enable SSL communication for all client connections. Because the federation


server proxy does not issue any tokens, it does not need the other two types of certificates. Web servers that are deployed as part of an AD FS deployment should also be configured with SSL server certificates to enable secure communications with client computers.

Choosing a Certification Authority

AD FS federation servers can use self-signed certificates, certificates from an internal, private Certification Authority (CA), or certificates that have been purchased from an external, public CA.

In most AD FS deployments, the most important factor when choosing the certificates is that the certificates be trusted by all parties involved. This means that if you are configuring an AD FS deployment that interacts with other organizations, you are almost certainly going to use a public CA for the SSL certificate on federation server proxy, because the certificates issued by the public CA are trusted by all partners automatically.

If you are deploying AD FS just for your organization, and all servers and client computers are under your control, consider using a certificate from an internal, private CA. If you deploy an internal Enterprise CA on Windows Server 2012, you can use Group Policy to ensure that all computers in the organization automatically trust the certificates issued by the internal CA. Using an internal CA can significantly decrease the cost of the certificates.

 **Note:** Deploying an internal CA using Active Directory Certificate Services (AD CS) is a straightforward process, but it is critical that the deployment be planned and implemented carefully.

Federation Server Roles

When you install the AD FS server role, you can configure the server as either a federation server or federation server proxy. After installing the federation server role, you can configure the server as either a Claims Provider, a Relying Party, or both. These server functions are as follows:

- **Claims provider:** A claims provider is a federation server that provides to users signed tokens that contain claims. Claims provider federation servers are deployed in organizations where user accounts are located. When a user requests a token, the claims provider federation server verifies the user authentication using AD DS, and then collects information from an attribute store, such as AD DS or AD LDS, to populate the user claim with the attributes required by the partner organization. The server issues tokens in SAML format. The claims provider federation server also protects the contents of security tokens in transit, by signing and optionally encrypting them.
- **Relying Party:** A relying party is a federation server that receives security tokens from a trusted claims provider. The relying party federation servers are deployed in organizations that provide application access to claims provider organizations. The relying party accepts and validates the claim, and then issues new security tokens that the web server can use to provide appropriate access to the application.

Claims provider federation server:

- Authenticates internal users
- Issues signed tokens containing user claims

Relying party federation server:

- Consumes tokens from the claims provider
- Issues tokens for application access

Federation server proxy:

- Is deployed in a perimeter network
- Provides a layer of security for internal federation servers



Note: A single AD FS server can operate as both a claims provider and a relying party, even with the same partner organizations. The AD FS server functions as a claims provider when it is authenticating users and providing tokens for another organization, but it can also accept tokens from the same or another organization in a relying party role.

- **Federation Server Proxy:** A federation server proxy provides an extra level of security for AD FS traffic that is coming from the Internet to the internal AD FS federation servers. Federation server proxies can be deployed in both the claims provider and relying party organizations. On the claims provider side, the proxy collects the authentication information from client computers and passes it to the claims provider federation server for processing. The federation server issues a security token to the proxy, which sends it to the relying party proxy. The relying party federation server proxy accepts these tokens, and then passes them on to the internal federation server. The relying party federation server then issues a security token for the web application, and then sends the token to the federation server proxy, which then forwards the token to the client. The federation server proxy does not provide any tokens or create claims; it only forwards requests from clients to internal AD FS servers. All communication between the federation server proxy and the federation server uses HTTPS.



Note: A federation server proxy cannot be configured as a claims provider or a relying party. The claims provider and relying party must be members of an AD DS domain. The federation server proxy can be configured as a member of a workgroup, or as a member of an extranet forest, and deployed in a perimeter network.

Demonstration: Installing the AD FS Server Role

In this demonstration, you will see how to install and complete the initial configuration of the AD FS server role in Windows Server 2012. The instructor will install the server role, and then run the AD FS Federation Server Configuration Wizard to configure the server as a standalone federation server.

Demonstration Steps

Install the AD FS Server Role

- On LON-DC1, in Server Manager, add the Active Directory Federation Services server role.

Configure the AD FS Server Role

1. Run the AD FS Federation Server Configuration Wizard using the following parameters:
 - Create a new federation service
 - Create a stand-alone deployment
 - Use the LON-DC1.Adatum certificate.
 - Choose the service name LON-DC1.Adatum.com
2. Open Internet Explorer and connect to <https://lon-dc1.adatum.com/federationmetadata/2007-06/federationmetadata.xml>.

Lesson 3

Implementing AD FS for a Single Organization

The simplest deployment scenario for AD FS is within a single organization. In this scenario, a single AD FS server can operate both as the claims provider and as the relying party. All users in this scenario are internal to the organization, as is the application that the users are accessing.

This lesson provides details on the components that are required to configure AD FS in a single organization deployment of AD FS. These components include configuring claims, claim rules, claims provider trusts, and relying party trusts.

Lesson Objectives

After completing this lesson, you will be able to:

- Describe AD FS claims.
- Describe AD FS claim rules.
- Describe claims provider trusts.
- Describe relying party trusts.
- Configure claims provider and relying party trusts.

What Are AD FS Claims?

AD FS claims provide the link between the claims provider and relying party roles in an AD FS deployment. An *AD FS claim* is a statement made about a particular subject (such as a user) by a trusted entity (such as a claims provider). The claims provider creates the claims and the relying party consumes the claims. AD FS claims provide a standards-based and flexible way for claims provider organizations to provide specific information about users in their organizations, and a way for relying parties to define exactly what information they require to provide application access. The claim information provides the details required by applications to enable access to claims-aware applications.

Claims provide information about users from the claims provider to the relying partner

AD FS:


- Provides a default set of built-in claims
- Enables the creation of custom claims
- Requires that each claim have a unique URI

Claims can be:

- Retrieved from an attribute store
- Calculated based on retrieved values
- Transformed into alternate values

Claim Types

Each AD FS claim has a claim type, such as email address, UPN, or last name. Users can be issued claims based on any defined claim type. Therefore, a user might be issued a claim with a type of Last Name and a value of, for example, Weber. AD FS provides several built-in claim types. Optionally, you can create new ones based on the organization requirements.

 **Note:** In AD FS 1.0, you could configure claims as identity claims, group claims, or custom claims. These claim types do not apply to AD FS 2.0 or later. Essentially, all claims are now considered custom claims.

Each AD FS claim type is identified by a Uniform Resource Identifier (URI) that uniquely identifies the claim type. This information is provided as part of the AD FS server metadata. For example, if the claims

provider organization and the relying party organization decide to use a claim type of AccountNumber, both organizations must configure a claim type with this name. The claim type is published and the claim type URI must be identical on both AD FS servers.

How Claim Values are Populated

The claims issued by a claims provider contain the information that is required by the relying party to enable appropriate application access. One of the first steps in planning an AD FS deployment is to define exactly what information the applications must have about each user, to provide that user access to the application. Once this information is defined, the claims are then defined on the claims provider federation server. The information required to populate the claim can be obtained in several ways:

- The claim can be retrieved from an attribute store. Frequently, the information required for the claim is already stored in an attribute store that is available to the federation server. For example, an organization might decide that the claim should include the user's UPN, email address, and specific group memberships. This information is already stored in AD DS, so the federation server can just retrieve this information from AD DS when creating the claim. Because AD FS can use AD DS, AD LDS, SQL Server, a non-Microsoft Lightweight Directory Access Protocol (LDAP) directory, or a custom attribute store to populate claims, you can define almost any value within the claim.
- The claim can be calculated based on collected information. Claims provider federation servers can also calculate information based on information that is gathered from an attribute store. For example, you may want to provide information about a person's salary within a claim. This information is likely stored in a Human Resources database, but the actual value may be considered confidential. You can define a claim that categorizes salaries within an organization, and then have the AD FS server calculate to which category a specific user belongs. In this way, the claim only includes the salary category information, not the actual user salary.
- The claim can be transformed from one value to another. In some cases, the information that is stored in an attribute store does not exactly match the information required by the application when making authorization information. For example, the application may have different user roles defined that do not directly match the attributes that are stored in any attribute store. However, the application role may correlate to AD DS group membership. For example, users in the Sales group may correlate to one application role, while users in the Sales Management group may correlate to a different application role. To establish the correlation in AD FS, you can configure a claims transformation that takes the value provided by the claims provider and translates the value into to a claim that is useful to the application in the relying party.
- If you have deployed DAC, a DAC device claim can be transformed into an AD FS claim. This can be used to ensure that users can access AD FS Web site only from trusted workstations that have been issued a valid device claim.

What Are AD FS Claim Rules?

Claim rules define how claims are sent and consumed by AD FS servers. Claim rules define the business logic that is applied to claims that are provided by claims providers, and to claims that are accepted by the relying parties. You can use claim rules to:

- Define which incoming claims are accepted from one or more claims providers.
- Define which outbound claims are provided to one or more relying parties.
- Apply authorization rules to enable access to a specific relying party for one or more users or groups of users.

- Claim rules define how claims are sent and consumed by AD FS servers
- Claims provider rules are acceptance transform rules
- Relying party rules can be:
 - Issuance transform rules
 - Issuance authorization rules
 - Delegation authorization rules
- AD FS servers provide default claim rules, templates, and a syntax for creating claim rules

You can define two types of claim rules:

- Claim rules for a claims provider trust. A claims provider trust is the AD FS trust relationship that is configured between an AD FS server and a claims provider. You can configure claim rules to define how the claims provider processes and issues claims.
- Claim rules for a relying party trust. A relying party trust is the AD FS trust relationship that is configured between an AD FS server and a relying party. You can configure claim rules that define how the relying party accepts claims from the claims provider.

Claim rules on an AD FS claims provider are all considered acceptance transform rules. These rules determine what claim types are accepted from the claims provider, and then sent to a relying party trust. When configuring AD FS within a single organization, there is a default claims provider trust that is configured with the local AD DS domain. This rule set defines the claims that are accepted from AD DS.

There are three types of claim rules for a relying party trust:

- Issuance Transform Rules: These rules define the claims that are sent to the relying party that has been defined in the relying party trust.
- Issuance Authorization Rules: These rules define which users are permitted or denied access to the relying party defined in the relying party trust. This rule set can include rules that explicitly permit access to a relying party, and/or rules that explicitly deny access to a relying party.
- Delegation Authorization Rules: These rules define the claims that specify which users can act on behalf of other users when accessing the relying party. This rule set can include rules that explicitly permit delegates for a relying party, or rules that explicitly deny delegates to a relying party.



Note: A single claim rule can only be associated with a single federated trust relationship. This means that you cannot create a set of rules for one trust and then re-use those rules for other trusts that you configure on your federation server.

AD FS servers are preconfigured with a set of default rules and several default templates that you can use to create the most common claim rules. You can also create custom claim rules using the AD FS claim rule language.

What Is a Claims Provider Trust?

A claims provider trust is configured on the relying party federation server. The claims provider trust identifies the claims provider, and describes how the relying party consumes the claims that the claims provider issues. You must configure a claims provider trust for each claims provider.

By default, an AD FS server is configured with a claims provider trust named Active Directory. This trust defines the claim rules, which are all acceptance transform rules that define how the AD FS server accepts AD DS credentials. For example, the default claim rules on the claims provider trust include rules that pass through the user names, security identifiers (SIDs) and group SIDs to the relying party. In a single organization AD FS deployment, where AD DS authenticates all users, the default claims provider trust may be the only required claims provider trust.

When you expand the AD FS deployment to include other organizations, you must create additional claims provider trusts for each federated organization. When configuring a claims provider trust, you have three options:

- Import data about the claims provider through the federation metadata. If the AD FS federation server or federation proxy server is accessible through the network from your AD FS federation server, you can enter the host name or URL for the partner federation server. Your AD FS federation server connects to the partner server, and downloads the federation metadata from the server. The federation metadata includes all the information that is required to configure the claims provider trust. As part of the federation metadata download, your federation server also downloads the SSL certificate that is used by the partner federation server.
- Import data about the claims provider from a file. Use this option if the partner federation server is not directly accessible from your federation server, but the partner organization has exported its configuration and provided you the information in a file. The configuration file must include the configuration information for the partner organization, as well as the SSL certificate that the partner federation server uses.
- Manually configure the claims provider trust. Use this option if you want to configure all of the settings for the claims provider trust directly. When you choose this option, you must provide the features that the claims provider supports, the URL used to access the claims provider AD FS servers, and add the SSL certificate that the partner organization uses.

- Claims provider trusts:
 - Are configured on the relying party federation server
 - Identify the claims provider
 - Configure the claim rules for the claims provider
- In a single organization scenario, a claims provider trust called Active Directory defines how AD DS user credentials are processed
- Additional claims provider trusts can be configured:
 - By importing the federation metadata
 - By importing a configuration file
 - By configuring the trust manually

What Is a Relying Party Trust?

A relying party trust is defined on the claims provider federation server. The relying party trust identifies the relying party, and also defines the claims rules that define how the relying party accepts and processes claims from the claims provider.

In a single organization scenario, the relying party trust defines how the AD FS server interacts with the applications deployed within the application. When you configure the relying party trust in a single organization, you provide the URL for the internal application, and configure settings such as

whether the application supports SAML 2.0 or whether it requires AD FS 1.0 tokens, the SSL certificate and URL used by the web server, and the issuance authorization rules for the application.

The process for configuring relying party trust is similar to that for the claims provider trust. When you expand the AD FS deployment to include other organizations, you must create additional relying party trusts for each federated organization. When configuring a relying party trust, you have three options:

- Import data about the relying party through the federation metadata. If the AD FS federation server or federation proxy server is accessible through the network from your AD FS federation server, you can enter the host name or URL for the partner federation server. Your AD FS federation server connects to the partner server, and then downloads the federation metadata from the server. The federation metadata includes all the information that is required to configure the relying party trust. As part of the federation metadata download, your federation server also downloads the SSL certificate that the partner federation server uses.
- Import data about the relying party from a file. Use this option if the partner federation server is not accessible from your federation server directly. In this case, the partner organization can export its configuration information to a file, and then provide it to you. The configuration file must include the configuration information for the partner organization, and the SSL certificate that the partner federation server uses.

Manually configure the claims provider trust. Use this option if you want to configure all of the settings for the claims provider trust directly.

- Relying party trusts:
 - Are configured on the claims provider federation server
 - Identify the relying party
 - Configure the claim rules for the relying party
- In a single organization scenario, a relying party trust defines the connection to internal applications
- Additional relying party trusts can be configured:
 - By importing the federation metadata
 - By importing a configuration file
 - By manually configuring the trust

Demonstration: Configuring Claims Provider and Relying Party Trusts

In this demonstration, you will see how to configure claims provider trusts and relying party trusts. The instructor will show how to edit the default Active Directory claims provider trust. The instructor will also create a new relying party trust, and demonstrate how to configure the trust.

Demonstration Steps

Configure a Claims Provider Trust

1. In the AD FS console, go to the **Claims Provider Trusts**, highlight the **Active Directory** store, and then click **Edit Claim Rules**.
2. In the **Edit Claim Rules for Active Directory** dialog box, on the **Acceptance Transform Rules** tab, start the **Add Transform Claim Rule Wizard** and complete the wizard with the following settings:
3. Under Claim rule template select **Send LDAP Attributes as Claims**.

4. Name the claim rule **Outbound LDAP Attribute Rule**.
5. Choose Active Directory as the Attribute Store.
6. In the **Mapping of LDAP attributes to outgoing claim types** select the following values:
 - E-Mail-Addresses to **E-Mail Address**
 - User-Principal-Name to **UPN**

Configure a Windows Identity Foundation Application for AD FS

1. On LON-SVR1, from the Start screen, start the Windows Identity Foundation Federation Utility.
2. Complete the wizard with the following settings:
 - Point to the web.config file sample application by browsing to **C:\inetpub\wwwroot\AdatumTestApp\web.config**.
 - Specify an Application URI box by typing **https://lon-svr1.adatum.com/AdatumTestApp/**.
 - Select Use an existing STS, and then enter the path **https://lon-dc1.adatum.com/federationmetadata/2007-06/federationmetadata.xml**.
 - Disable certificate chain validation.
 - Select **No encryption**.

Configure a Relying Party Trust

1. In the AD FS Management console, in the middle pane, click **Required: Add a trusted relying party**.
2. Complete the Add Relying Party Wizard with the following settings:
 - Select **Import data about the relying party published online or on a local network**, and type **https://lon-svr1.adatum.com/adatumtestapp**.
 - Specify a Display name of **ADatum Test App**.
 - Select **Permit all users to access this relying party**.
 - Ensure that the **Edit Claim Rules for ADatum Test App** check box is selected when the wizard is complete.

Lesson 4

Deploying AD FS in a B2B Federation Scenario

A second common scenario for implementing AD FS is in a B2B federation scenario. In this scenario, users in one organization require access to an application in another organization. AD FS in this scenario enables SSO. This way, users always log on to their home AD DS environment, but are granted access to the partner application based on the claims acquired from their local AD FS server.

Configuring AD FS in a B2B federation scenario is quite similar to configuring AD FS in a single organization scenario. The primary difference is that now both the claims provider trusts and the relying party trusts refer to external organizations, rather than internal AD DS or application.

Lesson Objectives

After completing this lesson, you will be able to:

- Configure the account partner in a B2B federation scenario.
- Configure the resource partner in a B2B federation scenario.

Explain how to configure claims rules for a B2B federation scenario.

- Explain how home realm discovery works.
- Configure claims rules.

Configuring an Account Partner

In a B2B AD FS scenario, the terminology that you use to describe the two partners involved in the AD FS deployment changes slightly. In this scenario, the claims provider organization is also called the account partner organization. An *account partner organization* is the organization in which the user accounts are stored in an attribute store. An account partner handles the following tasks:

- Gather credentials from users who are using a web-based service, and then authenticating those credentials.
- Build up claims for users, and then package the claims into security tokens. The tokens can then be presented across a federation trust to gain access to federation resources that are located at the resource partner organization.

Configuring the account partner organization to prepare for federation involves the following steps:

1. Implement the physical topology for the account partner deployment. This step could include deciding on the number of federation servers and federation server proxies to deploy the locations to deploy them to, and configuring the required DNS records and certificates.
2. Add an attribute store. Use the AD FS management console to add the attribute store. In most cases, you use the default Active Directory attribute store (which must be used for authentication), but you can also add other attribute stores if required to build the user claims.
3. Connect to a resource partner organization by creating a relying party trust. The easiest way to do this is to use the federation metadata URL that is provided by the resource partner organization. With

An account partner is a claims provider in a B2B federation scenario

To configure an account partner:

1. Implement the physical topology
2. Add an attribute store
3. Configure a relying party trust
4. Add a claim description
5. Prepare client computers for federation

this option, your AD FS server automatically collects the information required for the relying party trust.

4. Add a claim description. The claim description lists the claims that your organization provides to the relying partner. This information may include user names, email addresses, group membership information, or other identifying information about a user.
5. Prepare client computers for federation. This may involve two steps:
 - a. Add the account partner federation server. In the browser of client computers, add the account partner federation server to the Local Intranet list. By adding the account partner federation server to the Local Intranet list on the client computers, you enable Integrated Windows authentication, which means that users are not prompted for authentication if they are already logged into the domain. You can use Group Policy Objects (GPOs) to assign the URL to the Local Intranet site list.
 - b. Configure certificate trusts. This is an optional step that is required only if one or more of the servers that clients access do not have trusted certificates. The client computer may have to connect to the account federation servers, resource federation servers, or federation proxy servers, and the destination web servers. If any of these certificates are not from a trusted public CA, you may have to add the appropriate certificate or root certificate to the certificate store on the clients. You can do this by using GPOs.

Configuring a Resource Partner

The resource partner organization is the relying party in a B2B federation scenario. The resource partner organization is where the resources exist and are made accessible to account partner organizations. The resource partner handles the following tasks:


- Accepts security tokens that the account partner federation server produces, and validates them.
- Consumes the claims from the security tokens, and then provides new claims to its web servers after making an authorization decision.

A resource partner is a relying party in a B2B federation scenario

To configure an relying party:

1. Implement the physical topology
2. Add an attribute store
3. Configure a claims provider trust
4. Create claim rule sets for the claims provider trust

The web servers must have either Windows Identity Foundation or the AD FS 1.x Claims-Aware Web Agent role services installed to externalize the identity logic and accept claims.

 **Note:** Windows Identity Foundation provides a set of consistent development tools that enable developers to integrate claims-based authentication and authorization into their applications. Windows Identity Foundation also includes a Software Development Kit (SDK) and sample applications. You use a Windows Identity Foundation sample application in the lab for this module.

Configuring the resource partner organization is similar to configuring the account partner organization and consists of the following steps:

1. Implement the physical topology for the resource partner deployment. The planning and implementation steps are the same as the account partner, with the addition of planning the web server location and configuration.
2. Add an attribute store. On the resource partner, the attribute store is used to populate the claims that are offered to the client to present to the web server.
3. Connect to an account partner organization by creating a claims provider trust.
4. Create claim rule sets for the claims provider trust.

Configuring Claims Rules for B2B Scenarios

In a single organization deployment of AD FS, it may be quite easy to design and implement claims rules. In many cases, you may need to provide only the user name or group name that is collected from the claim and presented to the web server. In a B2B scenario, it is more likely that you will have to configure more complicated claims rules to define user access between widely varying systems.

Claim rules define how account partners (claims providers) create claims, and how resource partners (relying parties) consume claims. AD FS provides several templates that you can use when configuring claim rules:

- B2B scenarios may require more complex claims rules
- You can create claims rules using the following templates:
 - Send LDAP Attributes as Claims
 - Send Group Membership as a Claim
 - Pass Through or Filter an Incoming Claim
 - Transform an Incoming Claim
 - Permit or Deny Users Based on an Incoming Claim
- You can also create custom rules using the AD FS claim rule language

- **Send LDAP Attributes as Claims rule template.** Use this template when you select specific attributes in an LDAP attribute store to populate claims. You can configure multiple LDAP attributes as individual claims in a single claim rule that you create from this template. For example, you can create a rule that extracts the **sn** (surname) and **givenName** AD DS attributes from all authenticated users, and then send these values as outgoing claims to be sent to a relying party.
- **Send Group Membership as a Claim rule template.** Use this template to send a particular claim type and associated claim value that is based on the user's AD DS security group membership. For example, you might use this template to create a rule that sends a group claim type with a value of SalesAdmin, if the user is a member of the Sales Manager security group within their AD DS domain. This rule issues only a single claim, based on the AD DS group that you select as a part of the template.
- **Pass Through or Filter an Incoming Claim rule template.** Use this template to set additional restrictions on which claims are submitted to relying parties. For example, you might want to use a user email address as a claim, but only forward the email address if the domain suffix on the email address is adatum.com. When using this template, you can either pass through whatever claim you extract from the attribute store, or you can configure rules that filter whether the claim is passed on based on various criteria.
- **Transform an Incoming Claim rule template.** Use this template to map the value of an attribute in the claims provider attribute store to a different value in the relying party attribute store. For example, you may want to provide all members of the Marketing department at A. Datum Corporation limited access to a purchasing application at Trey Research. At Trey Research, the attribute used to define the limited access level may have an attribute of **LimitedPurchaser**. To address this scenario, you can

configure a claims rule that transforms an outgoing claim where the Department value is Marketing, to an incoming claim where the **ApplicationAccess** attribute is **LimitedPurchaser**. Rules created from this template must have a one-to-one relationship between the claim at the claims provider and the claim at the relying partner.

- Permit or Deny Users Based on an Incoming Claim rule template. This template is available only when you are configuring Issuance Authorization Rules or Delegation Authorization Rules on a relying party trust. Use this template to create rules that enable or deny access by users to a relying party, based on the type and value of an incoming claim. This claim rule template allows you to perform an authorization check on the claims provider before claims are sent to a relying party. For example, you can use this rule template to create a rule that only permits users from the Sales group to access a relying party, while authentication requests from members of other groups are not sent to the relying party.

If none of the built-in claim rule templates provide the functionality that you require, you can create more complex rules using the AD FS claim rule language. By creating a custom rule, you can extract claims information from multiple attribute stores and also combine claim types into a single claim rule.

How Home Realm Discovery Works

Some resource partner organizations that are hosting claims-aware applications may want to enable multiple account partners to access their applications. In this scenario, when users connect to the web application, there must be some mechanism for directing the users to the AD FS federation server in their home domain, rather than to another organization's federation server. The process for directing clients to the appropriate account partner is called *home realm discovery*.

Home realm discovery occurs after the client connects to the relying party's website and the client has been redirected to the relying party's federation server. At this point, the relying party's federation server must redirect the client to the federation server in the client's home realm so that the user can be authenticated. If there are multiple claims providers configured on the relying party federation server, it has to know to which federation server to redirect the client.

At a high level, there are three ways in which to implement home realm discovery:

1. Ask users to select their home realm. With this option, when the user is redirected to the relying party's federation server, the federation server can display a web page requesting that the user identify for which company they work. Once the user selects the appropriate company, the federation server can use that information to redirect the client computer to the appropriate home federation server for authentication.
2. Modify the link for the web application to include a WHR string that specifies the user's home realm. The relying party's federation server uses this string to redirect the user to the appropriate home realm automatically. This means that the user does not have to be prompted to select the home realm, because the WHR string in the URL that the user clicks relays the needed information to the relying party's federation server. The modified link might look something like <https://www.adatum.com/OrderApp/?whr=urn:federation:TreyResearch>.
3. If the remote application is SAML 2.0-compliant, users can use a SAML profile called IdPInitiated SSO. This SAML profile configures users to access their local claims provider first, which can prepare the

- Home realm discovery is required on resource partners when AD FS federations are configured with account partners
- To enable home realm discovery, you can:
 - Prompt the user for home realm information
 - Modify the URL for the web application to specify the home realm
 - Configure a SAML profile called IdPinitiated SSO, to direct users first to the account partner site

user's token with the claims required to access the partner's web application. This process changes the normal process for accessing the web application, by having the users log on to the claims provider federation server first, and then prompting them to select which application they want to access, so that their token can be created with the appropriate information.



Note: The home realm discovery process occurs the first time the user tries to access a web application. After the user authenticates successfully, a home realm discovery cookie is issued to the client so that the user does not have to go through the process the next time. This home realm discovery cookie expires after a month, unless the cookie cache is cleared prior to expiration.

Demonstration: Configuring Claims Rules

In this demonstration, you will see how to configure claims rules on a relying party trust that forwards a group name as part of the claim. You will also see how to configure a claims rule that limits access to the application only to members of a particular group.

Demonstration Steps

Configure Claims Rules

1. On LON-DC1, edit the **Adatum Test App** relying party trust by creating a new Issuance Transform Rule that passes through or filters an incoming claim. Name the rule **Send Group Name Rule**, and configure the rule to use an incoming claim type of group.
2. Delete the Issuance Authorization Rule that grants access to all users.
3. Create a new Issuance Authorization Rule that permits or denies user access based on the incoming claim. Configure the rule with the name **Permit Production Group Rule**, an **Incoming claim type** of **Group**, an **Incoming claim value** of **Production**, and select the option to **Permit access to users with this incoming claim**
4. Create a new Issuance Authorization Rule that permits or denies user access based on the incoming claim. Configure the rule with the name **Allow A Datum Users**, an **Incoming claim type** of **UPN**, an **Incoming claim value**, of **@adatum.com**, and select the option to **Permit access to users with this incoming claim**, and then click **Finish**.
5. Open the **Allow A Datum Users** rule properties, and show the claims rule language to the students.

Lab: Implementing AD FS

Scenario

A. Datum Corporation has set up a variety of business relationships with other companies and customers. Some of these partner companies and customers must access business applications that are running on the A. Datum network. The business groups at A. Datum want to provide a maximum level of functionality and access to these companies. The Security and Operations departments want to ensure that the partners and customers can only access the resources to which they require access, and that implementing the solution does not significantly increase the workload for the Operations team. A. Datum is also working on migrating some parts of their network infrastructure to online services, including Windows Azure and Office 365.

To meet these business requirements, A. Datum plans to implement AD FS. In the initial deployment, the company plans to use AD FS to implement SSO for internal users who access an application on a web server. A. Datum also has entered into a partnership with another company, Trey Research. Trey Research users must be able to access the same application.

As one of the senior network administrators at A. Datum, it is your responsibility to implement the AD FS solution. As a proof of concept, you plan to deploy a sample claims-aware application, and configure AD FS to enable both internal users and Trey Research users to access the same application.

Objectives

- Configure the AD FS prerequisites.
- Install and configure AD FS.
- Configure AD FS for single organization.
- Configure and validate SSO for a business federation scenario.

Lab Setup

Estimated Time: 90 minutes

- 20412A-LON-DC1
- 20412A-LON-SVR1
- 20412A-LON-CL1
- 20412A-MUN-DC1

User name: **Adatum\Administrator**

Password: **Pa\$\$w0rd**

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

1. On the host computer, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In Hyper-V Manager, click **20412A-LON-DC1**, and in the Actions pane, click **Start**.
3. In the Actions pane, click **Connect**. Wait until the virtual machine starts.
4. Log on using the following credentials:
 - User name: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
5. Repeat steps 2 to 3 for **20412A-LON-SVR1**, **20412A-LON-CL1**, and **20412A-MUN-DC1**.
 - a. Log on to **20412A-LON-SVR1** as **Adatum\Administrator**.

- b. Do not log on to 20412A-LON-CL1 at this point.
- c. On 20412A-MUN-DC1, log in as **TreyResearch\Administrator** with the password **Pa\$\$w0rd**.

Exercise 1: Configuring AD FS Prerequisites

Scenario

To deploy AD FS at A. Datum Corporation, you must verify that all required components are configured. You plan to verify that AD CS is deployed in the organization, and then configure the certificates required for AD FS on the AD FS server and on the web servers. You also plan to configure the DNS forwarders to enable communication between Adatum.com and TreyResearch.net.

The main tasks for this exercise are as follows:

1. Configure DNS forwarders.
2. Exchange root certificates to enable certificate trusts.
3. Request and install a certificate for the web server.
4. Bind the certificate to the claims-aware application on the web server, and verify application access.

► Task 1: Configure DNS forwarders

1. On LON-DC1, create a new conditional forwarder for the TreyResearch.net domain, using the DNS server IP address of **172.16.10.10**.
2. On MUN-DC1, create a new conditional forwarder for the Adatum.com domain, using the DNS server IP address of **172.16.0.10**.

► Task 2: Exchange root certificates to enable certificate trusts

1. On LON-DC1, copy **MUN-DC1.TreyResearch.net_TreyResearchCA.crt** from **\\MUN-DC1.treyresearch.net\certenroll** to the **Documents** folder.
2. Create a new MMC, and add the **Group Policy Management Editor**.
3. Edit the Default Domain Policy Group Policy Object and import the copied root certificate to the Trusted Root Certification Authorities folder.
4. On MUN-DC1, copy the **LON-DC1.Adatum.com_Adatum-LON-DC1-CA.crt** from **\\LON-DC1.Adatum.com\certenroll** to the **Documents** folder.
5. Create a new MMC, and add the **Certificates** snap-in focused on the Local Computer.
6. Import the copied root certificate to the Trusted Root Certification Authorities folder.

► Task 3: Request and install a certificate for the web server

1. On LON-SVR1, open the **Internet Information Services (IIS) Manager**.
2. Request a new domain certificate for the server using the following parameters:
 - Common name: **LON-SVR1.adatum.com**
 - Organization: **A. Datum**
 - Organization unit: **IT**
 - City/locality: **London**
 - State/province: **England**
 - Country/region: **GB**
3. Request the certificate from **Adatum-LON-DC1-CA**.

► **Task 4: Bind the certificate to the claims-aware application on the web server, and verify application access**

1. On LON-SVR1, in IIS, create a new HTTPS site binding, and then select the newly created certificate.
2. On LON-DC1, open Internet Explorer and connect to **https://lon-svr1.adatum.com/adatumtestapp**.
3. Verify that you can connect to the site, but that you receive a 401 access denied error. This is expected, because you have not yet configured AD FS for authentication.
4. Close Internet Explorer.

Results: In this exercise, you configured DNS forwarding to enable name resolution between A. Datum and Trey Research, and you exchanged root certificates between the two organizations. You also installed and configured a web certificate on the application server.

Exercise 2: Installing and Configuring AD FS

Scenario

To start the AD FS implementation, you plan to install AD FS on the A. Datum Corporation's domain controller, and configure the server as a standalone federation server. You also plan to configure the server to use a CA-signed token signing certificate.

The main tasks for this exercise are as follows:

1. Install and configure AD FS.
2. Create a standalone federation server using the AD FS Federation Server Configuration Wizard.
3. Verify that FederationMetaData.xml is present and contains valid data.

► **Task 1: Install and configure AD FS**

- On LON-DC1, in Server Manager, add the Active Directory Federation Services server role.

► **Task 2: Create a standalone federation server using the AD FS Federation Server Configuration Wizard**

- On LON-DC1, run the AD FS Federation Server Configuration Wizard using the following parameters:
 - Create a new federation service.
 - Create a standalone deployment.
 - Use the LON-DC1.Adatum.com certificate.
 - Choose a service name of LON-DC1.Adatum.com

► **Task 3: Verify that FederationMetaData.xml is present and contains valid data**

1. On LON-CL1, log on as **Adatum\Brad**, using the password **Pa\$\$w0rd**.
2. Open Internet Explorer.
3. Open Internet Options, and add **https://LON-DC1.Adatum.com**, and **https://LON-SVR1.adatum.com** to the Local intranet zone.
4. Connect to **https://lon-dc1.adatum.com/federationmetadata/2007-06/federationmetadata.xml**.
5. Verify that the xml file opens successfully, and then scroll through its contents.

6. Close Internet Explorer.

Results: In this exercise, you installed and configured the AD FS server role, and verified a successful installation by viewing the Federation Meta Data .xml contents.

Exercise 3: Configuring AD FS for a Single Organization

Scenario

The first scenario for implementing the proof of concept AD FS application is to ensure that internal users can use SSO to access the web application. You plan to configure the AD FS server and a web application to enable this scenario. You also want to verify that internal users can access the application.

The main tasks for this exercise are as follows:

1. Configure a Token-signing certificate for LON-DC1.Adatum.com.
2. Configure the Active Directory claims provider trust.
3. Configure the claims application to trust incoming claims by running the Windows Identity Foundation Federation Utility.
4. Configure a relying party trust for the claims-aware application.
5. Configure claim rules for the relying party trust.
6. Test access to the claims-aware application.

► Task 1: Configure a Token-signing certificate for LON-DC1.Adatum.com

1. On LON-DC1, use the **set-ADFSProperties -AutoCertificateRollover \$False** command to enable modification of the assigned certificates.
2. In the AD FS Management console, add the LON-DC1.Adatum.com certificate as a new token-signing certificate. Verify that the certificate has a subject of **CN=LON-DC1.Adatum.com**, and purposes of **Proves your identity to a remote computer** and **Ensures the identity of a remote computer**.
3. Make the new certificate the primary certificate, and remove the old certificate.

► Task 2: Configure the Active Directory claims provider trust

1. On LON-DC1, in the AD FS Management console, go to the **Claims Provider Trusts**, highlight the **Active Directory** store and then go to **Edit Claim Rules**.
2. In the **Edit Claim Rules for Active Directory** dialog box, on the **Acceptance Transform Rules** tab, launch the Add Transform Claim Rule Wizard and complete the wizard with the following settings:
 - Select **Send LDAP Attributes** as **Claims under Claim rule template**.
 - Name the claim rule **Outbound LDAP Attribute Rule**.
 - Choose **Active Directory** as the **Attribute Store**.
3. In the Mapping of LDAP attributes to outgoing claim types select the following values:
 - E-Mail-Addresses to **E-Mail Address**
 - User-Principal-Name to **UPN**
 - Display-Name to **Name**

► **Task 3: Configure the claims application to trust incoming claims by running the Windows Identity Foundation Federation Utility**

1. On LON-SVR1, from the Start screen, launch the Windows Identity Foundation Federation Utility.
2. Complete the wizard with the following settings:
 - Point to the web.config file of the Windows Identity Foundation sample application by pointing to **C:\inetpub\wwwroot\AdatumTestApp\web.config**.
 - Specify an Application URI box by typing **https://lon-svr1.adatum.com/AdatumTestApp/**.
 - Select to Use an existing STS, and enter a path **https://lon-dc1.adatum.com/federationmetadata/2007-06/federationmetadata.xml**.
 - Select **No encryption**.

► **Task 4: Configure a relying party trust for the claims-aware application**

1. In the AD FS Management console, in the middle pane, click **Required: Add a trusted relying party**.
2. Complete the Add Relying Party Wizard with the following settings:
 - Choose to **Import data about the relying party published online or on a local network**, and then type **https://lon-svr1.adatum.com/adatumtestapp**.
 - Specify a Display name of **ADatum Test App**.
 - Choose to **Permit all users to access this relying party**.
 - When the wizard completes, accept the option to open the **Edit Claims Rules for ADatum Test App**.

► **Task 5: Configure claim rules for the relying party trust**

1. In the **Edit Claim Rules for Adatum Test App** properties dialog box, choose to add a rule on the **Issuance Transform Rules** tab.
2. Complete the Add Transform Claim Rule Wizard with the following settings:
 - In the **Claim rule template** drop-down list, click **Pass through or Filter an Incoming Claim**.
 - Name the claim rule **Pass through Windows Account name rule**.
 - In the **Incoming claim type** drop-down list, click **Windows account name**.
3. Create three more rules to pass through **E-Mail Address**, **UPN**, and **Name type claim**.

► **Task 6: Test access to the claims-aware application**

1. On LON-CL1, open Internet Explorer, and connect to **https://lon-svr1.adatum.com/AdatumTestApp/**
2. Verify that you can access the application.

Results: In this exercise, you configured a Token signing certificate and configured a claims provider trust for Adatum.com. You also should have configured the sample application to trust incoming claims, and configured a relying party trust and associated claim rules. You also tested access to the sample Windows Identity Foundation application in a single organization scenario.

Exercise 4: Configuring AD FS for Federated Business Partners

Scenario

The second deployment scenario is to enable TreyResearch users to access the web application. You plan to configure the integration of AD FS at TreyResearch with AD FS at A. Datum Corporation, and then verify that TreyResearch users can access the application. You also want to confirm that you can configure access that is based on user groups. You must ensure that all users at A. Datum, and only users who are in the Production group at TreyResearch, can access the application.


The main tasks for this exercise are as follows:

1. Add a claims provider trust for the TreyResearch.net AD FS server.
2. Configure a relying party trust on MUN-DC1 for the A. Datum claims-aware application.
3. Verify access to the A. Datum test application for Trey Research users.
4. Configure claim rules for the claim provider trust and the relying party trust to allow access only for a specific group.
5. Verify restrictions and accessibility to the claims-aware application.

► Task 1: Add a claims provider trust for the TreyResearch.net AD FS server

1. On LON-DC1, in the AD FS Management console, go to **Trust Relationships**, go to **Claims Provider Trusts**, and then choose to **Add Claims Provider Trust**.
2. Complete the Add Claims Provider Trust Wizard with the following settings:
 - Choose **Import data about the claims provider published online or on a local network**, and enter **https://mun-dc1.treyresearch.net** as the data source.
 - In **Display Name**, type **mun-dc1.treyresearch.net**.
 - Complete the wizard.
3. In the **Edit Claim Rules for the mun-dc1.treyresearch.net** properties dialog box, use the following values:
 - **Add a Rule** to the Acceptance Transform Rules.
 - Choose **Pass Through or Filter an Incoming Claim** in the **Claim rule template** list.
 - Use **Pass through Windows account name rule** as the claim rule name.
 - Choose **Windows account name** as the incoming claim type, and then choose to **Pass through all claim values**.
 - Complete the rule.
4. On LON-DC1, run the following command in Windows PowerShell.

```
Set-ADFSClaimsProviderTrust -TargetName "mun-dc1.treyresearch.net" -
SigningCertificateRevocationCheck None
```

 **Note:** You should disable certificate revocation checking only in test environments. In a production environment, certificate revocation checking should be enabled.


► Task 2: Configure a relying party trust on MUN-DC1 for the A. Datum claims-aware application

1. On MUN-DC1, in the AD FS Management console, open the Add Relying Party Trust Wizard, and complete it with the following settings:

- Choose to **Import data about the relying party published online or on a local network** and type in **https://lon-dc1.adatum.com**.
 - Specify a Display name of Adatum TestApp.
 - Choose to **Permit all users to access this relying party**.
 - Accept the option to open the Edit Claim Rules for Adatum TestApp when the wizard completes.
2. In the **Edit Claim Rules for Adatum TestApp** properties dialog box, on the **Issuance Transform Rules** tab, click to add a rule with the following settings:
 - In the claim rule template list, choose **Pass Through or Filter an Incoming claim**.
 - In the **Claim rule name** box, type **Pass through Windows account name rule**.
 - Choose **Windows account name** in **Incoming claim type**.
 - Choose to **Pass through all claim values**.
 - Complete the wizard.

► **Task 3: Verify access to the A. Datum test application for Trey Research users**

1. On MUN-DC1, open Internet Explorer and connect to **https://lon-svr1.adatum.com/adatumtestapp/**
2. Select **mun-dc1.treyresearch.net** as the home realm, and log on as **TreyResearch\April**, with the password **Pa\$\$w0rd**.
3. Verify that you can access the application.
4. Close Internet Explorer, and connect to the same web site. Verify that this time you are not prompted for a home realm.

 **Note:** You are not prompted for a home realm again. Once users have selected a home realm and been authenticated by a realm authority, they are issued an `_LSRealm` cookie by the relying party federation server. The default lifetime for the cookie is 30 days. Therefore, to log on multiple times, you should delete that cookie after each logon attempt to return to a clean state.

► **Task 4: Configure claim rules for the claim provider trust and the relying party trust to allow access only for a specific group**

1. On MUN-DC1, open the AD FS Management Console, access the Adatum TestApp relying party trust.
2. Add a new Issuance Transform Rule that sends the group membership as a claim. Name the rule **Permit Production Group Rule**, configure the User's Group as **Production**, configure the Outgoing claim type as **Group**, and the Outgoing claim value as **Production**.
3. On LON-DC1, in the AD FS Management Console, edit the `mun-dc1.treyresearch.net` Claims Provider Rule to create a new rule that passes through or filters an incoming claim with the rule name of **Send Production Group Rule**. Configure the rule with an incoming claim type of **Group**.
4. Edit the Adatum Test App relying party trust by creating a new Issuance Transform Rule that passes through or filters an incoming claim. Name the rule **Send TreyResearch Group Name Rule**, and configure the rule to use an incoming claim type of **Group**.
5. Delete the Issuance Authorization Rule that grants access to all users.
6. Create a new Issuance Authorization Rule that permits or denies user access based on the incoming claim. Configure the rule with the name **Permit TreyResearch Production Group Rule**, an **Incoming claim type** of **Group**, an **Incoming claim value** of **Production**, and select the option to **Permit access to users with this incoming claim**.

7. Create a new Issuance Authorization Rule that permits or denies user access based on the incoming claim. Configure the rule with the name **Temp**, an **Incoming claim type** of **UPN**, an **Incoming claim value**, of **@adatum.com**, and select the option to **Permit access to users with this incoming claim**, and then click **Finish**.
8. Edit the Temp rule and copy the claim rule language into the clipboard.
9. Delete the Temp rule.
10. Create a new rule that sends claims using a custom rule named **ADatum User Access Rule**.
11. Click in the **Custom rule** box, and then press Ctrl + V to paste the clipboard contents into the box. Edit the first URL to match the following text, and then click **Finish**.

```
c:[Type == "http://schemas.xmlsoap.org/ws/2005/05/identity/claims/upn", Value =~
"^(?i).+@adatum\.com$"]=> issue(Type =
"http://schemas.microsoft.com/authorization/claims/permit", Value = "PermitUsersWithClaim");
```

► Task 5: Verify restrictions and accessibility to the claims-aware application

1. On MUN-DC1, open Internet Explorer, and connect to **https://lon-svr1.adatum.com/adatumtestapp/**
2. Verify that **TreyResearch\April** no longer has access to the A. Datum test app.
3. Clear the browsing history in Internet Explorer.
4. Connect to **https://lon-svr1.adatum.com/adatumtestapp/**.
5. Verify that **TreyResearch\Morgan** does have access to the A. Datum test app. Morgan is a member of the Production group.

Results: In this exercise, you configured a claims provider trust for TreyResearch on Adatum.com. and a relying party trust for Adatum on TreyResearch. You verified access to the A. Datum claim-aware application. Then you configured the application to restrict access from TreyResearch to specific groups, and you verified appropriate access.

► To shut down the virtual machines

When you finish the lab, revert the virtual machines to their initial state.

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

Lab Review

Question: In this lab, you implemented access to a claims-aware application for both internal and external users. What extra steps did you have to take in the relying party to enable access for external users?

Question: How can you identify which claims are used to provide user access to the sample Windows Identity Foundation application that you used in the lab?

Module Review and Takeaways

Question: What are the benefits of deploying AD FS with a cloud-based application or service?

Question: Under what circumstances would you choose to deploy a federation proxy server? Under what circumstances, do you not need to deploy a federation proxy server?

Common Issues and Troubleshooting Tips

Common Issue	Troubleshooting Tip
Certificate errors on the federation server	
Certificate errors on the client	
Client application failed to authenticate with AD FS	

Real-world Issues and Scenarios

1. **Question:** Tailspin Toys is deploying a new claims-based web application. The web application needs to be accessible to both Tailspin Toys users and to TreyResearch users. What AD FS components will you need to deploy at Tailspin Toys to enable this level of access?

Answer:

2. **Question:** Fabrikam, Inc. is examining the requirements for AD FS. The company wants to use a federation proxy server for maximum security. Fabrikam, Inc. currently has an internal network with internal DNS servers, and their internet-facing DNS is hosted by a hosting company. The perimeter network uses the hosting company's DNS servers for DNS resolution. What must the company do to prepare for the deployment?

Answer:

Course Evaluation

Course Evaluation



Your evaluation of this course will help Microsoft understand the quality of your learning experience.

Please work with your training provider to access the course evaluation form.

Microsoft will keep your answers to this survey private and confidential and will use your responses to improve your future learning experience. Your open and honest feedback is valuable and appreciated.

MCT USE ONLY. STUDENT USE PROHIBITED

Module 1: Implementing Advanced Network Services

Lab: Implementing Advanced Network Services

Exercise 1: Configuring Advanced DHCP Settings

► Task 1: Configure a superscope

1. On LON-DC1, in Server Manager, click **Tools**, and then click **DHCP**.
2. In the DHCP console, click **LON-DC1.adatum.com**, select and then right-click **IPv4**, and then click **New Scope**.
3. In the New Scope Wizard, click **Next**.
4. On the **Scope Name** page, in the **Name** box, type **Scope1**, and then click **Next**.
5. On the **IP Address Range** page, in the **Start IP address** box, type **192.168.0.50**, and then in the **End IP address** box, type **192.168.0.100**.
6. In the **Subnet mask** box, ensure that **255.255.255.0** is entered, and then click **Next**.
7. On the **Add Exclusions and Delay** page, click **Next**.
8. On the **Lease Duration** page, click **Next**.
9. On the **Configure DHCP Options** page, select **Yes, I want to configure these options now**, and then click **Next**.
10. On the **Router (Default Gateway)** page, in the **IP address** box, type **192.168.0.1**, click **Add**, and then click **Next**.
11. On the **Domain Name and DNS Servers** page, ensure the parent domain is **Adatum.com**, and then click **Next**.
12. On the **WINS Servers** page, click **Next**.
13. On the **Activate Scope** page, click **No, I will activate this scope later**, and then click **Next**.
14. On the **Completing the New Scope Wizard** page, click **Finish**.
15. Right-click **IPv4**, and then click **New Scope**.
16. In the New Scope Wizard, click **Next**.
17. On the **Scope Name** page, in the **Name** box, type **Scope2**, and then click **Next**.
18. On the **IP Address Range** page, in the **Start IP address** box, type **192.168.1.50**, and then in the **End IP address** box, type **192.168.1.100**.
19. In the **Subnet mask** box, ensure that **255.255.255.0** is entered, and then click **Next**.
20. On the **Add Exclusions and Delay** page, click **Next**.
21. On the **Lease Duration** page, click **Next**.
22. On the **Configure DHCP Options** page, select **Yes, I want to configure these options now**, and then click **Next**.
23. On the **Router (Default Gateway)** page, in the **IP address** box, type **192.168.1.1**, click **Add**, and then click **Next**.

24. On the **Domain Name and DNS servers** page, ensure the parent domain is **Adatum.com**, and then click **Next**.
25. On the **WINS Servers** page, click **Next**.
26. On the **Activate Scope** page, click **No, I will activate this scope later**, and then click **Next**.
27. On the **Completing the New Scope Wizard** page, click **Finish**.
28. Right-click the **IPv4** node, and then click **New Superscope**.
29. In the New Superscope Wizard, click **Next**.
30. On the **Superscope Name** page, in the **Name** box, type **AdatumSuper**, and then click **Next**.
31. On the **Select Scopes** page, select **Scope1**, hold down the Ctrl key, select **Scope2**, and then click **Next**.
32. On the **Completing the New Superscope Wizard** page, click **Finish**.

► **Task 2: Configure DHCP name protection**

1. On LON-DC1, in the DHCP console, expand **Lon-DC1.adatum.com**.
2. Right-click **IPv4**, and then click **Properties**.
3. Click the **DNS** tab.
4. In the Name Protection pane, click **Configure**.
5. Select the **Enable Name Protection** check box, and then click **OK**.
6. Click **OK** again.

► **Task 3: Configure and verify DHCP failover**

1. On LON-SVR1, in Server Manager, click **Tools**, and then from the drop-down list, click **DHCP**. Note that the server is authorized, but that no scopes are configured.
2. On LON-DC1, in the DHCP console, right-click the **IPv4** node, and then click **Configure Failover**.
3. In the Configure Failover Wizard, click **Next**.
4. On the **Specify a partner server to use for failover** page, in the **Partner Server** box, enter **172.16.0.21**, and then click **Next**.
5. On the **Create a new failover relationship** page, in the **Relationship Name** box, enter **Adatum**.
6. In the **Maximum Client Lead Time** field, set the hours to **0**, and set the minutes to **15**.
7. Ensure that the **Mode** field is set to **Load balance**, and that the **Load Balance Percentage** is set to **50%**.
8. Select the **State Switchover Interval** check box. Keep the default value of 60 minutes.
9. In the **Enable Message Authentication Shared Secret** box, type **Pa\$\$w0rd**, and then click **Next**.
10. Click **Finish**, and then click **Close**.
11. On LON-SVR1, refresh the IPv4 node, and then note that the IPv4 node is active.
12. Expand the **IPv4** node, expand **Scope Adatum**, click the **Address Pool** node, and note that the address pool is configured.
13. Click the **Scope Options** node, and note that the scope options are configured.
14. Start **20412A-LON-CL1** and log on as **Adatum\Administrator** with a password of **Pa\$\$w0rd**.
15. On the Start screen, type **Control Panel**.

16. In the **Apps Results** box, click **Control Panel**.
17. In Control Panel, click **Network and Internet**, click **Network and Sharing Center**, click **Change adapter settings**, and then right-click **Local Area Connection**, and then click **Properties**.
18. In the **Local Area Connection Properties** dialog box, click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
19. In the Properties window, select the **Obtain an IP address automatically** radio button, click **Obtain DNS server address automatically**, and then click **OK**.
20. In the **Local Area Connection Properties** dialog box, click **Close**.
21. Hover over the bottom right corner to expose the fly-out menu, and then click **Search** Charm.
22. In the **Apps** search box, type **Cmd**, and then press Enter.
23. In the command prompt window, type **ipconfig**, and then press Enter. Record your IP address.
24. On **LON-DC1**, on the taskbar, click the **Server Manager** icon.
25. In Server Manager, click **Tools**, and then click **Services**.
26. In the Services window, locate the **DHCP Server** service, and then click **Stop the service**.
27. Close the Services window, and close the DHCP console.
28. On LON-CL1, in the command prompt window, type **ipconfig /release**, and then press Enter.
29. Type **ipconfig /renew**, and then press Enter.
30. Type **ipconfig**, and then press Enter. What is your IP address? Answers may vary.
31. Shut down the LON-SVR1 server.
32. On LON-DC1, in the Services console, start the DHCP server service.
33. Close the Services console.

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

Exercise 2: Configuring Advanced DNS Settings

► Task 1: Configure DNSSEC

1. On LON-DC1, in Server Manager, click **Tools**, and then in the drop-down list, click **DNS**.
2. Expand **LON-DC1**, expand **Forward Lookup Zones**, click **Adatum.com**, and then right-click **Adatum.com**.
3. On the menu, click **DNSSEC>Sign the Zone**.
4. In the Zone Signing Wizard, click **Next**.
5. On the **Signing options** page, click **Customize zone signing parameters**, and then click **Next**.
6. On the **Key Master** page, ensure that LON-DC1 is the Key Master, and then click **Next**.
7. On the **Key Signing Key (KSK)** page, click **Next**.
8. On the **Key Signing Key (KSK)** page, click **Add**.
9. On the **New Key Signing Key (KSK)** page, click **OK**.
10. On the **Key Signing Key (KSK)** page, click **Next**.

11. On the **Zone Signing Key (ZSK)** page, click **Next**.
12. On the **Zone Signing Key (ZSK)** page, click **Add**.
13. On the **New Zone Signing Key (ZSK)** page, click **OK**.
14. On the **Zone Signing Key (ZSK)** page, click **Next**.
15. On the **Next Secure (NSEC)** page, click **Next**.
16. On the **Trust Anchors** page, select the **Enable the distribution of trust anchors for this zone** check box. Click **Next**.
17. On the **Signing and Polling Parameters** page, click **Next**.
18. On the **DNS Security Extensions** page, click **Next**, and then click **Finish**.
19. In the DNS console, expand **Trust Points**, expand **com**, and then click **Adatum**. Ensure that the DNSKEY resource records display, and that their status is valid.
20. Minimize the DNS Manager.
21. In Server Manager, click **Tools**, and then on the drop-down list, click **Group Policy Management**.
22. Expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**, right-click **Default Domain Policy**, and then click **Edit**.
23. In the Group Policy Management Editor, under **Computer Configuration**, expand **Policies**, expand **Windows Settings**, and then click **Name Resolution Policy**.
24. In the right pane, under **Create Rules**, in the **Suffix** box, type **Adatum.com** to apply the rule to the suffix of the namespace.
25. Select the **Enable DNSSEC in this rule** check box, select the **Require DNS clients to check that the name and address data has been validated by the DNS server** check box, click **Create**.
26. Close the Group Policy Management Editor and Group Policy Management Console.

► **Task 2: Configure the DNS socket pool**

1. On LON-DC1, hover over the bottom right corner to expose the fly-out menu, and then click **Search**.
2. In the **Apps** search box, type **Cmd**, and then press Enter.
3. In the command prompt window, type the following command, and then press Enter to view the current size of the DNS socket pool. Note that the current size is 2,500.

```
dnscmd /info /socketpoolsize
```

4. Type the following command, and then press Enter to change the socket pool size to 3,000.

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

5. Type the following command, and then press Enter to stop the DNS server.

```
net stop dns
```

6. Type the following command, and then press Enter to restart the DNS server.

```
net start dns
```

7. Type the following command, and then press Enter to confirm the new socket pool size.

```
dnscmd /info /socketpoolsize
```

► Task 3: Configure DNS cache locking

1. In the command prompt window, type the following command, and then press Enter to display the current percentage value of the DNS cache lock.

```
dnscmd /info /CacheLockingPercent
```

Note that the current value is 100 percent.

2. Type the following command, and then press Enter to change the cache lock value to 75 percent.

```
dnscmd /config /CacheLockingPercent 75
```

3. Type the following command, and then press Enter to stop the DNS server.

```
net stop dns
```

4. Type the following command, and then press Enter to restart the DNS server.

```
net start dns
```

5. Type the following command, and then press Enter to display the current percentage value of the DNS cache lock.

```
dnscmd /info /CacheLockingPercent
```

Note the new value is 75 percent.

6. Leave the command prompt window open for the next task.

► Task 4: Configure a GlobalName Zone

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

```
Dnscmd LON-DC1 /ZoneAdd Contoso.com /DsPrimary /DP /forest
```

2. In the command prompt window, type the following command, and then press Enter to enable support for GlobalName zones:

```
dnscmd lon-dc1 /config /enablglobalnamesupport 1
```

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

```
Dnscmd LON-DC1 /ZoneAdd GlobalNames /DsPrimary /DP /forest
```

4. Minimize the command prompt window.
5. Restore the DNS console from the taskbar.
6. In the DNS console, click **Action**, and then click **Refresh** to refresh the view.
7. In the DNS console, expand **Forward Lookup Zones**, click the **Contoso.com** zone, right-click **Contoso.com**, and then click **New Host (A or AAAA)**.
8. In the **New Host** dialog box, in the **Name** box, type **App1**.



Note: The **Name** box uses the parent domain name if it is left blank.

9. In the **IP address** box, type **192.168.1.200**, and then click **Add Host**.
10. Click **OK**, and then click **Done**.
11. Right-click the **GlobalNames** zone, and then click **New Alias (CNAME)**.
12. In the **New Resource Record** dialog box, in the **Alias name** box, type **App1**.
13. In the **Fully qualified domain name (FQDN) for target host** box, type **App1.Contoso.com**, and then click.
14. Close DNS Manager and close the command prompt.

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

Exercise 3: Configuring IP Address Management

► Task 1: Install the IPAM feature

1. On LON-SVR2, in the Server Manager Dashboard, click **Add roles and features**.
2. In the Add Roles and Features Wizard, click **Next**.
3. On the **Select installation type** page, click **Next**.
4. On the **Select destination server** page, click **Next**.
5. On the **Select server roles** page, click **Next**.
6. On the **Select features** page, select the **IP Address Management (IPAM) Server** check box.
7. In the **Add features that are required for IP Address Management (IPAM) Server** popup, click **Add Features**, and then click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. Close the Add Roles and Features Wizard when complete.

► Task 2: Configure IPAM–related GPOs

1. In the Server Manager navigation pane, click **IPAM**.
2. In the IPAM Overview pane, click **Connect to IPAM server**, and then select **LON-SVR2.Adatum.com** and then click **OK**.
3. Click **Provision the IPAM server**.
4. In the Provision IPAM Wizard, click **Next**.
5. On the **Select provisioning method** page, ensure that the **Group Policy Based** method is selected, in the **GPO name prefix** box, type **IPAM**, and then click **Next**.
6. On the **Confirm the Settings** page, click **Apply**. Provisioning will take a few moments to complete.
7. When provisioning completes, click **Close**.

► Task 3: Configure IP management server discovery

1. On the IPAM Overview pane, click **Configure server discovery**.
2. In the **Configure Server Discovery** settings dialog box, click **Add**, and then click **OK**.
3. In the IPAM Overview pane, click **Start server discovery**. Discovery may take 5-10 minutes to run. The yellow bar will indicate when discovery is complete.

► Task 4: Configure managed servers

1. In the IPAM Overview pane, click **Select or add servers to manage and verify IPAM access**. Notice that the IPAM Access Status is blocked.
2. Scroll down to the Details view, and note the status report, which is that the IPAM server has not yet been granted permission to manage LON-DC1 via Group Policy.
3. On the taskbar, right-click the **Windows PowerShell** icon, right-click **Windows PowerShell** and then click **Run as Administrator**.
4. At the command prompt, type the following command, and then click Enter:

```
Invoke-IPAMGpoProvisioning -Domain Adatum.com -GpoPrefixName IPAM -IPAMServerFqdn LON-SVR2.adatum.com -DelegatedGpoUser Administrator
```

5. When you are prompted to confirm the action, type **Y**, and then press Enter. The command will take a few moments to complete.
6. Close Windows PowerShell.
7. In Server Manager, in the details pane, right-click **LON-DC1**, and then click **Edit Server**.
8. In the **Add or Edit Server** dialog box, set the **Manageability status** to **Managed**, and then click **OK**.
9. Switch to LON-DC1, open the PowerShell prompt from the Taskbar.
10. Type **Gpupdate /force**, and then press Enter.
11. Close the command prompt window.
12. Switch to LON-SVR2, in Server Manager, in the IPAM console, right-click the **LON-DC1** entry, and then click **Refresh Server Access Status**. After the refresh completes, click the IPv4 console refresh button. It may take up to 10 minutes for the status to change. If necessary, repeat both refresh tasks as needed until a green check mark displays next to LON-DC1 and the IPAM Access Status shows Unblocked.
13. In the IPAM Overview pane, click **Retrieve data from managed servers**. This action will take a few moments to complete.

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

1. On LON-SVR2, in the IPAM navigation pane, under **MONITOR AND MANAGE**, click **DNS and DHCP Servers**.
2. In the details pane, right-click the instance of **LON-DC1.Adatum.com** that holds the DHCP server role, and then click **Create DHCP Scope**.
3. In the **Create DHCP Scope** dialog box, in the **Scope Name** box, type **TestScope**.
4. In the **Start IP address** box, type **10.0.0.50**.
5. In the **End IP address** box, type **10.0.0.100**.
6. Ensure the subnet mask is 255.0.0.0
7. In the Create scope pane, click **Options**.
8. In the Configure options pane, click the **Option** drop-down arrow, and then select **003 Router**.
9. Under **Values**, in the **IP Address** box, type **10.0.0.1**, click **Add to list**, and then click **OK**.
10. On LON-DC1, in the Server Manager toolbar, click **Tools**, and then click **DHCP**.
11. In the DHCP console, expand **LON-DC1**, expand **IPv4**, and confirm that the **TestScope** exists.

12. Minimize the DHCP console.

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

1. On LON-SVR2, in Server Manager, in the IPAM console tree, click **IP Address Blocks**.
2. In the right pane, click the **Tasks** drop-down arrow, and then click **Add IP Address Block**.
3. In the **Add or Edit IPv4 Address Block** dialog box, provide the following values, and then click **OK**:
 - Network ID: **172.16.0.0**
 - Prefix length: **16**
 - Description: **Head Office**
4. In the IPAM console tree, click **IP Address Inventory**.
5. In the right pane, click the **Tasks** drop-down arrow, and then click **Add IP Address**.
6. In the **Add IP Address** dialog box, under **Basic Configurations**, provide the following values, and then click **OK**:
 - IP address: **172.16.0.1**
 - MAC address: **112233445566**
 - Device type: **Routers**
 - Description: **Head Office Router**
7. Click the **Tasks** drop-down arrow, and then click **Add IP Address**.
8. In the **Add IP Address** dialog box, under **Basic Configuration**, provide the following values:
 - IP address: **172.16.0.10**
 - MAC address: **223344556677**
 - Device type: **Host**
9. In the Add IPv4 Address pane, click **DHCP Reservation**, and then enter the following values:
 - Reservation server name: **LON-DC1.Adatum.com**
 - Reservation name: **Webserver**
 - Reservation type: **Both**
10. In the Add IPv4 Address pane, click **DNS Record**, enter the following values, and then click **OK**:
 - Device name: **Webserver**
 - Forward lookup zone: **Adatum.com**
 - Forward lookup primary server: **LON-DC1.adatum.com**
11. When the entry displays in the IPv4 details pane, right-click the entry, and then click **Create DHCP Reservation**.
12. Right-click the entry again, and then click **Create DNS Host Record**.
13. On LON-DC1, open the DHCP console, expand **IPv4**, expand **Scope (172.16.0.0) Adatum**, and then click **Reservations**. Ensure that the 172.16.0.10 Webserver reservation displays.
14. Open the DNS console, expand **Forward Lookup Zones**, and then click **Adatum.com**. Ensure that a host record displays for Webserver.

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

► **To prepare for the next module**

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

MCT USE ONLY. STUDENT USE PROHIBITED

Module 2: Implementing Advanced File Services

Lab A: Implementing Advanced File Services

Exercise 1: Configuring iSCSI Storage

► Task 1: Install the iSCSI target feature

1. Log on to **LON-DC1** with username of **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 (Installed)**, expand **File and iSCSI Services**, select the **iSCSI Target Server** check box, and then click **Next**.
7. On the **Select features** page, click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. When the installation completes, click **Close**.

► Task 2: Configure the iSCSI targets

1. On LON-DC1, in Server Manager, in the navigation pane, click **File and Storage Services**.
2. In the File and Storage Services pane, click **iSCSI**.
3. In the iSCSI VIRTUAL DISKS pane, click **TASKS**, and then in the **TASKS** drop-down list 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** box, type **LON-DC1**, and then click **Next**.
9. On the **Specify access servers** page, click **Add**.
10. In the **Select a method to identify the initiator** dialog box, click **Enter a value for the selected type**, in the **Type** drop-down list box, click **IP Address**, in the **Value** text box, type **172.16.0.22**, and then click **OK**.
11. On the **Specify access servers** page, click **Add**.
12. In the **Select a method to identify the initiator** dialog box, click **Enter a value for the selected type**, in the **Type** drop-down list box, click **IP Address**, in the **Value** text box, type **131.107.0.2**, and then click **OK**.
13. On the **Specify access servers** page, click **Next**.
14. On the **Enable Authentication** page, click **Next**.

15. On the **Confirm selections** page, click **Create**.
16. On the **View results** page, wait until creation completes, and then click **Close**.
17. In the iSCSI VIRTUAL DISKS pane, click **TASKS**, and then in the **TASKS** drop-down list box, click **New iSCSI Virtual Disk**.
18. In the New iSCSI Virtual Disk Wizard, on the **Select iSCSI virtual disk location** page, under **Storage location**, click **C:**, and then click **Next**.
19. On the **Specify iSCSI virtual disk name** page, in the **Name** box, type **iSCSIDisk2**, and then click **Next**.
20. On the **Specify iSCSI virtual disk size** page, in the **Size** box, type **5**, in the drop-down list box, ensure **GB** is selected, and then click **Next**.
21. On the **Assign iSCSI target** page, click **lon-dc1**, and then click **Next**.
22. On the **Confirm selection** page, click **Create**.
23. On the **View results** page, wait until creation completes, and then click **Close**.
24. In the iSCSI VIRTUAL DISKS pane, click **TASKS**, and then in the **TASKS** drop-down list box, click **New iSCSI Virtual Disk**.
25. In the New iSCSI Virtual Disk Wizard, on the **Select iSCSI virtual disk location** page, under **Storage**, click **C:**, and then click **Next**.
26. On the **Specify iSCSI virtual disk name** page, in the **Name** text box, type **iSCSIDisk3**, and then click **Next**.
27. On the **Specify iSCSI virtual disk size** page, in the **Size** text box, type **5**, in the drop-down list box, ensure **GB** is selected, and then click **Next**.
28. On the **Assign iSCSI target** page, click **lon-dc1**, and then click **Next**.
29. On the **Confirm selection** page, click **Create**.
30. On the **View results** page, wait until creation completes, and then click **Close**.
31. In the iSCSI VIRTUAL DISKS pane, click **TASKS**, and then in the **TASKS** drop-down list box, click **New iSCSI Virtual Disk**.
32. In the New iSCSI Virtual Disk Wizard, on the **Select iSCSI virtual disk location** page, under **Storage**, click **C:**, and then click **Next**.
33. On the **Specify iSCSI virtual disk name** page, in the **Name** text box, type **iSCSIDisk4**, and then click **Next**.
34. 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**.
35. On the **Assign iSCSI target** page, click **lon-dc1**, and then click **Next**.
36. On the **Confirm selection** page, click **Create**.
37. On the **View results** page, wait until creation completes, and then click **Close**.
38. In the iSCSI VIRTUAL DISKS pane, click **TASKS**, and then in the **TASKS** drop-down list box, click **New iSCSI Virtual Disk**.
39. In the New iSCSI Virtual Disk Wizard, on the **Select iSCSI virtual disk location** page, under **Storage**, click **C:**, and then click **Next**.
40. On the **Specify iSCSI virtual disk name** page, in the **Name** text box, type **iSCSIDisk5**, and then click **Next**.

41. 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**.
42. On the **Assign iSCSI target** page, click **lon-dc1**, and then click **Next**.
43. On the **Confirm selection** page, click **Create**.
44. On the **View results** page, wait until creation completes, and then click **Close**.

► **Task 3: Configure MPIO**

1. Log on to **LON-SVR2** with username of **Adatum\Administrator** and the password of **Pa\$\$w0rd**.
2. In Server Manager, on the menu bar, click **Tools** and then in the **Tools** drop-down list, select **Routing and Remote access**.
3. In the Enable DirectAccess Wizard, click **Cancel**, and then click **OK** on the **Confirmation** dialog box.
4. Right-click **LON-SVR2** and then click **Disable Routing and Remote Access**. Click **Yes** and then close the Routing and Remote Access console.
5. In Server Manager click **Add roles and features**.
6. In the Add Roles and Features Wizard, on the **Before you begin** page, click **Next**.
7. On the **Select installation type** page, click **Next**.
8. On the **Select destination server** page, make sure that **Select server from the server pool** is selected, and then click **Next**.
9. On the **Select server roles** page, click **Next**.
10. On the **Select features** page, click **Multipath I/O**, and then click **Next**.
11. On the **Confirm installation selections** page, click **Install**.
12. When installation is complete, click **Close**.
13. In Server Manager, on the menu bar, click **Tools** and then in the **Tools** drop-down list, select **iSCSI Initiator**.
14. In the **Microsoft iSCSI** dialog box, click **Yes**.
15. In the **iSCSI Initiator Properties** dialog box, on the **Targets** tab, in the **Target** box, type **LON-DC1**, and then click **Quick Connect**. In the **Quick Connect** box, click **Done**.
16. Click **OK** to close the iSCSI Initiator Properties dialog box.
17. In Server Manager, on the menu bar, click **Tools**, and then in the **Tools** drop-down list, select **MPIO**.
18. In **MPIO Properties** dialog box, click the **Discover Multi-Paths** tab.
19. Select the **Add support for iSCSI devices** check box, and then click **Add**. When you are prompted to reboot the computer, click **Yes**.
20. After the computer restarts, log on to **LON-SVR2** with username of **Adatum\Administrator** and password of **Pa\$\$w0rd**.
21. In Server Manager, on the menu bar, click **Tools**, and then in the **Tools** drop-down list, select **MPIO**.
22. In the **MPIO Properties** dialog box, on the **MPIO Devices** tab, notice that additional **Device Hardware ID MSFT2005iSCSIBusType_0x9** is added to the list.
23. Click **OK** to close the MPIO Properties dialog box.

► Task 4: Connect to and configure the iSCSI targets

1. On LON-SVR2, in Server Manager, on the menu bar, click **Tools** and then in the **Tools** drop-down list, select **iSCSI Initiator**.
2. In the **iSCSI Initiator Properties** dialog box, on the **Targets** tab, click **Disconnect**.
3. In the **Disconnect From All Sessions** dialog box, click **Yes**.
4. In the **iSCSI Initiator Properties** dialog box, on the **Targets** tab, click **Connect**.
5. In the Connect to Target window, click **Enable multi-path**, verify that the **Add this connection to the list of Favorite Targets** check box is selected, and then click the **Advanced** button.
6. In the **Advanced Settings** dialog box, on the **General** tab, change the **Local Adapter** from **Default** to **Microsoft iSCSI Initiator**. In the **Initiator IP** drop-down list, click **172.16.0.22** and in the **Target Portal IP** drop-down list, click **172.16.0.10 / 3260**.
7. In the **Advanced Settings** dialog box, click **OK**.
8. In the Connect to Target window, click **OK**.
9. In the **iSCSI Initiator Properties** dialog box, on the **Targets** tab, click **Connect**.
10. In Connect to Target window, click **Enable multi-path**, verify that the **Add this connection to the list of Favorite Targets** check box is selected, and then click the **Advanced** button.
11. In the **Advanced Settings** dialog box, on the **General** tab, change the **Local Adapter** from **Default** to **Microsoft iSCSI Initiator**. In the Initiator IP drop-down list, select **131.107.0.2** and in the **Target Portal IP** drop-down list, select **131.107.0.1 / 3260**.
12. In the **Advanced Settings** dialog box, click **OK**.
13. In the Connect to Target window, click **OK**.
14. In the **iSCSI Initiator Properties** dialog box, click the **Volumes and Devices** tab.
15. In the **iSCSI Initiator Properties** dialog box, on the **Volumes and Devices** tab, click **Auto Configure**.
16. In the **iSCSI Initiator Properties** dialog box, click the **Targets** tab.
17. In the **Targets** list, select **iqn.1991-05.com.microsoft:lon-dc1-lon-dc1-target**, and then click **Devices**.
18. In the **Devices** dialog box, click the **MPIO** button.
19. Verify that in **Load balance policy**, **Round Robin** is selected. Under **This device has the following paths**, notice that two paths are listed. Select the first path and then click the **Details** button.
20. Note the IP address of the **Source** and **Target** portals, and then click **OK**.
21. Select the second path and then click the **Details** button.
22. Verify that the Source IP address is of the second network adapter, and then click **OK**.
23. Click **OK** to close the **Device Details** dialog box.
24. Click **OK** to close the **Devices** dialog box.
25. Click **OK** to close the **iSCSI Initiator Properties** dialog box.

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

Exercise 2: Configuring the File Classification Infrastructure

► Task 1: Create a classification property for corporate documentation

1. On LON-SVR1, in Server Manager, in the upper-right corner, click **Tools**, and then click **File Server Resource Manager**.
2. In the File Server Resource Manager window, expand **Classification Management**, select and then right-click **Classification Properties**, and then click **Create Local Property**.
3. In the Create Local Classification Property window, in the **Name** text box, type **Corporate Documentation**, in the **Property Type** drop-down list box ensure that **Yes/No** is selected, and then click **OK**.
4. Leave the File Server Resource Manager console open.

► Task 2: Create a classification rule for corporate documentation

1. In File Server Resource Manager, expand **Classification Management**, click **Classification Rules**, and then in the Actions pane, click **Create Classification Rule**.
2. In the Create Classification Rule window, on the **General** tab, in the **Rule name** text box, type **Corporate Documents Rule**, and then ensure that the **Enable** checkbox is selected.
3. Click the **Scope** tab, and then click **Add**.
4. In the Browse For Folder window, expand **Allfiles (E:\)**, expand **Labfiles**, click **Corporate Documentation**, and then click **OK**.
5. In the Create Classification Rule window, on the **Classification** tab, in the **Classification method** drop-down list box, click **Folder Classifier**, in the **Property-Choose a property to assign to files** drop-down list box, click **Corporate Documentation**, and then in the **Property-Specify a value** drop-down list box, click **Yes**.
6. Click the **Evaluation type** tab, click **Re-evaluate existing property values**, ensure that the **Aggregate the values** radio button is selected, and then click **OK**.
7. In File Server Resource Manager, in the Actions pane, click **Run classification with all rules now**.
8. In the Run classification window, select the **Wait for classification to complete** radio button, and then click **OK**.
9. Review the Automatic classification report that displays in Windows® Internet Explorer®, and ensure that the report lists the same number of classified files as in the Corporate Documentation folder.
10. Close Internet Explorer, but leave the File Server Resource Manager open.

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

1. In File Server Resource Manager, expand **Classification Management**, right-click **Classification Properties**, and then click **Create Local Property**.
2. In the Create Local Classification Property window, in the **Name** text box, type **Expiration Date**, in the **Property Type** drop-down list box, ensure that **Date-Time** is selected, and then click **OK**.
3. In File Server Resource Manager, expand **Classification Management**, click **Classification Rules**, and then in the Actions pane, click **Create Classification Rule**.
4. In the Create Classification Rule window, on the **General** tab, in the **Rule name** text box, type **Expiration Rule**, and ensure that the **Enable** check box is selected.
5. Click the **Scope** tab, and then click **Add**.

6. In the Browse For Folder window, expand **Allfiles (E:\)**, expand **Labfiles**, click **Corporate Documentation**, and then click **OK**.
7. Click the **Classification** tab, in the **Classification method** drop-down list box, click **Folder Classifier**, and then in the **Property-Choose a property to assign to files** drop-down list box, click **Expiration Date**.
8. Click the **Evaluation type** tab, click **Re-evaluate existing property values**, ensure that the **Aggregate the values** radio button is selected, and then click **OK**.
9. In File Server Resource Manager, in the Actions pane, click **Run classification with all rules now**.
10. In the Run classification window, select the **Wait for classification to complete** radio button, and then click **OK**.
11. Review the Automatic classification report that displays in Internet Explorer, and ensure that the report lists the same number of classified files as in the Corporate Documentation folder.
12. Close Internet Explorer, but leave the File Server Resource Manager open.

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

1. In File Server Resource Manager, select and then right-click **File Management Tasks**, and then click **Create File Management Task**.
2. In the Create File Management Task window, on the **General** tab, in the **Task name** text box, type **Expired Corporate Documents**, and then ensure that the **Enable** check box is selected.
3. Click the **Scope** tab, and then click **Add**.
4. In the Browse For Folder window, select **E:\Labfiles\Corporate Documentation**, and then click **OK**.
5. In the Create File Management Task window, on the **Action** tab, in the **Type** drop-down list box, ensure that **File expiration** is selected, and then in the **Expiration directory** box, type **E:\Labfiles\Expired**.
6. Click the **Notification** tab, and then click **Add**.
7. In the Add Notification window, in the **Event Log** tab, select the **Send warning to event log** check box, and then click **OK**.
8. Click the **Condition** tab, select the **Days since the file was last modified** check box, and then in the same row, replace the default value of 0 with **1**.



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

9. Click the **Schedule** tab, ensure that the **Weekly** radio button is selected, select the **Sunday** check box, and then click **OK**.
10. Leave the File Server Resource Manager open.

► **Task 5: Verify that corporate documents are expired**

1. In File Server Resource Manager, click **File Management Tasks**, right-click **Expired Corporate Documents**, and then click **Run File Management Task Now**.
2. In the Run File Management Task window, click **Wait for the task to complete**, and then click **OK**.
3. Review the File management task report that displays in Internet Explorer, and ensure that report lists the same number of classified files as in the Corporate Documentation folder.

4. In Server Manager, click **Tools**, and then click **Event Viewer**.
5. In the Event Viewer console, expand **Windows Logs**, and then click **Application**.
6. Review events with numbers **908** and **909**. Notice that 908 – FSRM started a file management job, and 909 – FSRM finished a file management job.

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

► **To prepare for the next lab**

When you finish the lab, revert 20417A-LON-SVR2. To do this, complete the following steps.

1. On the host computer, start **Hyper-V Manager**.
 2. In the **Virtual Machines** list, right-click **20417A-LON-SVR2**, and then click **Revert**.
 3. In the **Revert Virtual Machine** dialog box, click **Revert**.
- Keep all other virtual machines running for the next lab.

Lab B: Implementing BranchCache

Exercise 1: Configuring the Main Office Servers for BranchCache

► Task 1: Configure LON-DC1 to use BranchCache

1. On LON-DC1, on the taskbar, click the **Server Manager** icon.
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 (Installed)**, expand **File and iSCSI Services**, select the **BranchCache for Network Files** check box, and then click **Next**.
7. On the **Select features** page, click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. Click **Close** and then close Server Manager.
10. Point to the lower-right corner of the screen, click **Search**, in the **Search** text box, type **gpedit.msc**, and then press Enter.
11. In the Local Group Policy Editor console, in the navigation pane, under **Computer Configuration**, expand **Administrative Templates**, expand **Network**, and then click **Lanman Server**.
12. On the Lanman Server result pane, in the **Setting** list, right-click **Hash Publication for BranchCache**, and then click **Edit**.
13. In the **Hash Publication for BranchCache** dialog box, click **Enabled**, in the **Hash publication actions** list, select the **Allow hash publication only for shared folders on which BranchCache is enabled** check box, and then click **OK**.

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

1. In the Local Group Policy Editor console, in the navigation pane, under **Computer Configuration**, expand **Windows Settings**, right-click **Policy-based QoS**, and then click **Create new policy**.
2. In the Policy-based QoS Wizard, on the **Create a QoS policy** page, in the **Policy name** text box, type **Limit to 100 Kbps**, click the **Specify Outbound Throttle Rate** check box, and in the **Specify Outbound Throttle Rate** text box type **100**, and then click **Next**.
3. On the **This QoS policy applies to** page, click **Next**.
4. On the **Specify the source and destination IP addresses** page, click **Next**.
5. On the **Specify the protocol and port numbers** page, click **Finish**.
6. Close the Local Group Policy Editor.

► Task 3: Enable a file share for BranchCache

1. On the taskbar, click the Windows Explorer icon.
2. In the Windows Explorer window, browse to Local Disk (**C:**).
3. In the Local Disk (C:) window, on the menu, click the **Home** tab, and then click **New Folder**.
4. Type **Share**, and then press Enter.

5. Right-click **Share**, and then click **Properties**.
6. In the **Share Properties** dialog box, on the **Sharing** tab, click **Advanced Sharing**.
7. Select the **Share this folder** check box, and then click **Caching**.
8. In the **Offline Settings** dialog box, select the **Enable BranchCache** check box, and then click **OK**.
9. In the **Advanced Sharing** dialog box, click **OK**.
10. In the **Share Properties** dialog box, click **Close**.
11. Point to the lower-right corner of the screen, click **Search**, in the **Search** text box, type **cmd**, and then press Enter.
12. At the command prompt, type the following command, and then press Enter:

```
Copy C:\windows\system32\write.exe c:\share
```

13. Close the command prompt.
14. Close Windows Explorer.

► Task 4: Configure client firewall rules for BranchCache

1. On LON-DC1, on the taskbar, click the **Server Manager** icon.
2. In Server Manager, on the menu bar, click **Tools**, and then from the **Tools** drop-down list box, click **Group Policy Management**.
3. In Group Policy Management, expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**, right-click **Default Domain Policy**, and then click **Edit**.
4. In the Group Policy Management Editor, in the navigation pane, under **Computer Configuration**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, and then expand **Windows Firewall with Advanced Security**.
5. In Windows Firewall with Advanced Security, in the navigation pane, expand **Windows Firewall with Advanced Security**, and then click **Inbound Rules**.
6. In the Group Policy Management Editor, on the **Action** menu, click **New Rule**.
7. In the New Inbound Rule Wizard, on the **Rule Type** page, click **Predefined**, click **BranchCache – Content Retrieval (Uses HTTP)**, and then click **Next**.
8. On the **Predefined Rules** page, click **Next**.
9. On the **Action** page, click **Finish** to create the firewall inbound rule.
10. In the Group Policy Management Editor, in the navigation pane, click **Inbound Rules**, and then in the Group Policy Management Editor console, on the **Action** menu, click **New Rule**.
11. On the **Rule Type** page, click **Predefined**, click **BranchCache – Peer Discovery (Uses WSD)**, and then click **Next**.
12. On the **Predefined Rules** page, click **Next**.
13. On the **Action** page, click **Finish**.
14. Close the Group Policy Management Editor and Group Policy Management console.

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

Exercise 2: Configuring the Branch Office Servers for BranchCache

► Task 1: Install the BranchCache feature on LON-SVR1

1. On LON-SVR1, in Server Manager, click **Add roles and features**.
2. In the Add Roles and Features Wizard, on the **Before you begin** page, click **Next**.
3. On the **Select installation type** page, click **Next**.
4. On the **Select destination server** page, ensure that **Select server from the server pool** is selected, and then click **Next**.
5. On the **Select server roles** page, expand **File And Storage Services (Installed)**, expand **File and iSCSI Services**, and then select the **BranchCache for Network Files** check box.
6. On the **Select server roles** page, click **Next**.
7. On the **Select features** page, click **BranchCache**, and then click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. Click **Close**.

► Task 2: Start the BranchCache host server

1. Switch to LON-DC1.
2. In Server Manager, on the menu bar, click **Tools**, and then from the **Tools** drop-down list, click **Active Directory Users and Computers**.
3. In Active Directory Users and Computers, right-click **Adatum.com**, point to **New**, and then click **Organizational Unit**.
4. In the New Object - Organization Unit window, type **BranchCacheHost**, and then click **OK**.
5. Click the **Computers** container.
6. Click **LON-SVR1**, and then drag it to **BranchCacheHost**.
7. Click **Yes** to clear the warning about moving objects.
8. Close Active Directory Users and Computers.
9. In Server Manager, on the menu bar, click **Tools**, and then from the **Tools** drop-down list, click **Group Policy Management**.
10. In Group Policy Management, under **Domains**, expand **Adatum.com**, right-click **BranchCacheHost**, and then click **Block Inheritance**.
11. On LON-DC1, close all open windows.
12. Restart **LON-SVR1** and log on as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
13. On LON-SVR1, on the taskbar, click the **Windows PowerShell** icon.
14. In the Windows PowerShell window, type the following cmdlet, and then press Enter:

```
Enable-BCHostedServer  
-RegisterSCP
```

15. In the Windows PowerShell window, type the following cmdlet, and then press Enter:

```
Get-BCStatus
```

16. Close Windows PowerShell.

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

Exercise 3: Configuring Client Computers for BranchCache

► Task 1: Configure client computers to use BranchCache in Hosted Cache mode

1. On LON-DC1, on the taskbar, click **Server Manager**.
2. In Server Manager, on the menu bar, click **Tools** and then in the **Tools** drop-down list, select **Group Policy Management**.
3. In the Group Policy Management console, in the navigation pane, expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**, right-click **Default Domain Policy**, and then click **Edit**.
4. In the Group Policy Management Editor, in the navigation pane, under **Computer Configuration**, expand **Policies**, expand **Administrative Templates**, expand **Network**, and then click **BranchCache**.
5. In the BranchCache results pane, in the **Setting** list, right-click **Turn on BranchCache**, and then click **Edit**.
6. In the **Turn on BranchCache** dialog box, click **Enabled**, and then click **OK**.
7. In the BranchCache results pane, in the **Setting** list, right-click **Enable Automatic Hosted Cache Discovery by Service Connection Point**, and then click **Edit**.
8. In the **Enable Automatic Hosted Cache Discovery by Service Connection Point** dialog box, click **Enabled**, and then click **OK**.
9. In the BranchCache results pane, in the **Setting** list, right-click **Configure BranchCache for network files**, and then click **Edit**.
10. In the **Configure BranchCache for network files** dialog box, click **Enabled**, in the **Type the maximum round trip network latency (milliseconds) after which caching begins** text box, type **0**, and then click **OK**. This setting is required to simulate access from a branch office and is not typically required.
11. Close the Group Policy Management Editor.
12. Close the Group Policy Management Console.
13. Start **20412A-LON-CL1**, and log on as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
14. On the Start screen, in the lower-right corner of the screen, click **Search**, in the **Search** text box, type **cmd**, and then press Enter.
15. At the command prompt, type the following command, and then press Enter:

```
gpupdate /force
```
16. At the command prompt, type the following command, and then press Enter:

```
netsh branchcache show status all
```
17. Start **20412A-LON-CL2**, and log on as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
18. On the Start screen, in the lower-right corner of the screen, click **Search**, in the **Search** text box, type **cmd**, and then press Enter.
19. At the command prompt, type the following command, and then press Enter:

```
gpupdate /force
```

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

```
netsh branchcache show status all
```

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

Exercise 4: Monitoring BranchCache

► Task 1: Configure Performance Monitor on LON-SVR1

1. Switch to LON-SVR1.
2. In Server Manager, on the menu bar, click **Tools**, and then from the **Tools** drop-down list box, click **Performance Monitor**.
3. In the Performance Monitor console, in the navigation pane, under **Monitoring Tools**, click **Performance Monitor**.
4. In the Performance Monitor results pane, click the **Delete (Delete Key)** icon.
5. In the Performance Monitor results pane, click the **Add (Ctrl+N)** icon.
6. In the **Add Counters** dialog box, under **Select counters from computer**, click **BranchCache**, click **Add**, and then click **OK**.
7. On the **Change Graph type** button, select **Report**.

► Task 2: View performance statistics on LON-CL1

1. Switch to LON-CL1.
2. Point to the lower-right corner of the screen, click **Search**, in the **Search** text box, type **perfmon**, and then press Enter.
3. In the Performance Monitor console, in the navigation pane, under **Monitoring Tools**, click **Performance Monitor**.
4. In the Performance Monitor results pane, click the **Delete (Delete Key)** icon.
5. In the Performance Monitor results pane, click the **Add (Ctrl+N)** icon.
6. In the **Add Counters** dialog box, under **Select counters from computer**, click **BranchCache**, click **Add**, and then click **OK**.
7. Change graph type to **Report**. Notice that the value of all performance statistics is zero.

► Task 3: View performance statistics on LON-CL2

1. Switch to LON-CL2.
2. Point to the lower-right corner of the screen, click **Search**, in the **Search** text box, type **perfmon**, and then press Enter.
3. In the Performance Monitor console, in the navigation pane, under **Monitoring Tools**, click **Performance Monitor**.
4. In the Performance Monitor results pane, click the **Delete (Delete Key)** icon.
5. In the Performance Monitor results pane, click the **Add (Ctrl+N)** icon.

6. In the **Add Counters** dialog box, under **Select counters from computer**, click **BranchCache**, click **Add**, and then click **OK**.
7. Change graph type to **Report**. Notice that the value for all performance statistics is zero.

► **Task 4: Test BranchCache in the Hosted Cache mode**

1. Switch to LON-CL1.
2. On the taskbar, click the **Windows Explorer** icon.
3. In Windows Explorer, navigate to `\\LON-DC1.adatum.com\Share`, and then press Enter.
4. In the Share window, in the **Name** list, right-click **write.exe**, and then click **Copy**.
5. In the Share window, click **Minimize**.
6. On the desktop, right-click anywhere, and then click **Paste**.
7. Read the performance statistics on LON-CL1. This file was retrieved from LON-DC1 (Retrieval: Bytes from Server). After the file was cached locally, it was passed up to the hosted cache. (Retrieval: Bytes Served)
8. Switch to LON-CL2.
9. On the taskbar, click the **Windows Explorer** icon.
10. In the Windows Explorer address bar, type `\\LON-DC1.adatum.com\Share`, and then press Enter.
11. In the Share window, in the **Name** list, right-click **write.exe**, and then click **Copy**.
12. In the Share window, click **Minimize**.
13. On the **desktop**, right-click anywhere, and then click **Paste**.
14. Read the performance statistics on LON-CL2. This file was obtained from the hosted cache (Retrieval: Bytes from Cache).
15. Read the performance statistics on LON-SVR1. This server has offered cached data to clients (Hosted Cache: Client file segment offers made).

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

► **To prepare for the next module**

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

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20412A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20412A-LON-SVR1**, **20412A-LON-CL1**, and **20412A-LON-CL2**.

MCT USE ONLY. STUDENT USE PROHIBITED

Module 3: Implementing Dynamic Access Control

Lab: Implementing Dynamic Access Control

Exercise 1: Planning the Dynamic Access Control Implementation

► Task 1: Plan the Dynamic Access Control deployment

The scenario requires following:

1. Folders that belong to the Research team should be accessible and modifiable only by employees that belong to the Research team.
2. Files classified with classification **High** should be accessible only to managers.
3. Managers should access confidential files only from workstations that belong to the ManagersWKS security group.

You can meet these requirements by implementing claims, resource properties, and file classifications, as follows:

1. Create the appropriate claims for users and devices. The user claim uses department as the source attribute, and the device claim uses description as source attribute.
2. Configure the resource property for the research department.
3. Configure central access rules and central access policies to protect the resources. At the same time, you should configure file classification for confidential documents.
4. Apply a central access policy to folders in which files for the research departments and managers are located.
5. As a solution for users who receive error messages, you should implement Access Denied Assistance.

► Task 2: Prepare AD DS to support Dynamic Access Control

1. On LON-DC1, in Server Manager, click **Tools**, and then click **Active Directory Users and Computers**.
2. In the Active Directory Users and Computers console, right-click **Adatum.com**, click **New**, and then click **Organizational Unit**.
3. In the **New Object – Organizational Unit** dialog box, in the **Name** field, type **Test**, and then click **OK**.
4. Click the **Computers** container.
5. Press and hold the Ctrl key, click the **LON-SVR1**, **LON-CL1**, and **LON-CL2** computers, right-click, and then select **Move**.
6. In the Move window, click **Test**, and then click **OK**.
7. Close the Active Directory Users and Computers console.
8. On LON-DC1, in Server Manager, click **Tools**, and then click **Group Policy Management**.
9. Expand **Forest: Adatum.com**, expand **Domains**, and then expand **Adatum.com**.
10. Right-click the **Managers** OU, and then click **Block Inheritance**. This is to remove the block inheritance setting used in a later module in the course.
11. Click the **Group Policy Objects** container.
12. In the results pane, right-click **Default Domain Controllers Policy**, and then click **Edit**.

13. In the Group Policy Management Editor, under **Computer Configuration**, expand **Policies**, expand **Administrative Templates**, expand **System**, and then click **KDC**.
14. In the right pane, double-click **KDC support for claims, compound authentication and Kerberos armoring**.
15. In the KDC support for claims, compound authentication and Kerberos armoring window, select **Enabled**, in the **Options** section, click the drop-down list, select **Supported**, and then click **OK**.
16. Close the Group Policy Management Editor console and Group Policy Management Console.
17. On the taskbar, click the **Windows PowerShell®** icon.
18. At the Windows PowerShell command-line interface, type **gpupdate /force**, and then press Enter. After Group Policy updates, close the Windows PowerShell window.
19. On LON-DC1, in Server Manager, click **Tools**, and then click **Active Directory Users and Computers**.
20. Expand **Adatum.com**, right-click **Users**, click **New**, and then click **Group**.
21. In the **Group name** field, type **ManagersWKS**, and then click **OK**.
22. Click the **Test** OU.
23. Right-click **LON-CL1**, and then click **Properties**.
24. Click the **Member Of** tab, and then click **Add**.
25. In the Select Groups window, type **ManagersWKS**, click **Check Names**, click **OK**, and then click **OK** again.
26. Click the **Managers** organization unit (OU).
27. Right-click **Aidan Delaney**, and then select **Properties**.
28. Click the **Organization** tab. Ensure that the **Department** field is populated with the value **Managers**, and then click **Cancel**.
29. Click the **Research** OU.
30. Right-click **Allie Bellew**, and select **Properties**.
31. Click the **Organization** tab. Ensure that the **Department** field is populated with the value **Research**, and then click **Cancel**.

Results: After completing this exercise, you will have planned for Dynamic Access Control deployment, and you will have prepared AD DS for Dynamic Access Control implementation.

Exercise 2: Configuring User and Device Claims

► Task 1: Review the default claim types

1. On LON-DC1, in Server Manager, click **Tools**, and then click **Active Directory Administrative Center**.
2. In the Active Directory Administrative Center, in the navigation pane, click **Dynamic Access Control**.
3. In the central pane, double-click **Claim Types**.
4. Verify that there are no default claims defined.
5. In the navigation pane, click **Dynamic Access Control**, and then double-click **Resource Properties**.
6. Review the default resource properties. Note that all properties are disabled by default.

7. In the navigation pane, click **Dynamic Access Control**, and then double-click **Resource Property Lists**.
8. In the central pane, right-click **Global Resource Property List**, and then click **Properties**.
9. In the Global Resource Property List, in the Resource Properties section, review the available resource properties, and then click **Cancel**.

► **Task 2: Configure claims for users**

1. In the Active Directory Administrative Center, in the navigation pane, click **Dynamic Access Control**, and then double-click **Claim Types**.
2. In the Claim Types container, in the **Tasks** pane, click **New**, and then click **Claim Type**.
3. In the Create Claim Type window, in the Source Attribute section, select **department**.
4. In the **Display name** text box, type **Company Department**.
5. Select both **User** and **Computer** check boxes, and then click **OK**.

► **Task 3: Configure claims for devices**

1. In the Active Directory Administrative Center, in the **Tasks** pane, click **New**, and then select **Claim Type**.
2. In the Create Claim Type window, in the Source Attribute section, click **description**.
3. Clear the **User** check box, select the **Computer** check box, and then click **OK**.

Results: After completing this exercise, you will have reviewed the default claim types, configured claims for users, and configured claims for devices.

Exercise 3: Configuring Resource Property Definitions

► **Task 1: Configure resource property definitions**

1. In the Active Directory Administrative Center, click **Dynamic Access Control**.
2. In the central pane, double-click **Resource Properties**.
3. In the **Resource Properties** list, right-click **Department**, and then click **Enable**.
4. In the **Resource Properties** list, right-click **Confidentiality**, and then click **Enable**.
5. In the **Global Resource Property List**, ensure that both the **Department** and **Confidentiality** properties are enabled.
6. Double-click **Department**.
7. Scroll down to the Suggested Values section, and then click **Add**.
8. In the Add a suggested value window, in both **Value** and **Display name** text boxes, type **Research**, and then click **OK** two times.
9. Click **Dynamic Access Control**, and then double-click **Resource Property Lists**.
10. In the central pane, double-click **Global Resource Property List**, ensure that both **Department** and **Confidentiality** display and then click **Cancel**. If they do not display, click **Add**, add these two properties, and then click **OK**.
11. Close the Active Directory Administrative Center.

► Task 2: Classify files

1. On LON-SVR1, in Server Manager, click **Tools**, and then click **File Server Resource Manager**.
2. In the File Server Resource Manager console, expand **Classification Management**.
3. Select and then right-click **Classification Properties**, and then click **Refresh**.
4. Verify that **Confidentiality** and **Department** properties are listed.
5. Click **Classification Rules**.
6. In the Actions pane, click **Create Classification Rule**.
7. In the Create Classification Rule window, for the **Rule name**, type **Set Confidentiality**.
8. Click the **Scope** tab, and then click **Add**.
9. In the **Browse For Folder** dialog box, expand **Local Disk (C:)**, click the **Docs** folder, and then click **OK**.
10. Click the **Classification** tab.
11. Make sure that following settings are set, and then click **Configure**:
 - Classification method: **Content Classifier**
 - Property: **Confidentiality**
 - Value: **High**
12. In the **Classification Parameters** dialog box, click the **Regular expression** drop-down list, and then click **String**.
13. In the **Expression** field (next to the word String,) type **secret**, and then click **OK**.
14. Click the **Evaluation Type** tab, select **Re-evaluate existing property values**, click **Overwrite the existing value**, and then click **OK**.
15. In the File Server Resource Manager, in the Actions pane, click **Run Classification with all rules now**.
16. Click **Wait for classification to complete**, and then click **OK**.
17. After the classification is complete, you will be presented with a report. Verify that two files were classified. You can confirm this in the Report Totals section.
18. Close the report.
19. Open a Windows Explorer window, and browse to the **C:\Docs** folder.
20. Right-click **Doc1.txt**, click **Properties**, and then click the **Classification** tab. Verify that **Confidentiality** is set to **High**.
21. Repeat step 20 on files **Doc2.txt** and **Doc3.txt**. Doc2.txt should have same Confidentiality as Doc1.txt, while Doc3.txt should have no value. This is because only Doc1 and Doc2 have the word secret in their content.

► Task 3: Assign properties to a folder

1. On LON-SVR1, open Windows Explorer, and browse to drive **C**.
2. In drive C, right-click the **Research** folder, and then click **Properties**.
3. Click the **Classification** tab, and then click **Department**. In the Value section, click **Research**, click **Apply**, and then click **OK**.

Results: After completing this exercise, you will have configured resource properties for files, classified files, and assigned properties to a folder.

Exercise 4: Configuring Central Access Rules and Central Access Policies

► Task 1: Configure central access rules

1. On LON-DC1, in Server Manager, click **Tools**, and then click **Active Directory Administrative Center**.
2. In the Active Directory Administrative Center, in the navigation pane, click **Dynamic Access Control**, and then double-click **Central Access Rules**.
3. In the Tasks pane, click **New**, and then click **Central Access Rule**.
4. In the **Central Access Rule** dialog box, in the **Name** field, type **Department Match**.
5. In the Target Resources section, click **Edit**.
6. In the **Central Access Rule** dialog box, click **Add a condition**.
7. Set a condition as follows: **Resource-Department-Equals-Value-Research**, and then click **OK**.
8. In the Permissions section, click **Use following permissions as current permissions**.
9. In the Permissions section, click **Edit**.
10. Remove permission for **Administrators**.
11. In **Advanced Security Settings for Permissions**, click **Add**.
12. In **Permission Entry for Permissions**, click **Select a principal**.
13. In the Select User, Computer, Service Account or Group window, type **Authenticated Users**, click **Check Names**, and then click **OK**.
14. In the Basic permissions section, select the **Modify, Read and Execute, Read** and **Write** check boxes.
15. Click **Add a condition**.
16. Click the **Group** drop-down list, and then click **Company Department**.
17. Click the **Value** drop-down list, and then select **Resource**.
18. In the last drop-down list, select **Department**.



Note: You should have this expression as a result: **User-Company Department-Equals-Resource-Department**.

19. Click **OK** three times.
20. In the tasks pane, click **New**, and then click **Central Access Rule**.
21. For the name of rule, type **Access Confidential Docs**.
22. In the Target Resources section, click **Edit**.
23. In the Central Access Rule window, click **Add a condition**.
24. In the last drop-down list, click **High**.



Note: You should have this expression as a result: **Resource-Confidentiality-Equals-Value-High**.

25. Click **OK**.
26. In the Permissions section, click **Use following permissions as current permissions**.
27. In the Permissions section, click **Edit**.
28. Remove permission for **Administrators**.
29. In Advanced Security Settings for Permissions, click **Add**.
30. In the Permission Entry for Permissions, click **Select a principal**.
31. In the Select User, Computer, Service Account or Group window, type **Authenticated Users**, click **Check Names**, and then click **OK**.
32. In the Basic permissions section, select the **Modify, Read and Execute, Read** and **Write** check boxes.
33. Click **Add a condition**.
34. Set the first condition to:
User-Group-Member of each-Value-Managers, and then click Add a condition.



Note: If you cannot find Managers in the last drop-down list, click **Add items**. Then in the Select User, Computer, Service Account or Group window, type **Managers**, click **Check Names**, and then click **OK**.

35. Set the second condition to: **Device-Group-Member of each-Value-ManagersWKS**.



Note: If you cannot find ManagersWKS in the last drop-down list, click **Add items**. Then in the Select User, Computer, Service Account or Group window, type **ManagersWKS**, click **Check Names**, and then click **OK**.

36. Click **OK** three times.

► **Task 2: Create a central access policy**

1. On LON-DC1, in the Active Directory Administrative Center, click **Dynamic Access Control**, and then double-click **Central Access Policies**.
2. In the tasks pane, click **New**, and then click **Central Access Policy**.
3. In the **Name** field, type **Protect confidential docs**, and then click **Add**.
4. Click the **Access Confidential Docs** rule, click >>, and then click **OK** twice.
5. In the tasks pane, click **New**, and then click **Central Access Policy**.
6. In the **Name** field, type **Department Match**, and then click **Add**.
7. Click the **Department Match** rule, click >>, and then click **OK** twice.
8. Close the Active Directory Administrative Center.

► **Task 3: Publish a central access policy by using Group Policy**

1. On LON-DC1, in Server Manager, click **Tools**, and then click **Group Policy Management**.
2. In the Group Policy Management Console, under **Domains**, expand **Adatum.com**, right click **Test**, and then click **Create a GPO in this domain, and link it here**.
3. Type **DAC Policy**, and then click **OK**.

4. Right-click **DAC Policy**, and then click **Edit**.
5. Expand Computer **Configuration**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, expand **File System**, right-click **Central Access Policy**, and then click **Manage Central Access Policies**.
6. Click both **Department Match** and **Protect confidential docs**, click **Add**, and then click **OK**.
7. Close the Group Policy Management Editor.
8. Close the Group Policy Management Console.

► **Task 4: Apply the central access policy to resources**

1. On LON-SVR1, on the taskbar, click the **Windows PowerShell** icon.
2. Type **gpupdate /force**, and then press Enter.
3. Close the command prompt window.
4. Open Windows Explorer, browse to drive **C**, right-click the **Docs** folder, and then click **Properties**.
5. In the **Properties** dialog box, click the **Security** tab, and then click **Advanced**.
6. In the Advanced Security Settings for Docs window, click the **Central Policy** tab, and then click **Change**.
7. On the drop-down list, select **Protect confidential docs**, and then click **OK** two times.
8. Right-click the **Research** folder, and then click **Properties**.
9. In the **Properties** dialog box, click the **Security** tab, and then click **Advanced**.
10. In the Advanced Security Settings for Research window, click the **Central Policy** tab, and then click **Change**.
11. In the drop-down list, select **Department Match**, and then click **OK** two times.

► **Task 5: Configure access denied remediation settings**

1. On LON-DC1, in Server Manager, click **Tools**, and then click **Group Policy Management**.
2. In the Group Policy Management Console, expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**, and then click **Group Policy objects**.
3. Right-click **DAC Policy**, and then select **Edit**.
4. Under **Computer Configuration**, expand **Policies**, expand **Administrative Templates**, expand **System**, and then click **Access-Denied Assistance**.
5. In the right pane, double-click **Customize Message for Access Denied errors**.
6. In the Customize Message for Access Denied errors window, click **Enabled**.
7. In the **Display the following message to users who are denied access** text box, type **You are denied access because of permission policy. Please request access**.
8. Select the **Enable users to request assistance** check box.
9. Review other options but do not make any changes, and then click **OK**.
10. In the right pane of the Group Policy Management Editor, double-click **Enable access-denied assistance on client for all file types**.
11. Click **Enabled**, and then click **OK**.
12. Close both the Group Policy Management Editor and the Group Policy Management Console.

13. Switch to LON-SVR1, on the taskbar click the Windows PowerShell icon.
14. At the Windows PowerShell command-line interface, type **gpupdate /force**, and then press Enter.

Results: After completing this exercise, you will have configured central access rules and central access policies for Dynamic Access Control.

Exercise 5: Validating and Remediating Dynamic Access Control

► Task 1: Validate Dynamic Access Control functionality

1. Start and then log on to **LON-CL1** as **Adatum\April** with the password **Pa\$\$w0rd**.
2. Click the Desktop tile, and then on the taskbar, click the **Windows Explorer** icon.
3. In the Windows Explorer address bar, type **\\LON-SVR1\Docs**, and then press Enter.
4. In the **Docs** folder, try to open **Doc3**. You should be able to open that document. Close notepad.
5. In the Windows Explorer address bar, type **\\LON-SVR1\Research**, and then press Enter. You should be unable to access folder.
6. Click **Request assistance**. Review options for sending messages, and then click **Close**.
7. Log off LON-CL1.
8. Log on to LON-CL1 as **Adatum>Allie** with the password **Pa\$\$w0rd**.
9. Click the Desktop tile, and then on the taskbar, click the **Windows Explorer** icon.
10. In the Windows Explorer address bar, type **\\LON-SVR1\Research**, and then press Enter.
11. Verify that you can access this folder and open documents inside, because Allie is a member of the Research team.
12. Log off LON-CL1.
13. Log on to **LON-CL1** as **Adatum\Aidan** with the password **Pa\$\$w0rd**.
14. Click the Desktop tile, and then on the taskbar, click the **Windows Explorer** icon.
15. In the Windows Explorer address bar, type **\\LON-SVR1\Docs**.
16. Verify that you can access this folder and open all files inside.
17. Log off LON-CL1.
18. Start and then log on to **LON-CL2** as **Adatum\Aidan** with the password **Pa\$\$w0rd**.
19. Click the Desktop tile, and then on the taskbar, click the **Windows Explorer** icon. In the Windows Explorer address bar, type **\\LON-SVR1\Docs**. You should be unable to view Doc1 or Doc2, because the LON-CL2 is not permitted to view secret documents.

Results: After completing this exercise, you will have validated Dynamic Access Control functionality.

Exercise 6: Implementing new resource policies

► Task 1: Configure staging for a central access policy

1. On LON-DC1, open Server Manager, click **Tools**, and then select **Group Policy Management**.

2. In the Group Policy Management Console, expand **Forest:adatum.com**, expand **Domains**, expand **Adatum.com**, and then click **Group Policy object**.
3. Right-click **DAC Policy**, and then select **Edit**.
4. In the Group Policy Management Editor, expand **Computer Configuration**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, expand **Advanced Audit Policy Configuration**, expand **Audit Policies**, and then click **Object Access**.
5. Double-click **Audit Central Access Policy Staging**, select all three check boxes, and then click **OK**.
6. Double-click **Audit File System**, select all three check boxes, and then click **OK**.
7. Close the Group Policy Management Editor and the Group Policy Management console.

► Task 2: Configure staging permissions

1. On LON-DC1, open Server Manager, and then open Active Directory Administrative Center.
2. In the navigation pane, click **Dynamic Access Control**.
3. Double-click **Central Access Rules**.
4. Right-click **Department Match**, and then select **Properties**.
5. Scroll down to the Proposed Permissions section, click **Enable permission staging configuration**, and then click **Edit**.
6. Click **Authenticated Users**, and then click **Edit**.
7. Change the condition to: **User-Company Department-Equals-Value-Marketing**, and then click **OK**.
8. Click **OK** two more times to close all windows.
9. Switch to LON-SVR1 and open Windows PowerShell.
10. Type **gpupdate /force** and press Enter.
11. Close Windows PowerShell window.

► Task 3: Verify staging

1. Log on to **LON-CL1** as **Adatum\Adam** with the password **Pa\$\$w0rd**.
2. Click the Desktop tile, and then on the taskbar, click the Windows **Explorer** icon. In the Windows Explorer address bar, click the yellow icon, and then type **\\LON-SVR1\Research**.
3. Try to open the **Research** folder and its files. You will not be able to open it.
4. Switch to LON-SVR1.
5. Open Server Manager, click **Tools**, and then select **Event Viewer**.
6. Expand **Windows Logs**, and then navigate to **Security Log**.
7. Look for Events with ID **4818**.
8. Read the content of these logs.

► Task 4: Use effective permissions to test Dynamic Access Control

1. On LON-SVR1, open Windows Explorer.
2. In the Windows Explorer window, navigate to **C:\Research**, right-click **Research**, and then click **Properties**.
3. In the **Properties** dialog box, click the **Security** tab, click **Advanced**, and then click **Effective Access**.

4. Click **select a user**.
5. In the Select User, Computer, Service Account, or Group window, type **April**, click **Check Names**, and then click **OK**.
6. Click **View effective access**.
7. Review the results. The user April should not have access to this folder.
8. Click **Include a user claim**.
9. On the drop-down list, click **Company Department**.
10. In the **Value** text box, type **Research**.
11. Click **View Effective access**. The user should now have access.
12. Close all open windows.

Results: After completing this exercise, you will have implemented new resource policies.

► **To prepare for the next module**

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20412A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20412A-LON-SVR1**, **20412A-LON-CL1**, and **20412A-LON-CL2**.

Module 4: Implementing Network Load Balancing

Lab: Implementing Network Load Balancing

Exercise 1: Implementing an NLB Cluster

► Task 1: Verify website functionality for standalone servers

1. On LON-SVR1, on the taskbar, click the Windows Explorer icon.
2. Navigate to the folder `c:\inetpub\wwwroot`.
3. Double-click the file `iis-8.png`. This will open the file in Microsoft® Paint.
4. Ensure that the **Paint Brush** tool is selected, and then in the palette, click the color **Red**.
5. Use the mouse to mark the IIS Logo distinctively, using the color red.
6. Save the changes that you made to `iis-8.png`, and then close Microsoft Paint.
7. Close Windows Explorer.
8. Switch to LON-DC1.
9. Click to the Start screen.
10. Click the **Internet Explorer**® icon.
11. In the Internet Explorer address bar, type the address `http://LON-SVR1` and then press Enter. Verify that the web page displays the IIS logo with the distinctive color red mark that you added.
12. In the Internet Explorer address bar, enter the address `http://LON-SVR2` and then press Enter. Verify that the web page does not display the marked IIS logo.
13. Close Internet Explorer.

► Task 2: Install the Windows Network Load Balancing feature

1. Switch to LON-SVR1.
2. In the Server Manager console, click the **Tools** menu, and then click **Windows PowerShell ISE**.
3. In the blue PowerShell ISE window, enter the following command and then press Enter:

```
Invoke-Command -ComputersName LON-SVR1,LON-SVR2 -command {Install-WindowsFeature  
NLB,RSAT-NLB}
```

► Task 3: Create a new Windows Server 2012 NLB cluster

1. On LON-SVR1, in the Windows PowerShell ISE window, type the following command, and then press Enter:

```
New-NlbCluster -InterfaceName "Local Area Connection" -OperationMode Multicast -  
ClusterPrimaryIP 172.16.0.42 -ClusterName LON-NLB
```

2. In the Windows PowerShell ISE window, type the following command, and then press Enter:

```
Invoke-Command -ComputersName LON-DC1 -command {Add-DNSServerResourceRecordA  
zonename adatum.com -name LON-NLB -Ipv4Address 172.16.0.42}
```

► **Task 4: Add a second host to the cluster**

- On LON-SVR1, in the Windows PowerShell ISE window, type the following command and then press Enter:

```
Add-NlbClusterNode -InterfaceName "Local Area Connection" -NewNodeName "LON-SVR2" -  
NewNodeInterface "Local Area Connection"
```

► **Task 5: Validate the NLB cluster**

1. On LON-SVR1, in the Server Manager Console, click the **Tools** menu, and then click **Network Load Balancing Manager**.
2. In the Network Load Balancing Manager, verify that nodes LON-SVR1 and LON-SVR2 display with the status of **Converged** for the LON-NLB cluster.
3. Right-click the **LON-NLB** cluster, and then click **Cluster properties**.
4. On the **Cluster Parameters** tab, verify that the cluster is set to use the **Multicast** operations mode.
5. On the **Port Rules** tab, verify that there is a single port rule named **All** that starts at port **0** and ends at port **65535** for both **TCP** and **UDP** protocols, and that it uses **Single** affinity.
6. Click **OK** to close the **Cluster Properties** dialog box.

Results: After this exercise, you should have successfully implemented an NLB cluster.

Exercise 2: Configuring and Managing the NLB Cluster

► Task 1: Configure port rules and affinity

1. On LON-SVR2, on the taskbar, click the **Windows PowerShell** icon.
2. In Windows PowerShell, type each of the following commands, pressing Enter after each command:

```
Cmd.exe
Mkdir c:\porttest
Xcopy /s c:\inetpub\wwwroot c:\porttest
Exit
New-Website -Name PortTest -PhysicalPath "C:\porttest" -Port 5678
New-NetFirewallRule -DisplayName PortTest -Protocol TCP -LocalPort 5678
```

3. On the taskbar, click the **Windows Explorer** icon.
4. Click drive **C**, double-click the **porttest** folder, and then double-click **iis-8.png**. This will open Microsoft Paint.
5. Select the color blue from the palette.
6. Use the **Blue** paintbrush to mark the IIS Logo in a distinctive manner.
7. Save the changes to iis-8.png, and then close Microsoft Paint.
8. Switch to LON-DC1.
9. Click to the Start screen.
10. On the Start screen, click the **Internet Explorer** icon.
11. In the Internet Explorer address bar, type **http://LON-SVR2:5678** and then press Enter.
12. Verify that the IIS Start page with the IIS logo distinctively marked with blue displays in Internet Explorer.
13. Switch to LON-SVR1.
14. On LON-SVR1, switch to **Network Load Balancing Manager**.
15. In the Network Load Balancing Manager, right click **LON-NLB**, and then click **Cluster Properties**.
16. In the **Cluster Properties** dialog box, on the **Port Rules** tab, select the **All port** rule, and then click **Remove**.
17. On the **Port Rules** tab, click **Add**.
18. In the **Add/Edit Port Rule** dialog box, enter the following information, and then click **OK**:
 - Port range: **80 to 80**
 - Protocols: **Both**
 - Filtering mode: **Multiple Host**
 - Affinity: **None**
19. Click **OK**.
20. On the **Port Rules** tab, click **Add**.
21. In the **Add/Edit Port Rule** dialog box, enter the following information, and then click **OK**:
 - Port range: **5678 to 5678**
 - Protocols: **Both**
 - Filtering mode: **Single Host**

22. Click **OK** to close the **Cluster Properties** dialog box.
23. In the Network Load Balancing Manager, right click **LON-SVR1**, and then click **Host Properties**.
24. In the **Port Rules** tab, click the port rule that has **5678** as the Start and End value, and then click **Edit**.
25. Click the **Handling priority** value, and change it to **10**.
26. Click **OK** twice to close both the **Add/Edit Port Rule** dialog box and the **Host Properties** dialog box.

► **Task 2: Validate port rules**

1. Switch to LON-DC1.
2. Switch to the Start screen.
3. On the Start screen, click the **Internet Explorer** icon.
4. In the Internet Explorer address bar, type **http://lon-nlb**, and then press Enter.
5. Click the **Refresh** icon 20 times. Verify that you see web pages with and without the distinctive red marking.
6. On LON-DC1, verify that you have Internet Explorer open.
7. In the address bar, enter the address **http://LON-NLB:5678**, and press Enter.
8. In the address bar, click the **Refresh** icon 20 times. Verify that you are able to view only the web page with the distinctive blue marking.

► **Task 3: Manage host availability in the NLB Cluster**

1. Switch to LON-SVR1.
2. Select the Network Load Balancing Manager.
3. Right-click **LON-SVR1**, click **Control Host**, and then click **Suspend**.
4. Click the **LON-NLB** node. Verify that node LON-SVR1 displays as **Suspended**, and that node LON-SVR2 displays as **Converged**.
5. Right-click **LON-SVR1**, click **Control Host**, and then click **Resume**.
6. Right-click **LON-SVR1**, click **Control Host**, and then click **Start**.
7. Click the **LON-NLB** node. Verify that both nodes LON-SVR1 and LON-SVR2 now display as **Converged**. You may have to refresh the view.

Results: After this exercise, you should have successfully configured and managed an NLB cluster.

Exercise 3: Validating High Availability for the NLB Cluster

► **Task 1: Validate website availability when the host is unavailable**

1. On LON-SVR1, on the taskbar, click the **Windows PowerShell** icon.
2. Type the following command, and then press Enter:

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

3. Switch to LON-DC1.
4. On LON-DC1, open Internet Explorer.
5. In the Internet Explorer address bar, type the address **http://LON-NLB**, and then press Enter.

6. Refresh the website 20 times. Verify that the website is available while LON-SVR1 reboots, but that it does not display the distinctive red mark on the IIS logo until LON-SVR1 has completed the reboot cycle.

► **Task 2: Configure and validate Drainstop**

1. Log on to **LON-SVR1** with the username **Adatum\Administrator** and the password **Pa\$\$word**.
2. In Server Manager, click the **Tools** menu, and then click **Network Load Balancing Manager**.
3. In the Network Load Balancing Manager console, right-click **LON-SVR2**, click **Control Host**, and then click **Drainstop**.
4. Switch to LON-DC1.
5. In Internet Explorer, in the address bar, type **http://lon-nlb**, and then press Enter.
6. Refresh the site 20 times, and verify that only the welcome page with the red IIS logo displays.

Results: After this exercise, you should have successfully validated high availability for the NLB cluster.

► **To prepare for the next module**

When you finish the lab, revert the virtual machines to their initial state.

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

MCT USE ONLY. STUDENT USE PROHIBITED

Module 5: Implementing Failover Clustering

Lab: Implementing Failover Clustering

Exercise 1: Configuring a Failover Cluster

► Task 1: Connect cluster nodes to the iSCSI targets

1. On LON-SVR3, in Server Manager, click **Tools**, and then click **iSCSI Initiator**.
2. In the **Microsoft iSCSI** dialog box, click **Yes**.
3. Click the **Discovery** tab.
4. Click **Discover Portal**.
5. In the **IP address or DNS name** box, type **172.16.0.21**, and then click **OK**.
6. Click the **Targets** tab.
7. Click **Refresh**.
8. In the **Targets** list, select **iqn.1991-05.com.microsoft:LON-SVR1-target1-target**, and then click **Connect**.
9. Select **Add this connection to the list of Favorite Targets**, and then click **OK** two times.
10. On LON-SVR4, in Server Manager, click **Tools**, and then click **iSCSI Initiator**.
11. In the **Microsoft iSCSI** dialog box, click **Yes**.
12. Click the **Discovery** tab.
13. Click **Discover Portal**.
14. In the **IP address or DNS name** box, type **172.16.0.21**, and then click **OK**.
15. Click the **Targets** tab.
16. Click **Refresh**.
17. In the **Targets** list, select **iqn.1991-05.com.microsoft:LON-SVR1-target1-target**, and then click **Connect**.
18. Select **Add this connection to the list of Favorite Targets**, and then click **OK** two times.
19. On LON-SVR3, in **Server Manager**, click **Tools**, and then click **Computer Management**.
20. Expand **Storage**, and then click **Disk Management**.
21. Right-click **Disk 1**, and then click **Online**.
22. Right-click **Disk 1**, and then click **Initialize disk**. In the **Initialize Disk** dialog box, click **OK**.
23. Right-click the unallocated space next to **Disk 1**, and then click **New Simple Volume**.
24. On the **Welcome** page, click **Next**.
25. On the **Specify Volume Size** page, click **Next**.
26. On the **Assign Drive Letter or Path** page, click **Next**.
27. On the **Format Partition** page, in the **Volume Label** box, type **Data**. Select the **Perform a quick format** check box, and then click **Next**.
28. Click **Finish**. (Note: If the Microsoft Windows window pops up with prompt to format the disk, click **Cancel**.)

29. Repeat steps 21 through 28 for **Disk 2 and Disk 3**. (Note: Use Data2 and Data3 for Volume Labels).
30. Close the Computer Management window.
31. On LON-SVR4, in **Server Manager**, click **Tools**, and then click **Computer Management**.
32. Expand **Storage**, and then click **Disk Management**.
33. Right-click **Disk Management**, and then click **Refresh**.
34. Right-click **Disk 1**, and then click **Online**.
35. Right-click **Disk 2**, and then click **Online**.
36. Right-click **Disk 3**, and then click **Online**.
37. Close the Computer Management window.

► **Task 2: Install the failover clustering feature**

1. On LON-SVR3, in Server Manager, click **Add roles and features**.
2. On the **Before you begin** page, click **Next**.
3. On the **Select installation type** page, click **Next**.
4. On the **Select destination server** page, make sure that **Select server from the server pool** is selected, and then click **Next**.
5. On the **Select server roles** page, click **Next**.
6. On the **Select features** page, in the **Features** list, click **Failover Clustering**. In the Add features that are required for Failover Clustering? window, click **Add Features**. Click **Next**.
7. On the **Confirm installation selections** page, click **Install**.
8. When installation is complete, click **Close**.
9. Repeat steps 1 through 8 on LON-SVR4.

► **Task 3: Validate the servers for failover clustering**

1. On LON-SVR3, in the Server Manager, click **Tools**, and then click **Failover Cluster Manager**.
2. In the Actions pane of the Failover Cluster Manager, click **Validate Configuration**.
3. In the Validate a Configuration Wizard, click **Next**.
4. In the **Enter Name** box, type **LON-SVR3**, and then click **Add**.
5. In the **Enter Name** box, type **LON-SVR4**.
6. Click **Add**, and then click **Next**.
7. Verify that **Run all tests (recommended)** is selected, and then click **Next**.
8. On the **Confirmation** page, click **Next**.
9. Wait for the validation tests to finish (it might take up to 5 minutes), and then on the **Summary** page, click **View Report**.
10. Verify that all tests completed without errors. Some warnings are expected.
11. Close Internet Explorer.
12. On the **Summary** page, remove the check mark next to **Create the cluster now using the validated nodes**, click **Finish**.

► Task 4: Create the failover cluster

1. On LON-SVR3, in Failover Cluster Manager, in the center pane, under **Management**, click **Create Cluster**.
2. In the Create Cluster Wizard on the **Before You Begin** page, read the information.
3. Click **Next**, in the **Enter server name** box, type **LON-SVR3**, and then click **Add**. Type **LON-SVR4**, and then click **Add**.
4. Verify the entries, and then click **Next**.
5. In **Access Point for Administering the Cluster**, in the **Cluster Name** box, type **Cluster1**.
6. Under **Address**, type **172.16.0.125**, and then click **Next**.
7. In the **Confirmation** dialog box, verify the information, and then click **Next**.
8. On the **Summary** page, click **Finish** to return to the Failover Cluster Manager.

► Task 5: Configure Cluster Shared Volumes

1. On LON-SVR3, in the Failover Cluster Manager console, expand **cluster1.Adatum.com**, and then expand **Storage**, and click **Disk**.
2. In the right pane, locate a disk that is assigned to **Available Storage** (you can see this is in **Assigned To** column). Right-click that disk, and select the **Add to Cluster Shared Volumes** option. (If possible use Cluster Disk 2).
3. Make sure that disk is assigned to **Cluster Shared Volume**.

Results: After this exercise, you will have installed and configured the failover clustering feature.

Exercise 2: Deploying and Configuring a Highly Available File Server

► Task 1: Add the File Server application to the failover cluster

1. On LON-SVR3, in **Server Manager**, click **Dashboard** and then click **Add roles and features**.
2. On the **before your begin** page click **Next**.
3. On the **Select installation type** page click **Next**.
4. On the **Select destination server** page click **Next**.
5. On the **Select server roles** page, expand **File and Storage Services (Installed)**, expand **File and iSCSI services** and select **File Server**.
6. Click **Next** two times.
7. On the **Confirmation** page, click **Install**.
8. When **installation succeeded** message appears click **Close**.
9. Repeat steps 1-8 on LON-SVR4.
10. On LON-SVR3, open the **Failover Cluster Manager**, and then expand **Cluster1.Adatum.com**.
11. Right-click **Roles**, and then click **Configure Role**.
12. On the **Before You Begin** page, click **Next**.
13. On the **Select Role** page, select **File Server**, and then click **Next**.
14. On the **File Server Type** page, click **Scale-Out File Server for application data**, and then click **Next**.

15. On the **Client Access Point** page, in the **Client Access Name** box, type **AdatumFS**, and then click **Next**.
16. On the **Confirmation** page, click **Next**.
17. On the **Summary** page, click **Finish**.

► **Task 2: Add a shared folder to a highly available file server**

1. On LON-SVR3, in the Failover Cluster Manager, click **Roles**, right-click **AdatumFS**, and then click **Add File Share**.
2. In the New Share Wizard, on the **Select the profile for this share** page, click **SMB Share – Quick**, and then click **Next**.
3. On the **Select the server and the path for this share** page, click **Select by volume**, and then click **Next**.
4. On the **Specify share name** page, in the **Share name** box, type **Data**, and then click **Next**.
5. On the **Configure share settings** page, verify that **Enable continuous availability** is selected, and then click **Next**.
6. On the **Specify permissions to control access** page, click **Next**.
7. On the **Confirmation** page, click **Create**.
8. On the **View results** page, click **Close**.

► **Task 3: Configure failover and failback settings**

1. On LON-SVR3, in the Failover Cluster Manager, click **Roles**, right-click **AdatumFS**, and then click **Properties**.
2. Click the **Failover** tab and then click **Allow failback**.
3. Click **Failback between**, and set values to **4** and **5** hours.
4. Click the **General** tab.
5. Select both **LON-SVR3** and **LON-SVR4** as preferred owners.
6. Move **LON-SVR4** up.
7. Click **OK**.

► **Task 4: Validate cluster quorum settings**

1. Open the Failover Cluster Manager console.
2. In the Failover Cluster Manager console, click **Cluster1.Adatum.com**.
3. In the central pane, review the value for **Quorum Configuration**. It should be set to **Node and Disk Majority**.

Results: After this exercise, you will have deployed and configured a highly available file server.

Exercise 3: Validating the Deployment of the Highly Available File Server

► **Task 1: Validate the highly available file server deployment**

1. On LON-DC1, open Windows Explorer, and in the Address bar, type **\\AdatumFS**, and then press Enter.

2. Verify that you can access the location and that you can open the **Data** folder. Create a test text document inside this folder.
3. On LON-SVR3, open the Failover Cluster Manager.
4. Expand **Cluster1.adatum.com**, and then click **Roles**. Note the current owner of AdatumFS.



Note: You can view the owner in the Owner node column. It will be either LON-SVR3 or LON-SVR4.

5. Right-click **AdatumFS**, and then click **Move**, and then click **Select Node**.
6. In the **Move Clustered Role** dialog box, click **OK**.
7. Verify that **AdatumFS** has moved to a new owner.
8. Switch to the LON-DC1 computer and verify that you can still access the **\\AdatumFS** location.

► **Task 2: Validate the failover and quorum configuration for the file server role**

1. On LON-SVR3, in the Failover Cluster Manager, click **Roles**.
2. Verify the current owner for the AdatumFS role.



Note: You can view the owner in the Owner node column, which will be either LON-SVR3 or LON-SVR4.

3. Expand **Nodes**, and then select the node that is the current owner of the AdatumFS role.
4. Right-click the node, click **More Actions**, and then click **Stop Cluster Service**. In the **Stop Cluster Service** dialog box, click **Yes**.
5. Verify that **AdatumFS** has moved to another node. To do this, click the other node, and verify that AdatumFS is running.
6. Switch to the LON-DC1 computer and verify that you can still access the **\\AdatumFS** location.
7. Switch to the LON-SVR3 computer. In the Failover Cluster Manager, right-click the stopped node, click **More Actions**, and then click **Start Cluster Service**.
8. Expand **Storage** and then click **Disks**. In the center pane, right-click the disk that is assigned to **Disk Witness in**.



Note: You can view can view this in the Assigned to column.


9. Click **Take Offline**, and then click **Yes**.
10. Switch to LON-DC1, and verify that you can still access the **\\AdatumFS** location. By doing this, you verify that the cluster is still running, even if the witness disk is offline.
11. Switch to LON-SVR3, and in the Failover Cluster Manager console, click **Storage**, right-click the disk that is in **Offline** status, and then click **Bring Online**.

Results: After this exercise, you will have tested the failover and failback scenarios.

Exercise 4: Configuring Cluster-Aware Updating on the Failover Cluster


► Task 1: Configure Cluster-Aware Updating

1. On LON-DC1, in Server Manager, click **Add roles and features**.
2. In the Add roles and features Wizard, on the **Before you begin** page, click **Next**.
3. On the **Select installation type** page, click **Next**.
4. On the **Select destination server** page, make sure that **Select server from the server pool** is selected, and then click **Next**.
5. On the **Select server roles** page, click **Next**.
6. On the **Select features** page, in the list of features, click **Failover Clustering**. In **Add features that are required for Failover Clustering?** dialog box, click **Add Features**. Click **Next**.
7. On the **Confirm installation selections** page, click **Install**.
8. When installation is complete, click **Close**.
9. On LON-DC1, in the **Server Manager** dashboard, click **Tools**, and then click **Cluster-Aware Updating**.
10. In the Cluster-Aware Updating window, in the **Connect to a failover cluster** drop-down list, select **Cluster1**. Click **Connect**.
11. In the Cluster Actions pane, click **Preview updates for this cluster**.
12. In the Cluster1-Preview Updates window, click **Generate Update Preview List**. After several minutes, updates will be shown in the list. Review updates and then click **Close**.


 **Note:** An Internet connection is required for this step to complete successfully. Make sure that MSL-TMG1 server is up and running and that you can access Internet from LON-DC1.

► Task 2: Update the failover cluster and configure self-updating

1. On LON-DC1, in the Cluster-Aware Updating console, click **Apply updates to this cluster**.
2. On the **Getting Started** page, click **Next**.
3. On the **Advanced options** page, review the options for updating, and then click **Next**.
4. On the **Additional Update Options** page, click **Next**.
5. On the **Confirmation** page, click **Update**, and then click **Close**.
6. In the Cluster nodes pane, you can review the progress of updating.

 **Note:** Remember that one node of the cluster is in Waiting state and the other node is restarting after it is updated.

7. Wait until the process is finished. Process is finished when both nodes have Succeeded in Last Run status column.

 **Note:** This may require a restart of both the nodes.

8. Log on to **LON-SVR3** with the username as **Adatum\Administrator** and password as **Pa\$\$w0rd**.

9. On LON-SVR3, in the **Server Manager**, click **Tools**, and then click **Cluster-Aware Updating**.
10. In the **Cluster-Aware Updating** dialog box, in the **Connect to a failover cluster** drop-down list, select **Cluster1**. Click **Connect**.
11. Click the **Configure cluster self-updating options** in the Cluster Actions pane.
12. On the **Getting Started** page, click **Next**.
13. On the **Add CAU Clustered Role with Self-Updating Enabled** page, click **Add the CAU clustered role, with self-updating mode enabled, to this cluster**, and then click **Next**.
14. On the **Specify self-updating schedule** page, click **Weekly**, in the **Time of day** box, select **4:00 AM**, and then in the **Day of the week** box, select **Sunday**. Click **Next**.
15. On the **Advanced Options** page, click **Next**.
16. On the **Additional Update Options** page, click **Next**.
17. On the **Confirmation** page, click **Apply**.
18. After the clustered role is added successfully, click **Close**.

Results: After this exercise, you will have configured Cluster-Aware Updating on the Failover Cluster.

► **To prepare for the next module**

When you finish the lab, revert the virtual machines to their initial state.

1. On the host computer, start Hyper-V Manager.
2. In the **Virtual Machines** list, right-click **20412A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20412A-LON-SVR1**, **20412A-LON-SVR3**, **20412A-LON-SVR4** and **MSL-TMG1**.

MCT USE ONLY. STUDENT USE PROHIBITED

Module 6: Implementing Failover Clustering with Hyper-V

Lab: Implementing Failover Clustering with Hyper-V

Exercise 1: Configuring Hyper-V Replicas

► Task 1: Boot the physical host machines from VHD

1. Restart the classroom computer, and in the **Windows Boot Manager**, select either **20412A-LON-HOST1** or **20412A-LON-HOST2**.



Note: If you start LON-HOST1, your partner must start LON-HOST2.

2. Log on to the server as **Adatum\Administrator** with password **Pa\$\$w0rd**.
3. On LON-HOST1, from Server Manager, click **Tools**, and then click **Hyper-V manager**.
4. Ensure that virtual machine **20412A-LON-DC1** is running.
5. Repeat steps 3 and 4 on LON-HOST2, and ensure that virtual machine **20412A-LON-SVR1** is running.

► Task 2: Import the LON-CORE virtual machine on LON-HOST1

1. On LON-HOST1, open the Hyper-V Manager console.
2. In the Actions pane, click **Import Virtual Machine**.
3. In the Import Virtual Machine Wizard, on the **Before You Begin** page, click **Next**.
4. On the **Locate Folder** page, click **Browse**.
5. Browse to folder **E:\Program Files\Microsoft Learning\20412\Drives\20412A-LON-CORE**, click **Select Folder**, and then click **Next**.



Note: The drive letter may differ based on the number of drives on the physical host machine.

6. On the **Select Virtual Machine** page, click **20412A-LON-CORE**, and then click **Next**.
7. On the **Choose Import Type** page, click **Next**.
8. On the **Summary** page, click **Finish**.

► Task 3: Configure a replica on both host machines

1. On LON-HOST2, open the Microsoft Hyper-V® Manager console.
2. In Hyper-V Manager, right-click **LON-HOST2**, and then click **Hyper-V Settings**.
3. In Hyper-V Settings for LON-HOST2, click **Replication Configuration**.
4. In the Replication Configuration pane, click **Enable this computer as a Replica server**.
5. In the Authentication and ports section, select **Use Kerberos (HTTP)**.
6. In the Authorization and storage section, click **Allow replication from any authenticated server**, and then click **Browse**.

7. Click **Computer**, double-click **Local Disk (E)**, click **New folder**, in the **Name** text box, type **VMReplica**, and then press Enter.
8. Select the **E:\VMReplica** folder, and then click **Select Folder**.
9. In Hyper-V Settings for LON-HOST2, click **OK**.
10. In the **Settings** dialog box, read the notice, and then click **OK**.
11. Point to the lower left-hand corner of the desktop, and then click **Settings**.
12. Click **Control Panel**, click **System and Security**, and then click **Windows Firewall**.
13. Click **Advanced settings**.
14. In Windows Firewall with Advanced Security, click **Inbound Rules**.
15. In the right pane, in the rule list, find the rule named **Hyper-V Replica HTTP Listener (TCP-In)**. Right-click the rule, and then click **Enable Rule**.
16. Close the Windows Firewall with Advanced Security console, and then close Windows Firewall.
17. Repeat steps 1-16 on LON-HOST1.

► **Task 4: Configure replication for the LON-CORE virtual machine**

1. On LON-HOST1, open Hyper-V Manager, click **LON-HOST1**, right-click **20412A-LON-CORE**, and then click **Enable Replication**.
2. On the **Before You Begin** page, click **Next**.
3. On the **Specify Replica Server** page, click **Browse**.
4. In the Select Computer window, type **LON-HOST2**, click **Check Names**, click **OK**, and then click **Next**.
5. On the **Specify Connection Parameters** page, review settings, ensure that **Use Kerberos authentication (HTTP)** is selected, and then click **Next**.
6. On the **Choose Replication VHDs** page, ensure that **20412A-LON-CORE.vhd** is selected, and then click **Next**.
7. On the **Configure Recovery History** page, select **Only the latest recovery point**, and then click **Next**.
8. On the **Choose Initial Replication Method** page, click **Send initial copy over the network**, select **Start replication immediately**, and then click **Next**.
9. On the **Completing the Enable Replication Wizard** page, click **Finish**.
10. Wait 10-15 minutes. In the Hyper-V Manager console, you can monitor the progress of the initial replication in the **Status** column.
11. When replication completes, ensure that **20412A-LON-CORE** appears on **LON-HOST2** in Hyper-V Manager.

► **Task 5: Validate a planned failover to the replica site**

1. On LON-HOST2, in Hyper-V Manager, right-click **20412A-LON-CORE**, select **Replication**, and then click **View Replication Health**.
2. Review the content in the window that appears, ensure that there are no errors, and then click **Close**.
3. On LON-HOST1, open Hyper-V Manager, and verify that 20412A-LON-CORE is turned off.
4. Right-click **20412A-LON-CORE**, select **Replication**, and then click **Planned Failover**.

5. In the Planned Failover window, ensure that **Start the Replica virtual machine after failover** is selected, and then click **Fail Over**.
6. In the Planned Failover window, click **Close**.
7. On LON-HOST2, in Hyper-V Manager, ensure that 20412A-LON-CORE is running.
8. On LON-HOST1, right-click **20412A-LON-CORE**, point to **Replication**, and then click **Remove replication**.
9. In the **Remove replication** dialog box, click **Remove Replication**.
10. On LON-HOST2, right-click **20412A-LON-CORE**, and then click **Shut Down**.
11. In the **Shut Down Machine** dialog box, click **Shut Down**.

Results: After completing this exercise, you will have configured a Hyper-V Replica.

Exercise 2: Configuring a Failover Cluster for Hyper-V

► Task 1: Connect to the iSCSI target from both host machines

1. On LON-HOST1, from the taskbar, click the **Server Manager** icon to open Server Manager, click **Tools**, click **iSCSI Initiator**, and then at the **Microsoft iSCSI** prompt, click **Yes**.
2. Click the **Discovery** tab.
3. Click **Discover Portal**.
4. In the **IP address or DNS name** box, type **172.16.0.21**, and then click **OK**.
5. Click the **Targets** tab, and then click **Refresh**.
6. In the **Targets** list, select **iqn.1991-05.com.microsoft:lon-svr1-target1-target**, and then click **Connect**.
7. Select **Add this connection to the list of Favorite Targets**, and then click **OK** twice.
8. On LON-HOST2, from the taskbar, click the **Server Manager** icon to open Server Manager, click **Tools**, and then click **iSCSI Initiator**.
9. In the **Microsoft iSCSI** dialog box, click **Yes**.
10. Click the **Discovery** tab, and then click **Discover Portal**.
11. In the **IP address or DNS name** box, type **172.16.0.21**, and then click **OK**.
12. Click the **Targets** tab, and then click **Refresh**.
13. In the **Discovered targets** list, select **iqn.1991-05.com.microsoft:lon-svr1-target1-target**, and then click **Connect**.
14. Select **Add this connection to the list of Favorite Targets**, and then click **OK** twice.
15. On LON-HOST2, in Server Manager, click **Tools**, and then click **Computer Management**.
16. Expand **Storage**, click **Disk Management**, right-click **Disk 2**, and then click **Online**.
17. Right-click **Disk 2** again, and then click **Initialize Disk**.
18. In the **Initialize Disk** dialog box, click **OK**.
19. Right-click the unallocated space next to Disk 2, and then click **New Simple Volume**.
20. On the **Welcome** page, click **Next**.

21. On the **Specify Volume Size** page, click **Next**.
22. On the **Assign Drive Letter or Path** page, click **Next**.
23. On the **Format Partition** page, in the **Volume label** box, type **ClusterDisk**, select the **Perform a quick format** check box, and then click **Next**.
24. Click **Finish**.
25. Repeat steps 17 through 24 for Disk 3 and Disk 4. In step 23, provide the name **ClusterVMs** for Disk 3, and the name **Quorum** for Disk 4.
26. On LON-HOST1, in Server Manager, click **Tools**, and then click **Computer Management**.
27. Expand **Storage**, and then click **Disk Management**.
28. Right-click **Disk Management**, and then click **Refresh**.
29. Right-click **Disk 2**, and then click **Online**.
30. Right-click **Disk 3**, and then click **Online**.
31. Right-click **Disk 4**, and then click **Online**.

► **Task 2: Configure failover clustering on both host machines**

1. On LON-HOST1, on the taskbar, click the **Server Manager** icon to open Server Manager.
2. From the Dashboard, click **Add roles and features**.
3. On the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, click **Next**.
5. On the **Select destination server** page, ensure that **Select server from the server pool** is selected, and then click **Next**.
6. On the **Select server roles** page, click **Next**.
7. On the **Select features** page, in the **Features** list, click **Failover Clustering**.
8. At the **Add features that are required for failover clustering** prompt, click **Add Features**, and then click **Next**.
9. On the **Confirm installation selections** page, click **Install**.
10. When installation completes, click **Close**.
11. Repeat steps 1 through 10 on LON-HOST2.
12. On LON-HOST1, in Server Manager, click **Tools**, and then click **Failover Cluster Manager**.
13. In the Failover Cluster Manager, in the center pane, under **Management**, click **Create Cluster**.
14. In the Create Cluster Wizard, on the **Before You Begin** page, read the information, and then click **Next**.
15. On the **Select Servers** page, in the **Enter server name** box, type **LON-HOST1**, and then click **Add**.
16. In the **Enter server name** box, type **LON-HOST2**, and then click **Add**.
17. Verify the entries, and then click **Next**.
18. On the **Validation Warning** page, click **No. I don't require support from Microsoft for this cluster**, and then click **Next**.
19. In the **Access Point for Administering the Cluster** page, in the **Cluster Name** box, type **VMCluster**.
20. In the **IP address name** box, under **Address**, type **172.16.0.126**, and then click **Next**.

21. In the **Confirmation** dialog box, verify the information, clear the **Add all eligible storage to the cluster** check box, and then click **Next**.
 22. On the **Summary** page, and then click **Finish**.
- **Task 3: Configure disks for the failover cluster**
1. On LON-HOST1, in the Failover Cluster Manager, expand **VMCluster.Adatum.com**, expand **Storage**, right-click **Disks**, and then click **Add Disk**.
 2. In the **Add Disks to a Cluster** dialog box, verify that all disks are selected, and then click **OK**.
 3. In the Failover Cluster Manager, verify that all disks appear available for cluster storage.
 4. Select the **ClusterVMs** disk, right-click **ClusterVMs**, and then select **Add to Cluster Shared Volumes**.
 5. Right-click **VMCluster.adatum.com**, select **More Actions**, click **Configure Cluster Quorum Settings**, and then click **Next**.
 6. On the **Select Quorum Configuration Option** page, click **Use typical settings**, and then click **Next**.
 7. On the **Confirmation** page, click **Next**.
 8. On the **Summary** page, click **Finish**.

Results: After completing this exercise, you will have configured a failover cluster for Hyper-V.

Exercise 3: Configuring a Highly Available Virtual Machine

► **Task 1: Move virtual machine storage to the iSCSI target**

1. In the Failover Cluster Manager, verify that **LON-HOST1** is the owner of the **ClusterVMs** disk. If it is not, move the ClusterVMs disk to LON-HOST1.
2. On LON-HOST1, open a Windows® Explorer window, and browse to **E:\Program Files\Microsoft Learning\20412\Drives\20412A-LON-CORE\Virtual Hard Disks**.
3. Move the **20412A-LON-CORE.vhd** virtual hard drive file to the **C:\ClusterStorage\Volume1** location.

► **Task 2: Configure the virtual machine as highly available**

1. On LON-HOST1, in the Failover Cluster Manager, click **Roles** and then in the Actions pane, click **Virtual Machines**.
2. Click **New Virtual Machine**.
3. Select **LON-HOST1**, and then click **OK**.
4. In the New Virtual Machine Wizard, click **Next**.
5. On the **Specify Name and Location** page, in the **Name** text box, type **TestClusterVM**, click **Store the virtual machine in a different location**, and then click **Browse**.
6. Browse to and select **C:\ClusterStorage\Volume1**, click **Select Folder**, and then click **Next**.
7. On the **Assign Memory** page, type **1536**, and then click **Next**.
8. On the **Configure Networking** page, click select **External Network**, and then click **Next**.
9. On the **Connect Virtual Hard Disk** page, click **Use an existing virtual hard disk**, and then click **Browse**.

10. Browse to **C:\ClusterStorage\Volume1**, click **20412A-LON-CORE.vhd**, and then click **Open**.
11. Click **Next**, and then click **Finish**.
12. In the High Availability Wizard, on the **Summary** page, click **Finish**.
13. In Failover Cluster Manager, from the Roles node, right-click the **TestClusterVM**, and then click **Start**.
14. Ensure that the machine starts successfully.

► **Task 3: Perform live migration for the virtual machine**

1. On LON-HOST1, open the Failover Cluster Manager., expand **VMCluster.Adatum.com**, and then click **Roles**.
2. Right-click **TestClusterVM**, click **Move**, click **Live Migration**, and then click **Select Node**.
3. Click **LON-HOST2**, and then click **OK**.
4. Right-click **TestClusterVM**, and then click **Connect**.
5. Ensure that you can access and operate the virtual machine while it is migrating to another host.
6. Wait until migration completes.

► **Task 4: Perform storage migration for the virtual machine**

1. On LON-HOST2, open Hyper-V Manager.
2. In the central pane, click **20412A-LON-SVR1-B**.
3. In the Actions pane, click **Move**.
4. On the **Before You Begin** page, click **Next**.
5. On the **Choose Move Type** page, click **Move the virtual machine's storage**, and then click **Next**.
6. On the **Choose Options for Moving Storage** page, click **Move all of the virtual machine's data to a single location**, and then click **Next**.
7. On the **Choose a new location for virtual machine** page, click **Browse**.
8. Browse to **C:**, create a new folder called **LON-SVR1**, click **Select Folder**, and then click **Next**.
9. On the **Summary** page, click **Finish**. While the virtual machine is migrating, connect to it and verify that it is fully operational.
10. After the move process completes, click **Close**.
11. Shut down all running virtual machines.

Results: After completing this exercise, you will have configured a highly available virtual machine.

► **To prepare for next module**

1. Restart **LON-HOST1**.
2. When you are prompted with the boot menu, select **Windows Server 2008 R2**, and then press Enter.
3. Log on to the host machine as directed by your instructor.
4. Repeat steps 1-3 on **LON-HOST2**.

Module 7: Implementing Disaster Recovery

Lab: Implementing Windows Server Backup and Restore

Exercise 1: Backing Up Data on a Windows Server 2012 Server

► Task 1: Install Windows Server Backup

1. Switch on LON-SVR1.
2. In Server Manager, in the Welcome pane, click **Add roles and features**.
3. In the Add Roles and Features Wizard, on the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, click **Next**.
5. On the **Select destination server** page, click **Next**.
6. On the **Select server roles** page, click **Next**.
7. On the **Select features** page, select **Windows Server Backup**, and then click **Next**.
8. On the **Confirm installation selections** page, click **Install**.
9. On the **Installation progress** page, wait until the **Installation succeeded on LON-SVR1.adatum.com** message displays, and then click **Close**.

► Task 2: Configure a scheduled backup

1. On LON-SVR1, in Server Manager, click **Tools**, and then click **Windows Server Backup**.
2. In the navigation pane, click **Local Backup**.
3. Click Backup Schedule.
4. In the Backup Schedule Wizard, on the **Getting Started** page, click **Next**.
5. On the **Select Backup Configuration** page, click **Full server (recommended)**, and then click **Next**.
6. On the **Specify Backup Time** page, next to **Select time of day**, select **1:00 AM**, and then click **Next**.
7. On the **Specify Destination Type** page, click **Backup to a shared network folder**, and then click **Next**. Review the warning, and then click **OK**.
8. On the **Specify Remote Shared Folder** page, in the **Location** text box, type **\\LON-DC1\Backup**, and then click **Next**.
9. In the **Register Backup Schedule** dialog box, in the **Username** text box, type **Administrator**, in the **Password** text box, type **Pa\$\$w0rd**, and then click **OK**.
10. Click **Finish**, and then click **Close**.

► Task 3: Complete an on-demand backup

1. On LON-SVR1, in Server Manager, click **Tools**, and then click **Windows Server Backup**.
2. In the Actions pane, click **Backup Once**.
3. In the Backup Once Wizard, on the **Backup Options** page, click **Different options**, and then click **Next**.
4. On the **Select Backup Configuration** page, click **Custom**, and then click **Next**.
5. On the **Select Items for Backup** page, click **Add Items**.

6. Expand **Local disk (C:)**, select the **Financial Data** check box, click **OK**, and then click **Next**.
7. On the **Specify Destination Type** page, click **Remote shared folder**, and then click **Next**.
8. On the **Specify Remote Folder** page, type `\\LON-DC1\Backup`, and then click **Next**.
9. On the **Confirmation** page, click **Backup**.
10. On the **Backup Progress** page, after the backup is complete, click **Close**.

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

Exercise 2: Restoring Files Using Windows Server Backup

► Task 1: Delete a file from the server

1. On LON-SVR1, on the task bar, click **Windows Explorer**.
2. In Windows Explorer, browse to **Local Disk (C:)**, right-click **Financial Data**, and then click **Delete**.

► Task 2: Restore a file from backup

1. In the Windows Server Backup console, in the Actions pane, click **Recover**.
2. On the **Getting Started** page, click **A backup stored on another location**, and then click **Next**.
3. On the **Specify Location Type** page, click **Remote shared folder**, and then click **Next**.
4. On the **Specify Remote Folder** page, type `\\LON-DC1\Backup`, and then click **Next**.
5. On the **Select Backup Date** page, click **Next**.
6. On the **Select Recovery Type** page, click **Next**.
7. On the **Select Items to Recover** page, expand **LON-SVR1**, click **Local Disk (C:)**, and on the right pane, select **Financial Data**, and then click **Next**.
8. On the **Specify Recovery Options** page, under **Another Location**, type `C:\`, and then click **Next**.
9. On the **Confirmation** page, click **Recover**.
10. On the **Recovery Progress** page, click **Close**.
11. Open drive `C:\` and ensure that the **Financial Data** folder is restored.

Results: After completing this exercise, you will have tested and validated the procedure for restoring a file from backup

Exercise 3: Implementing Microsoft Online Backup and Restore

► Task 1: Install the Microsoft Online Backup Service component


1. On LON-SVR1, on the taskbar, click the **Windows Explorer** icon.
2. In **Allfiles (E:)**, in the details pane, double-click **OBSInstaller.exe**, and then click **Run**.
3. In the **Microsoft Online Service Pre-Release Agreement** dialog box, select **I accept the Service Agreement terms and conditions**, and then click **OK**.
4. On the **Prerequisites Check** page, click **Next**.

5. On the **Installation Settings** page, specify the settings (if not default), and then click **Next**:
 - Installation Folder: **C:\Program Files**
 - Cache Location: **C:\Program Files\Microsoft Online Backup Service Agent**
6. On the **Microsoft Update Opt-In** page, select **I don't want to use Microsoft Update**, and then click **Install**.
7. On the **Installation** page, ensure that the **Microsoft Online Backup Service Agent installation has completed successfully** message displays, clear the **Check for newer updates** check box, and then click **Finish**.
8. On LON-SVR1, click **Start**, and then click **Microsoft Online Backup Service**.
9. On LON-SVR1, click **Start**, and then click **Microsoft Online Backup Service Shell**.

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

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

1. In the Server Manager window, on the **Welcome to Server Manager** page, click **1**. Configure this local server.
2. In the Server Manager window, on the **Local Server** page, click **LON-SVR1**.
3. In the System Properties window, click **Change**, in the **Computer Name** box, type *YOURCITYNAME-YOURNAME*, click **OK** twice, and then click **Close**.
4. In a window that displays the message that you should restart your computer, click **Restart Now**.
5. Wait until *YOURCITYNAME-YOURNAME* has restarted, and then log on as **Adatum\Administrator** with the password **Pa\$\$w0rd**.
6. Start the Microsoft Online Backup Service console, and then click **Register Server**.
7. In the Register Server Wizard, on the **Account Credentials** page, in the **Username** box, type **holuser@onlinebackupservice.onmicrosoft.com**, in the **Password** box, type **Pa\$\$w0rd**, and then click **Next**.

 **Note:** In a real-life scenario, you would type the username and password of your Microsoft Online Backup Service subscription account.

8. On the **Proxy Configuration** page, click **Next**.
9. On the **Encryption Settings** page, in the **Enter passphrase** and **Confirm passphrase** boxes, type **Pa\$\$w0rdPa\$\$w0rd**, and then click **Register**.
10. On the **Server Registration** page, ensure that the **Microsoft Online Backup Service is now available for this server** message displays, and then click **Close**.

► Task 3: Configure an online backup and start a backup

1. Switch to the Microsoft Online Backup Service console, and then click **Schedule Backup**.
2. On the **Getting started** page, click **Next**.
3. On the **Select Items to back up** page, click **Add Items**.
4. In the **Select Items** dialog box, expand **C:**, select **Financial Data**, click **OK**, and then click **Next**.

5. On the **Specify Backup Time** page, select **Saturday**, click **1:00 AM**, click **Add**, and then click **Next**.
6. On the **Specify Retention Setting** page, accept the default settings, and then click **Next**.
7. On the **Confirmation** page, click **Finish**.
8. On the **Modify Backup Progress** page, click **Close**.
9. In the Microsoft Online Backup Service console, click **Back Up Now**.
10. In the Back Up Now Wizard, on the **Confirmation** page, click **Back Up**.
11. On the **Backup progress** page, wait until **Backup is successfully completed** message displays, and then click **Close**.

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

1. On LON-SVR1, on the taskbar, click the **Windows Explorer** icon, and then in the Windows Explorer navigation pane, click **Local Disk (C:)**.
2. In the Local Disk (C:) window, right-click **Financial Data**, and then click **Delete**.
3. Switch to the Microsoft Online Backup Service console, and then click **Recover Data**.
4. In the Recover Data Wizard, on the **Getting Started** page, select **This server**, and then click **Next**.
5. On the **Select Recovery Mode** page, select **Browse for files**, and then click **Next**.
6. On the **Select Volume and Date** page, in the **Select the volume** drop-down list box, select **C:**. In the calendar, click the date when you performed the backup, in the **Time** drop-down list, click the time when you performed backup, and then click **Next**.
7. On the **Select Items to Recover** page, expand **C:**, click the **Financial Data** folder, and then click **Next**.
8. On the **Specify Recovery Options** page, select **Original location** and **Create copies so that you have both versions**, and then click **Next**.
9. On the **Confirmation** page, click **Recover**.
10. On the **Recovery Progress** page, ensure that **File(s) recovery job succeeded** status message displays, and then click **Close**.
11. In Windows Explorer, expand drive **C:**, and ensure that the **Financial Data** folder is restored to drive C.

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

1. Switch to the Microsoft Online Backup Service console, and then click **Unregister Server**.
2. On the **Getting started** page, click **Unregister this server**, and then click **Next**.
3. On the **Account Credentials** page, provide the following credentials:
 - Username: **holuser@onlinebackupservice.onmicrosoft.com**
 - Password: **Pa\$\$w0rd**
4. Click **Unregister**.
5. On the **Server Unregistration** page, click **Close**.

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

► **To prepare for the next module**

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20412A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20412A-LON-SVR1**, and **MSL-TMG1**.

MCT USE ONLY. STUDENT USE PROHIBITED

Module 8: Implementing Distributed Active Directory Domain Services Deployments

Lab: Implementing Complex AD DS Deployments

Exercise 1: Implementing Child Domains in AD DS

► Task 1: Configure Domain Name System (DNS) for domain delegation

1. On LON-DC1, in Server Manager, click **Tools**, and then click **DNS**.
2. In DNS Manager, expand **LON-DC1**, expand **Forward Lookup Zones**, select and then right-click **Adatum.com**, and then click **New Delegation**.
3. In the New Delegation Wizard, click **Next**, in the **Delegated domain** text box, type **na**, and then click **Next**.
4. In the **Name Servers** box, click **Add**.
5. In the **Server fully qualified domain name (FQDN)** text box, type **TOR-DC1.adatum.com**, clear **<Click here to add an IP Address>**, type **172.16.0.25**, and then click **OK**.
6. In the Name Servers window, click **Next**.
7. In the Complete the New Delegation Wizard window, click **Finish**.

► Task 2: Install a domain controller in a child domain

1. On TOR-DC1, in Server Manager, click **Manage**, and from the drop-down list box, click **Add Roles and Features**.
2. On the **Before you begin** page, click **Next**.
3. On the **Select installation type** page, confirm that **Role-based or feature-based installation** is selected, and then click **Next**.
4. On the **Select destination server** page, ensure that **Select a server from the server pool** is selected, and that **TOR-DC1.adatum.com** is highlighted, and then click **Next**.
5. On the **Select server roles** page, click **Active Directory Domain Services**.
6. On the **Add features that are required for Active Directory Domain Services?** page, click **Add Features**.
7. On the **Select server roles** page, click **Next**.
8. On the **Select features** page, click **Next**.
9. On the **Active Directory Domain Services** page, click **Next**.
10. On the **Confirm installation selections** page, click **Install**. (This may take a few minutes to complete.)
11. When the Active Directory® Domain Services (AD DS) binaries have installed, click the blue **Promote this server to a domain controller** link.
12. In the Deployment Configuration window, click **Add a new domain to an existing forest**.
13. Verify that **Select domain type** is set to **Child Domain**, and that **Parent** domain name is set to **Adatum.com**. In the **New domain name** text box, type **na**.

14. Confirm that **Supply the credentials to perform this operation** is set to **ADATUM\Administrator (Current user)**, and then click **Next**. (If this is not the case, then use the **Change** button to enter the credentials **Adatum\Administrator**, and the password **Pa\$\$w0rd**).
15. In the Domain Controller Options window, ensure that **Domain functional level** is set to **Windows Server 2012 Release Candidate**.
16. Ensure that both the **Domain Name system (DNS) server** and **Global Catalog (GC)** check boxes are selected.
17. Confirm that **Site name:** is set to **Default-First-Site-Name**.
18. Under **Type the Directory Services Restore Mode (DSRM) password**, type **Pa\$\$w0rd** in both text boxes and then click **Next**.
19. On the **DNS Options** page, click **Next**.
20. On the **Additional Options** page, click **Next**.
21. On the Paths window, click **Next**.
22. On the Review Options window, click **Next**.
23. On the Prerequisites Check window, confirm that there are no issues, and then click **Install**.

► **Task 3: Verify the default trust configuration**

1. Log on to **TOR-DC1** as **NA\Administrator** using the password **Pa\$\$w0rd**.
2. When Server Manager opens, click **Local Server**. Verify that **Windows Firewall** shows **Domain: On**. If it does not, then next to **Local Area Connection** click **172.16.0.25, IPv6 enabled**. Right-click **Local Area Connection** and then click **Disable**. Right-click **Local Area Connection** and then click **Enable**. The Local Area Connection should now show **Adatum.com**.
3. In Server Manager, from the **Tools** menu, click **Active Directory Domains and Trusts**.
4. In the Active Directory Domains and Trusts console, expand **Adatum.com**, right-click **na.adatum.com**, and then click **Properties**.
5. Select the **Trusts** tab.
6. In the **Domain trusted by this domain (outgoing trusts)** box, click **Adatum.com**, and then click **Properties**.
7. In the Adatum.com Properties window, click **Validate** and select **Yes, validate the incoming trust**.
8. In the **User name** text box, type **administrator**, and in the **Password** text box, type **Pa\$\$w0rd**, and then click **OK**.
9. A message will display: The trust has been validated. It is in place and active.

Note: If you receive a message that the trust cannot be validated, or that the secure channel (SC) verification has failed, ensure that you have completed step 2 and then wait for at least 10-15 minutes. You can continue with the lab and come back later to verify this step.

10. Click **OK**.
11. Click **OK** twice to close the Adatum.com Properties dialog box.

Results: After completing this exercise, you will have implemented child domains in AD DS.

Exercise 2: Implementing Forest Trusts

► Task 1: Configure stub zones for DNS name resolution

1. On LON-DC1, in Server Manager, click the **Tools** menu, and then from the drop-down menu, click **DNS**.
2. In the DNS tree pane, expand **LON-DC1**, right-click **Forward Lookup Zones**, and then click **New Zone**.
3. In the New Zone Wizard, click **Next**.
4. On the Zone Type window, click **Stub zone**, and then click **Next**.
5. On the Active Directory Zone Replication Scope window, click **To all DNS servers running on domain controllers in this forest: adatum.com**, and then click **Next**.
6. In the **Zone name:** text box, type **treyresearch.net**, and then click **Next**.
7. On the Master DNS Servers window, click **<Click here to add an IP Address or DNS Name>**, type **172.16.10.10**, click on the free space, and then click **Next**.
8. On the Completing the New Zone Wizard window, click **Next** and then **Finish**.
9. Select and then right-click the new stub zone **treyresearch.net** and then click **Transfer from Master**.
10. Right-click **treyresearch.net** and then click **Refresh**.
11. Confirm that the treyresearch.net stub zone has some records.
12. Switch to MUN-DC1.
13. In Server Manager, click the **Tools** menu, and from the drop-down menu, click **DNS**.
14. In the tree pane, expand **MUN-DC1**, select and then right-click **Forward Lookup Zones**, and then click **New Zone**.
15. In the New Zone Wizard, click **Next**.
16. On the Zone Type window, click **Stub zone**, and then click **Next**.
17. In the Active Directory Zone Replication Scope window, select **To all DNS servers running on domain controllers in this forest: Treyresearch.net** and then click **Next**.
18. In the **Zone name:** text box, type **adatum.com**, and then click **Next**.
19. In the Master DNS Servers window, click **<Click here to add an IP Address or DNS Name>**, type **172.16.0.10**, click on the free space, and then click **Next**.
20. In the Completing the New Zone Wizard window, click **Next** and then click **Finish**.
21. Select and then right-click the new stub zone **adatum.com**, and then click **Transfer from Master**.
22. Right-click **adatum.com**, and then click **Refresh**.
23. Confirm that the adatum.com stub zone has some records.
24. Close DNS Manager.

► Task 2: Configure a forest trust with selective authentication

1. On LON-DC1, from the **Tools** menu, click **Active Directory Domain and Trusts**.
2. In the Active Directory Domains and Trusts management console window, right-click **Adatum.com**, and then click **Properties**.
3. In the Adatum.com Properties window, click the **Trusts** tab, and then click **New Trust**.

4. On the New Trust Wizard window, click **Next**.
5. In the **Name** text box, type **tresearch.net** and then click **Next**.
6. In the Trust Type window, select **Forest trust** and click **Next**.
7. In the Direction of Trust window, select **One-way: outgoing**, and then click **Next**.
8. In the Sides of Trust window, select **Both this domain and the specified domain** and then click **Next**.
9. In the User Name and Password window type **Administrator** as the user name and **Pa\$\$w0rd** as the password in the appropriate boxes, and then click **Next**.
10. In the Outgoing Trust Authentication Level-Local Forest window, select **Selective authentication**, and then click **Next**.
11. In the **Trust Selections Complete** page, click **Next**.
12. On the **Trust Creation Complete** page, click **Next**.
13. On the **Confirm Outgoing Trust** page, click **Next**.
14. Click **Finish**.
15. In the Adatum.com Properties window, click the **Trusts** tab.
16. On the **Trusts** tab, under **Domains trusted by this domain (outgoing trusts)**, select **tresearch.net** and click **Properties**.
17. In the tresearch.net Properties window, click **Validate**.
18. Review the message that appears: The trust has been validated. It is in place and active.
19. Click **OK**, and then click **No** at the prompt.
20. Click **OK** twice.
21. Close Active Directory Domains and Trusts.

► **Task 3: Configure a server for selective authentication**

1. On LON-DC1, in Server Manager, from the **Tools** menu, click **Active Directory Users and Computers**.
2. In the Active Directory Users and Computers console, from the **View** menu, click **Advanced Features**.
3. Expand **Adatum.com**, and then click **Computers**.
4. Right-click **LON-SVR1** and then click **Properties**.
5. Click the **Security** tab, and then click **Add**.
6. On the Select Users, Computers, Service Accounts, or Groups window, click **Locations**.
7. Click **tresearch.net** and then click **OK**.
8. In the **Enter the object name to select (examples:)** text box, type **tresearch\it**, and then click **Check Names**. When prompted for credentials, type **tresearch\administrator** with the password of **Pa\$\$w0rd**. Click **OK**.
9. On the Select Users, Computers, Service Accounts, or Groups window, click **OK**.
10. In the LON-SVR1 Properties window, ensure that **tresearch\it** is highlighted, and select the **Allow** checkbox that is in line with **Allowed to authenticate**.
11. Click **OK**.

12. Switch to LON-SVR1.
13. On the taskbar, click **Windows Explorer**.
14. Click Local Disk (C).
15. Right-click in the details pane, click **New**, and then click **Folder**.
16. In the **Name** text box, type **IT-Data**, and then press Enter.
17. Right-click **IT-Data**, and then click **Properties**.
18. In the IT-Data Properties window, click the **Security** tab, and then click **Edit**.
19. On the Permission for IT-Data window, click **Add**.
20. In the **Enter the object names to select (examples:)** text box, type **tresearch\it**, and then click **Check Names**. When the name resolves, click **OK**. If you are prompted for credentials, type **Tresearch\administrator** with the password of **Pa\$\$w0rd**. Click **OK** twice.
21. In the Permissions for IT-Data window, select **tresearch\it**, click the **Allow** that is opposite the **Modify** permission and then click **OK**.
22. Click the **Sharing** tab, select **Advanced Sharing**, and then click **Share this folder**.
23. Click **Permissions**, confirming that **Everyone** is highlighted, and then click **Full Control**.
24. Click **OK** twice, and then click **Close**.
25. Log off of MUN-DC1.
26. Log on to **MUN-DC1** as **tresearch\Alice** with the password **Pa\$\$w0rd**.
27. Hover your pointer in the lower-right corner of the desktop, and when the sidebar displays, click **Search**.
28. In the **Search** text box, type **\\LON-SVR1\IT-Data**. The folder will open.

Results: After completing this exercise, you will have implemented forest trusts.

► To prepare for the next module

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

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

MCT USE ONLY. STUDENT USE PROHIBITED

Module 9: Implementing Active Directory Domain Services Sites and Replication

Lab: Implementing AD DS Sites and Replication

Exercise 1: Modifying the Default Site

► Task 1: Install the Toronto domain controller

1. On TOR-DC1, click **Manage**, and from the drop-down list box, click **Add Roles and Features**.
2. On the **Before you begin** page, click **Next**.
3. On the **Select installation type** page, confirm that **Role-based or feature-based installation** is selected, and then click **Next**.
4. On the **Select destination server** page, ensure that **Select a server from the server pool** is selected, and that **TOR-DC1.adatum.com** is highlighted, and then click **Next**.
5. On the **Select server roles** page, click **Active Directory Domain Services**.
6. On the **Add features that are required for Active Directory Domain Services?** page, click **Add Features**. Click **Next**.
7. On the **Select features** page, click **Next**.
8. On the **Active Directory Domain Services** page, click **Next**.
9. On the **Confirm installation selections** page, click **Install**. (This may take a few minutes to complete.)
10. When the AD DS binaries have installed, click the blue **Promote this server to a domain controller** link.
11. In the Deployment Configuration window, click **Add a domain controller to an existing domain**. Click **Next**.
12. In the Domain Controller Options window, ensure that both the **Domain Name system (DNS) server** and **Global Catalog (GC)** check boxes are selected.
13. Confirm that **Site** name: is set to **Default-First-Site-Name**, and then under **Type the Directory Services Restore Mode (DSRM) password**, type **Pa\$\$w0rd** in both the **Password** and **Confirm password** boxes. Click **Next**.
14. On the **DNS Options** page, click **Next**.
15. On the **Additional Options** page, click **Next**.
16. On the Paths window, click **Next**.
17. On the Review Options window, click **Next**.
18. On the Prerequisites Check window, confirm that there are no issues, and then click **Install**. The server will automatically restart.
19. After TOR-DC1 restarts, log on as **Adatum\Administrator** with the password **Pa\$\$w0rd**.

► Task 2: Rename the default site

1. If necessary, on LON-DC1 open the Server Manager console.
2. In Server Manager, click **Tools**, and then click **Active Directory Sites and Services**.
3. In the Active Directory Sites and Services console, in the navigation pane, expand **Sites**.
4. Right-click **Default-First-Site-Name**, and then click **Rename**.
5. Type **LondonHQ**, and then press Enter.
6. Expand **LondonHQ**. expand the **Servers** folder, and then verify that both **LON-DC1** and **TOR-DC1** belong to the **LondonHQ** site.

► Task 3: Configure IP subnets associated with the default site

1. On LON-DC1, in the Active Directory Sites and Services console, in the navigation pane, expand **Sites**, and then click the **Subnets** folder.
2. Right-click **Subnets**, and then click **New Subnet**.
3. In the **New Object – Subnet** dialog box, under **Prefix**, type **172.16.0.0/24**.
4. Under **Select a site object for this prefix**, select **LondonHQ**, and then click **OK**.

Results: After completing this exercise, you will have reconfigured the default site and assigned IP address subnets to the site.

Exercise 2: Creating Additional Sites and Subnets

► Task 1: Create the AD DS sites for Toronto

1. On LON-DC1, in the Active Directory Sites and Services console, in the navigation pane, right-click **Sites**, and then click **New Site**.
2. In the **New Object – Site** dialog box, next to **Name**, type **Toronto**.
3. Under **Select a site link object for this site**, select **DEFAULTIPSITELINK**, and then click **OK**.
4. In the **Active Directory Domain Services** dialog box, click **OK**. The Toronto site displays in the navigation pane.
5. In the Active Directory Sites and Services console, in the navigation pane, right-click **Sites**, and then click **New Site**.
6. In the **New Object – Site** dialog box, next to **Name**, type **TestSite**.
7. Under **Select a site link object for this site**, select **DEFAULTIPSITELINK**, and then click **OK**. The TestSite site displays in the navigation pane.

► Task 2: Create IP subnets associated with the Toronto sites

1. On LON-DC1, in the Active Directory Sites and Services console, in the navigation pane, expand **Sites**, and then click the **Subnets** folder.
2. Right-click **Subnets**, and then click **New Subnet**.
3. In the **New Object – Subnet** dialog box, under **Prefix**, type **172.16.1.0/24**.
4. Under **Select a site object for this prefix**, select **Toronto**, and then click **OK**.
5. Right-click **Subnets**, and then click **New Subnet**.

6. In the **New Object – Subnet** dialog box, under **Prefix**, type **172.16.100.0/24**.
7. Under **Select a site object for this prefix**, select **TestSite**, and then click **OK**.
8. In the navigation pane, click the **Subnets** folder. Verify the three subnets were created and associated with their appropriate site as displayed in the details pane.

Results: After this exercise, you will have created two additional sites representing the IP subnet addresses located in Toronto.

Exercise 3: Configuring AD DS Replication

► Task 1: Configure site links between AD DS sites

1. On LON-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 **IP**, and then click **New Site Link**.
3. In the **New Object – Site Link** dialog box, next to **Name**, type **TOR-TEST**.
4. Under **Sites not in this site link**, select **Toronto**, select **TestSite**, click **Add**, and then click **OK**.
5. Right-click **TOR-TEST**, and then click **Properties**.
6. In the **TOR-TEST Properties** dialog box, click **Change Schedule**.
7. In the **Schedule for TOR-TEST** dialog box, highlight the range from **Monday 9am to Friday 3pm**.
8. Select **Replication Not Available**, and then click **OK**.
9. Click **OK** to close TOR-TEST Properties.
10. Right-click **DEFAULTIPSITELINK**, and then click **Rename**.
11. Type **LON-TOR**, and then press Enter.
12. Right-click **LON-TOR**, and then click **Properties**.
13. Under **Sites in this link**, click **TestSite**, and then click **Remove**.
14. Next to **Replicate Every**, change the value to **60** minutes, and then click **OK**.

► Task 2: Move TOR-DC1 to the Toronto site

1. On LON-DC1, in the Active Directory Sites and Services console, in the navigation pane, expand **Sites**, expand **LondonHQ**, and then expand the **Servers** folder.
2. Right-click **TOR-DC1**, and then click **Move**.
3. In the **Move Server** dialog box, click **Toronto**, and then click **OK**.
4. In the navigation pane, expand the **Toronto** site, expand **Servers**, and then click **TOR-DC1**.

► Task 3: Monitor AD DS site replication

1. On LON-DC1, on the taskbar, click the **Windows PowerShell** button.
2. At the command prompt, type the following, and then press Enter:

```
Repadmin /kcc
```

This command recalculates the inbound replication topology for the server.

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

```
Repadmin /showrep1
```

Verify that the last replication with TOR-DC1 was successful.

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

```
Repadmin /bridgeheads
```

This command displays the bridgehead servers for the site topology.

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

```
Repadmin /rep1summary
```

This command displays a summary of replication tasks. Verify that no errors appear.

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

```
DCDiag /test:replications
```

Verify that all connectivity and replication test pass successfully.

7. Switch to TOR-DC1, and then repeat steps 1 through 6 to view information from the TOR-DC1 perspective.

Results: After this exercise, you will have configured site links and monitored replication.

► To prepare for the next module

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

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

Module 10: Implementing Active Directory Certificate Services

Lab: Implementing Active Directory Certificate Services

Exercise 1: Deploying a standalone root CA

► **Task 1: Install the Active Directory® Certificate Services (AD CS) server role on non-domain joined server**

1. Log on to **LON-CA1** as **Administrator** with the password **Pa\$\$w0rd**.
2. In the Server Manager console, click **Add roles and features**.
3. On the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, click **Next**.
5. On the **Select destination server** page, click **Next**.
6. On the **Select server roles** page, select **Active Directory Certificate Services**. When Add Roles and Features Wizard window displays, click **Add Features**, and then click **Next**.
7. On the **Select features** page, click **Next**.
8. On the **Active Directory Certificate Services** page, click **Next**.
9. On the **Select role services** page, ensure that **Certification Authority** is selected, and then click **Next**.
10. On the **Confirm installation selections** page, click **Install**.
11. On the **Installation progress** page, after installation completes successfully, click the text **Configure Active Directory Certificate Services** on the destination server.
12. In the AD CS Configuration Wizard, on the **Credentials** page, click **Next**.
13. On the **Role Services** page, select **Certification Authority**. Click **Next**.
14. On the **Setup Type** page, select **Standalone CA**, and then click **Next**.
15. On the **CA Type** page, ensure that **Root CA** is selected, and then click **Next**.
16. On the **Private Key** page, ensure that **Create a new private key** is selected, and then click **Next**.
17. On the **Cryptography for CA** page, keep the default selections for Cryptographic Service Provider (CSP) and Hash Algorithm, but set the **Key length** to **4096**, and then click **Next**.
18. On the **CA Name** page, in the **Common name** for this CA box, type **AdatumRootCA**, and then click **Next**.
19. On the **Validity Period** page, click **Next**.
20. On the **CA Database** page, click **Next**.
21. On the **Confirmation** page, click **Configure**.
22. On the **Results** page, click **Close**.
23. On the **Installation progress** page, click **Close**.

► Task 2: Configure a new certificate revocation location

1. On LON-CA1, in the Server Manager console, click **Tools**, and then click **Certification Authority**.
2. In the **certsrv – [Certification Authority (Local)]** console, right-click **AdatumRootCA**, and then click **Properties**.
3. In the AdatumRootCA Properties window, click the **Extensions** tab.
4. In the **Extensions** tab, in the **Select extension:** drop-down list, select **CRL Distribution Point (CDP)** and then click the **Add** button.
5. In the **Location** text box, type **http://lon-svr1.adatum.com/CertData/**, in the **Variable** drop-down list box, click **<CaName>**, and then click **Insert**.
6. In the **Variable** drop-down list box, click **<CRLNameSuffix>**, and then click **Insert**. In the **Variable** drop-down list box, click **<DeltaCRLAllowed>**, and then click **Insert**.
7. In the **Location** text box, position the cursor at the end of URL, type **.crl**, and then click **OK**.
8. Select options: **Include in the CDP extensions of issued certificates** and **Include in CRLs. Clients use this to find Delta CRL locations**. Click **Apply**. In the Certification Authority pop-up window, click **No**.
9. In the **Select extension:** drop-down list box, click **Authority Information Access (AIA)**, and then click **Add**.
10. In the **Location** text box, type **http://lon-svr1.adatum.com/CertData/**, then in **Variable** drop-down box click **<ServerDNSName>**, and then click **Insert**.
11. In the **Location** text box, type an underscore (**_**), in the Variable drop-down list box, click **<CaName>**, and then click **Insert**.
12. In the **Variable** drop-down list box, click **<CertificateName>**, and then click **Insert**.
13. In the **Location** text box, position the cursor at the end of URL, type **.crt**, and then click **OK**.
14. Select the **Include in the AIA extension of issued certificates** check box, and then click **OK**.
15. Click **Yes** to restart Certification Authority service.
16. In the **Certification Authority** console, expand **AdatumRootCA**, right-click **Revoked Certificates**, point to **All Tasks**, and then click **Publish**.
17. In the Publish CRL window, click **OK**.
18. Right-click **AdatumRootCA**, and then click **Properties**.
19. In the AdatumRootCA **Properties** dialog box, click **View Certificate**.
20. In the Certificate window, click the **Details** tab.
21. On the **Details** tab, click **Copy to File**.
22. On the Certificate Export Wizard **Welcome** page, click **Next**.
23. On the **Export File Format** page, select **DER encoded binary X.509 (.CER)**, and then click **Next**.
24. On the **File to Export** page, click **Browse**. In the **File name** text box, type **\\lon-svr1\C\$**, and then press Enter.
25. In the **File name** text box, type **RootCA**, click **Save**, and then click **Next**.
26. Click **Finish**, and then click **OK** three times.
27. Open a Windows® Explorer window, and browse to **C:\Windows\System32\CertSrv\CertEnroll**.

28. In the Cert Enroll folder, select both files, right-click the highlighted files, and then click **Copy**.
29. In the Windows Explorer address bar, type `\\lon-svr1\C$`, and then press Enter.
30. Right-click the empty space, and then click **Paste**.
31. Close Windows Explorer.

Results: After completing this exercise, you will have installed and configured a standalone root CA.

Exercise 2: Deploying an Enterprise Subordinate CA

► Task 1: Install and configure AD CS role on LON-SVR1

1. Log on to **LON-SVR1** as **Adatum\Administrator** with the password of **Pa\$\$w0rd**.
2. In the Server Manager console, click **Add roles and features**.
3. On the **Before you begin** page, click **Next**.
4. On the **Select installation type** page, click **Next**.
5. On the **Select destination server** page, click **Next**.
6. On the **Select server roles** page, select **Active Directory Certificate Services**.
7. When the Add Roles and Features Wizard window displays, click **Add Features**, and then click **Next**.
8. On the **Select features** page, click **Next**.
9. On the **Active Directory Certificate Services** page, click **Next**.
10. On the **Select role services** page, ensure that **Certification Authority** is selected already, and select **Certification Authority Web Enrollment**.
11. When the Add Roles and Features Wizard window displays, click **Add Features**, and then click **Next**.
12. On the **Confirm installation selections** page, click **Install**.
13. On the **Installation progress** page, after installation is successful, click the text **Configure Active Directory Certificate Services** on the destination server.
14. In the AD CS Configuration Wizard, on the **Credentials** page, click **Next**.
15. On the **Role Services** page, select both **Certification Authority** and **Certification Authority Web Enrollment**, and then click **Next**.
16. On the **Setup Type** page, select **Enterprise CA**, and then click **Next**.
17. On the **CA Type** page, click **Subordinate CA**, and then click **Next**.
18. On the **Private Key** page, ensure that **Create a new private key** is selected, and then click **Next**.
19. On the **Cryptography for CA** page, keep the default selections, and then click **Next**.
20. On the **CA Name** page, in the **Common name for this CA** text box, type **Adatum-IssuingCA**, and then click **Next**.
21. On the **Certificate Request** page, ensure that **Save a certificate request to file on the target machine** is selected, and then click **Next**.
22. On the **CA Database** page, click **Next**.
23. On the **Confirmation** page, click **Configure**.
24. On the **Results** page, click **Close**.

25. On the **Installation progress** page, click **Close**.

► **Task 2: Install a subordinate Certification Authority (CA) certificate**

1. On LON-SVR1, open a Windows Explorer window, and navigate to **Local Disk (C:)**.
2. Right-click **RootCA.cer**, and then click **Install Certificate**.
3. In the Certificate Import Wizard, click **Local Machine**, and then click **Next**.
4. On the **Certificate Store** page, click **Place all certificates in the following store**, and then click **Browse**.
5. Select **Trusted Root Certification Authorities**, and then click **OK**.
6. Click **Next**, and then click **Finish**. Click **OK**.
7. In the Windows Explorer window, select the **adatumRootCA.crl** and **LON-CA1_AdatumRootCA.crt** files, right-click the files, and then click **Copy**.
8. Double-click **inetpub**.
9. Double-click **wwwroot**.
10. Create a new folder, and name it **CertData**.
11. Paste the two copied files into that folder.
12. Switch to **Local Disk (C:)**.
13. Right-click the file **LON-SVR1.Adatum.com_Adatum- IssuingCA.req**, and then click **Copy**.
14. In the Windows Explorer address bar, type **\\LON-CA1\C\$**, and then press Enter.
15. In the Windows Explorer window, right-click an empty space, and then click **Paste**. Make sure that request file is copied to LON-CA1.
16. Switch to the LON-CA1 server.
17. In the Certificate Authority console, right-click **AdatumRootCA**, point to **All Tasks**, and then click **Submit new request**.
18. In the Open Request File window, navigate to **Local Disk (C:)**, select file **LON-SVR1.Adatum.com_Adatum- IssuingCA.req**, and then click **Open**.
19. In the Certification Authority console, click the **Pending Requests** container. Right click **Pending Requests** item and click **Refresh**.
20. In the right pane, right-click the request (with ID 2), point to **All Tasks**, and then click **Issue**.
21. Click the **Issued Certificates** container.
22. In the right pane, double-click the certificate, and then click the **Details** tab.
23. Click **Copy to File**.
24. On the Certificate Export Wizard **Welcome** page, click **Next**.
25. On the **Export File Format** page, select **Cryptographic Message Syntax Standard – PKCS #7 Certificates (.P7B)**, select **Include all certificates in the certification path if possible** and then click **Next**.
26. On the **File to Export** page, click **Browse**.
27. In the **File name** text box, type **\\lon-svr1\C\$**, and then press Enter.
28. In the **File name** text box, type **SubCA**, click **Save**, and then click **Next**.

29. Click **Finish**, and then click **OK** twice.
30. Switch to **LON-SVR1**.
31. In Server Manager, click **Tools**, and then click **Certification Authority**.
32. In the Certification Authority console, right-click **Adatum-IssuingCA**, point to **All Tasks**, and then click **Install CA Certificate**.
33. Navigate to **Local Disk (C:)**, click the **SubCA.p7b** file, and then click **Open**.
34. Wait for 15-20 seconds, and then on the toolbar, click the green icon to start the CA service.
35. Ensure that CA starts successfully.

► **Task 3: Publish the RootCA certificate through Group Policy**

1. On LON-DC1, open **Server Manager**, click **Tools**, and then click **Group Policy Management**.
2. In the Group Policy Management Console, expand **Forest:Adatum.com**, expand **Domains**, expand **Adatum.com**, right-click **Default Domain Policy**, and then click **Edit**.
3. In the **Computer Configuration** node, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, expand **Public Key Policies**, right-click **Trusted Root Certification Authorities**, and then click **Import**.
4. Click **Next**.
5. On the **File to Import** page, click **Browse**.
6. In the **file name** text field, type **\\lon-svr1\C\$**, and then press Enter.
7. Select file **RootCA.cer**, and then click **Open**.
8. Click **Next** two times, and then click **Finish**.
9. Click **OK**.
10. Close the Group Policy Management Editor.
11. Close the Group Policy Management Console.

Results: After completing this exercise, you will have deployed and configured an enterprise subordinate CA

Exercise 3: Configuring Certificate Templates

► **Task 1: Create a new template based on the Web server template**

1. On LON-SVR1, in the Certification Authority console, expand **Adatum-IssuingCA**, right-click **Certificate Templates**, and then select **Manage**.
2. In the Certificate Templates Console, locate the **Web Server** template in the list, right-click it, and then select **Duplicate Template**.
3. Click the **General** tab.
4. In the **Template** display name field, type **Adatum Web Server**, and set the **Validity period** to **3 years**
5. Click the **Request Handling** tab, select **Allow private key to be exported**, and then click **OK**.

► **Task 2: Create a new template for users that includes smart card logon**

1. In the Certificate Templates Console, right-click the **User** certificate template, and then click **Duplicate Template**.
2. In the **Properties of New Template** dialog box, click the **General** tab, and in the **Template display name** text box, type **Adatum Smart Card User**.
3. On the **Subject Name** tab, clear both the **Include e-mail name in subject name** and the **E-mail name** check boxes.
4. On the **Extensions** tab, click **Application Policies**, and then click **Edit**.
5. In the **Edit Application Policies Extension** dialog box, click **Add**.
6. In the **Add Application Policy** dialog box, select **Smart Card Logon**, and then click **OK** twice.
7. Click the **Superseded Templates** tab, and then click **Add**.
8. Click the **User** template, and then click **OK**.
9. On the **Security** tab, click **Authenticated Users**. Under **Permissions for Authenticated Users**, select the **Allow** check box for **Read**, **Enroll** and **Autoenroll**, and then click **OK**.
10. Close the Certificate Templates Console.

► **Task 3: Configure the templates so they can be issued**

1. On LON-SVR1, in the Certification Authority console, right-click **Certificate Templates**, point to **New**, and then click **Certificate Template to Issue**.
2. In the Enable Certificate Templates window, select **Adatum Smart Card User** and **Adatum Web Server**, and then click **OK**.

► **Task 4: Update the Web server certificate on the LON-SVR2 Web Server**

1. Log on to **LON-SVR2** as **Adatum\Administrator** with the password of **Pa\$\$w0rd**.
2. Open Windows PowerShell window from taskbar and type **gpupdate /force** and press Enter. If prompted to do so, restart the server, and logon with same credentials as in step 1.
3. From Server Manager, click **Tools**, and then click **Internet Information Services (IIS) Manager**.
4. In the IIS console, click **LON-SVR2**, click **No** at the **Internet Information Services (IIS) Manager** prompt, and then in the central pane, double-click **Server Certificates**.
5. In the Actions pane, click **Create Domain Certificate**.
6. On the **Distinguished Name Properties** page, complete the following fields, and then click **Next**:
 - Common name: **lon-svr2.adatum.com**
 - Organization: **Adatum**
 - Organizational Unit: **IT**
 - City/locality: **Seattle**
 - State/province: **WA**
 - Country/region: **US**
7. On the **Online Certification Authority** page, click **Select**.
8. Click **Adatum-IssuingCA**, and then click **OK**.
9. In the friendly name text box, type **lon-svr2**, and then click **Finish**.

10. Ensure that the certificate displays in the Server Certificates console.
11. In the IIS console, expand **LON-SVR2**, expand **Sites**, and then click **Default Web Site**.
12. In the Actions pane, click **Bindings**.
13. In the Site Bindings window, click **Add**.
14. In the **Type** drop-down list box, click **https**.
15. In the **SSL certificate** drop-down list box, click **lon-svr2**, click **OK**, and then click **Close**.
16. Close the IIS console.

Results: After completing this exercise, you will have created and published new certificate templates.

Exercise 4: Configuring Certificate Enrollment

► Task 1: Configure autoenrollment for users

1. On LON-DC1, in Server Manager, click **Tools**, and then click **Group Policy Management**.
2. Expand **Forest: Adatum.com**, expand **Domains**, expand **Adatum.com**, right-click **Default Domain Policy**, and then click **Edit**.
3. Expand **User Configuration**, expand **Policies**, expand **Windows Settings**, expand **Security Settings**, and then click to highlight **Public Key Policies**.
4. In the right pane, double-click **Certificate Services Client – Auto-Enrollment**.
5. In the **Configuration Model** drop-down list box, click **Enabled**.
6. Select the **Renew expired certificates, update pending certificates, and remove revoked certificates** option.
7. Select the **Update certificates that use certificate templates** option.
8. Click **OK** to close the properties window.
9. In the right pane, double-click the **Certificate Services Client – Certificate Enrollment Policy** object.
10. On the **Enrollment Policy** tab, set the **Configuration Model** to **Enabled**, and ensure that the certificate enrollment policy list displays the **Active Directory Enrollment Policy** (it should have a checkmark next to it, and a status of **Enabled**).
11. Click **OK** to close the window.
12. Close both the Group Policy Management Editor and the Group Policy Management console.

► Task 2: Verify autoenrollment

1. On LON-SVR1, open Windows PowerShell® from task bar.
2. Type **gpupdate /force**, and then press Enter.
3. After the policy is refreshed, type **mmc.exe**, and then press Enter.
4. In Console1, click **File**, and then in the **File** menu, click **Add/Remove Snap-in**.
5. Click **Certificates**, and then click **Add>**.
6. Click **Finish**, and then click **OK**.
7. Expand **Certificates – Current User**, expand **Personal**, and then click **Certificates**.

8. Verify that certificate based on **Adatum Smart Card User** template is issued for administrator.
9. Close Console1.

► **Task 3: Configure the Enrollment Agent for smart card certificates**

1. On LON-SVR1, in the Server Manager console, click **Tools**, and then open **Certification Authority**.
2. In the certsrv console, expand **Adatum-IssuingCA**, right-click **Certificate Templates**, and then click **Manage**.
3. In the Certificate Templates console, double-click **Enrollment Agent**.
4. Click the **Security** tab, and then click **Add**.
5. In the Select Users, Computers, Service Accounts, or Groups window, type **Allie**, click **Check Names**, and then click **OK**.
6. On the **Security** tab, click **Allie Bellew**, select **Allow** for **Read** and **Enroll** permissions, and then click **OK**.
7. Close the Certificate Templates Console.
8. In the certsrv console, right-click **Certificate Templates**, point to **New**, and then click **Certificate Template to Issue**.
9. In the list of templates, click **Enrollment Agent**, and then click **OK**.
10. Switch to LON-CL1, and log on as **Adatum\Allie** with the password **Pa\$\$w0rd**.
11. Open a command prompt window, and at a command prompt, type **mmc.exe**, and then press Enter.
12. In Console1, click **File**, and then click **Add/Remove Snap-in**.
13. Click **Certificates**, and then click **Add>**.
14. Click **OK**.
15. Expand **Certificates – Current User**, expand **Personal**, click **Certificates**, right-click **Certificates**, point to **All Tasks**, and then click **Request New Certificate**.
16. In the Certificate Enrollment Wizard, on the **Before You Begin** page, click **Next**.
17. On the **Select Certificate Enrollment Policy** page, click **Next**.
18. On the **Request Certificates** page, select **Enrollment Agent**, and then click **Enroll**.
19. Click **Finish**.
20. Switch to LON-SVR1.
21. In the Certification Authority console, right-click **Adatum-IssuingCA**, and then click **Properties**.
22. Click the **Enrollment Agents** tab.
23. Click **Restrict Enrollment agents**.
24. On the pop-up window that displays, click **OK**.
25. In the Enrollment agents section, click **Add**.
26. In the **Select User, Computer or Group** field, type **Allie**, click **Check Names**, and then click **OK**.
27. Click **Everyone**, and then click **Remove**.
28. In the certificate templates section, click **Add**.
29. In the list of templates, select **Adatum Smart Card User**, and then click **OK**.

30. In the Certificate Templates section, click **<All>**, and then click **Remove**.
31. In the Permission section, click **Add**.
32. In the **Select User, Computer or Group** field, type **Marketing**, click **Check Names**, and then click **OK**.
33. In the Permission section, click **Everyone**, and then click **Remove**.
34. Click **OK**.

Results: After completing this exercise, you will have configured and verified autoenrollment for users, and configured an enrollment agent for smart cards.

Exercise 5: Configuring Certificate Revocation

► Task 1: Configure Certified Revocation List (CRL) distribution

1. On LON-SVR1, in the Certification Authority console, right-click **Revoked Certificates**, and then click **Properties**.
2. In the Revoked Certificates Properties window, set the **CRL publication interval** to **1 Day** and the **Delta CRL** publication interval to **1 hour**, and then click **OK**.
3. Right-click **Adatum-IssuingCA**, and then click **Properties**.
4. In the Properties window, click the **Extensions** tab.
5. On the **Extensions** tab, review the values for **CDP**.
6. Click **Cancel**.

► Task 2: Install and configure an Online Responder

1. On LON-SVR1, open Server Manager.
2. In Server Manager, click **Add roles and features**.
3. Click **Next** three times.
4. On the **Select server roles** page, expand **Active Directory Certificate Services (Installed)**, and then select **Online Responder**.
5. Click **Add Features**.
6. Click **Next** two times, and then click **Install**.
7. When the message displays that installation succeeded, click **Configure Active Directory Certificate Services on the destination server**.
8. In AD CS Configuration Wizard, click **Next**.
9. Select **Online Responder**, and then click **Next**.
10. Click **Configure**, and then click **Close** two times.
11. On LON-SVR1, open the Certification Authority console.
12. In the Certification Authority console, right-click **Adatum-IssuingCA**, and then click **Properties**.
13. In the **Adatum-IssuingCA Properties** dialog box, on the **Extensions** tab, in the **Select extension** list, click **Authority Information Access (AIA)**, and then click **Add**.
14. In the **Add Location** dialog box, type **http://LON-SVR1/ocsp**, and then click **OK**.

15. Select the **Include in the AIA extension of issued certificates** check box.
16. Select the **Include in the online certificate status protocol (OCSP) extension** check box, and then click **OK**.
17. In the **Certificate Authority** dialog box, restart AD CS by clicking **Yes**.
18. In the certsrv console, expand **Adatum-IssuingCA**, right-click the **Certificate Templates** folder, and then click **Manage**.
19. In the Certificate Templates console, double-click the **OCSP Response Signing** template.
20. In the **OCSP Response Signing Properties** dialog box, click the **Security** tab, under **Permissions for Authenticated Users**, select the **Allow** for **Enroll** check box, and then click **OK**.
21. Close the Certificate Templates console.
22. In the Certification Authority console, right-click the **Certificate Templates** folder, point to **New**, and then click **Certificate Template to Issue**.
23. In the **Enable Certificate Templates** dialog box, select the **OCSP Response Signing** template, and then click **OK**.
24. On LON-SVR1, in Server Manager, click **Tools**, and then click **Online Responder Management**.
25. In the ocsf Management console, right-click **Revocation Configuration**, and then click **Add Revocation Configuration**.
26. In the Add Revocation Configuration Wizard, click **Next**.
27. On the **Name the Revocation Configuration** page, in the **Name** box, type **AdatumCA Online Responder**, and then click **Next**.
28. On the **Select CA Certificate Location** page, click **Next**.
29. On the **Choose CA Certificate** page, click **Browse**, click the **Adatum-IssuingCA** certificate, click **OK**, and then click **Next**.
30. On the **Select Signing Certificate** page, verify that **Automatically select a signing certificate is selected**, and **Auto-Enroll for an OCSP signing certificate** are both selected, and then click **Next**.
31. On the **Revocation Provider** page, click **Finish**. The revocation configuration status will appear as **Working**.
32. Close the **Online Responder** console.

Results: After completing this exercise, you will have configured certificate revocation settings.

Exercise 6: Configuring Key Recovery

► Task 1: Configure the CA to issue Key Recovery Agent (KRA) certificates

1. On LON-SVR1, open the Certification Authority console.
2. In the Certification Authority console, expand the **Adatum-IssuingCA** node, right-click the **Certificates Templates** folder, and then click **Manage**.
3. In the Details pane, right-click the **Key Recovery Agent** certificate, and then click **Properties**.
4. In the **Key Recovery Agent Properties** dialog box, click the **Issuance Requirements** tab.
5. Clear the **CA certificate manager approval** check box.

6. Click the **Security** tab. Notice that Domain Admins and Enterprise Admins are the only groups that have the Enroll permission, and then click **OK**.
7. Close the Certificate Templates Console.
8. In the Certification Authority console, right-click **Certificate Templates**, point to **New**, and then click **Certificate Template to Issue**.
9. In the **Enable Certificate Templates** dialog box, select the **Key Recovery Agent** template, and then click **OK**.
10. Close the Certification Authority console.

► **Task 2: Acquire the KRA certificate**

1. On LON-SVR1, open Windows PowerShell window. At a command prompt, type **MMC.exe**, and then press Enter.
2. In the Console1-[Console Root] console, click **File**, and then click **Add/Remove Snap-in**.
3. In the **Add or Remove Snap-ins** dialog box, click **Certificates**, and then click **Add**.
4. In the **Certificates snap-in** dialog box, select **My user account**, click **Finish**, and then click **OK**.
5. Expand the **Certificates - Current User** node, and right-click **Personal**.
6. Point to **All Tasks**, and then click **Request New Certificate**.
7. In the Certificate Enrollment Wizard, on the **Before You Begin** page, click **Next**.
8. On the **Select Certificate Enrollment Policy** page, click **Next**.
9. On the **Request Certificates** page, select the **Key Recovery Agent** check box. Click **Enroll**, and then click **Finish**.
10. Refresh the console, and view the KRA in the personal store; that is, scroll across the certificate properties and verify that the Certificate Template Key Recovery Agent is present.
11. Close Console1 without saving changes.

► **Task 3: Configure the CA to allow key recovery**

1. On LON-SVR1, in the Certification Authority console, right-click **Adatum-IssuingCA**, and then click **Properties**.
2. In the **Adatum-IssuingCA Properties** dialog box, click the **Recovery Agents** tab, and then select **Archive the key**.
3. Under **Key recovery agent certificates**, click **Add**.
4. In the **Key Recovery Agent Selection** dialog box, click the certificate that is for Key Recovery Agent purpose (it will most likely be last on the list), and then click **OK** twice.
5. When prompted to restart the CA, click **Yes**.

► **Task 4: Configure a custom template for key archival**

1. On LON-SVR1, in the Certification Authority console, right-click the **Certificates Templates** folder, and then click **Manage**.
2. In the Certificate Templates console, right-click the **User** certificate, and then click **Duplicate Template**.
3. In the **Properties of New Template** dialog box, on the **General** tab, in the **Template display name** box, type **Archive User**.

4. On the **Request Handling** tab, select the **Archive subject's encryption private key** check box. Click **OK** on the popup window.
5. Click the **Subject Name** tab, clear the **E-mail name** and **Include e-mail name in subject name** check boxes, and then click **OK**.
6. Close the Certificate Templates Console.
7. In the Certification Authority console, right-click the **Certificates Templates** folder, point to **New**, and then click **Certificate Template to Issue**.
8. In the **Enable Certificate Templates** dialog box, select the **Archive User** template, and then click **OK**.
9. Close the Certification Authority console.

► **Task 5: Verify key archival functionality**

1. Log on to the **LON-CL1** virtual computer as **Adatum\Aidan**, using the password **Pa\$\$w0rd**.
2. On the Start screen, type **mmc.exe** and then press Enter.
3. In the Console1-[Console Root] console, click **File**, and then click **Add/Remove Snap-in**.
4. In the **Add or Remove Snap-ins** dialog box, click **Certificates**, and then click **Add**. Click **OK**.
5. Expand the **Certificates - Current User** node, right click Personal, click **All Tasks**, and then click **Request New Certificate**.
6. In the Certificate Enrollment Wizard, on the **Before You Begin** page, click **Next**.
7. Click **Next**.
8. On the **Request Certificate** page, select the **Archive User** check box, click **Enroll**, and then click **Finish**.
9. Refresh the console, and view that a certificate is issued to Aidan, based on the **Archive User** certificate template.
10. Simulate the loss of a private key by deleting the certificate. In the central pane, right-click the certificate that you just enrolled, select **Delete**, and then click **Yes** to confirm.
11. Switch to LON-SVR1.
12. Open the Certification Authority console, expand **Adatum-IssuingCA**, and then click **Issued Certificates** store.
13. In the details pane, double-click a certificate with Requestor Name **Adatum\Aidan**, and Certificate Template name of **Archive User**.
14. Click the **Details** tab, copy the **Serial Number**, and then click **OK**. (You may either copy the number to Notepad (select it and press CTRL+C), or write it down on paper.)
15. Open Windows PowerShell console from task bar.
16. In the command prompt window that appears, type the following command (where *<serial number>* is the serial number that you copied), and then press Enter:

```
certutil -getkey <serial number> outputblob
```



Note: If you paste the serial number from Notepad, remove spaces between numbers.

17. Verify that **outputblob** file now displays in the C:\Users\Administrator.Adatum folder.

18. To convert the outputblob file into a .pfx file, in the command prompt window, type the following command, and press Enter:

```
Certutil -recoverkey outputblob aidan.pfx
```

19. When prompted, type **Pa\$\$w0rd** as the new password, and then confirm the password.
20. After the command executes, close the Windows PowerShell window.
21. Browse to **C:\Users\Administrator.ADATUM**, and then verify that **aidan.pfx**—the recovered key—is created.
22. Copy **aidan.pfx** file to **\\lon-cl1\C\$**.
23. Switch to LON-CL1, and ensure that you are still logged on as **Aidan**.
24. Browse to drive **C** and double-click the **aidan.pfx** file.
25. On the **Welcome to the Certificate Import Wizard** page, click **Next**.
26. On the **File to Import** page, click **Next**.
27. On the **Password** page, enter **Pa\$\$w0rd** as password, and then click **Next**.
28. On the **certificate store** page, click **Next**, click **Finish**, and then click **Ok**.
29. Expand the **Certificates - Current User** node, expand **Personal**, and then click **Certificates**.
30. Refresh the console, and verify that the certificate for **Aidan** is restored.

Results: After completing this exercise, you will have implemented key archival, and tested private key recovery.

► To prepare for the next module

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

1. On the host computer, start **Hyper-V Manager**.
2. In the **Virtual Machines** list, right-click **20412A-LON-DC1**, and then click **Revert**.
3. In the **Revert Virtual Machine** dialog box, click **Revert**.
4. Repeat steps 2 and 3 for **20412A-LON-CL1**, **20412A-LON-SVR1**, **20412A-LON-CA1** and **20412A-LON-SVR2**.

MCT USE ONLY. STUDENT USE PROHIBITED

Module 11: Implementing Active Directory Rights Management Services

Lab: Implementing AD RMS

Exercise 1: Installing and Configuring AD RMS

► Task 1: Configure Domain Name System (DNS) and the Active Directory® Rights Management Services (AD RMS) service account

1. Log on to **LON-DC1** with the **Adatum\Administrator** account and the password **Pa\$\$w0rd**.
2. In Server Manager, click **Tools**, and then click **Active Directory Administrative Center**.
3. Select and then right-click **Adatum (local)**, click **New**, and then click **Organizational Unit**.
4. In the **Create Organizational Unit** dialog box, in the **Name** field, type **Service Accounts**, and then click **OK**.
5. Right-click the **Service Accounts** OU, click **New**, and then click **User**.
6. On the **Create User** dialog box, enter the following details, and then click **OK**:
 - First name: **ADRMSSVC**
 - User UPN logon: **ADRMSSVC**
 - Password: **Pa\$\$w0rd**
 - Password never expires: **Enabled**
 - User cannot change password: **Enabled**
7. Right-click the **Users** container, click **New**, and then click **Group**.
8. In the **Create Group** dialog box, enter the following details, and then click **OK**:
 - Group name: **AD RMS_SuperUsers**
 - E-mail: **AD RMS_SuperUsers@adatum.com**
9. Right-click the **Users** container, click **New**, and then click **Group**.
10. In the **Create Group** dialog box, enter the following details, and then click **OK**.
 - Group name: **Executives**
 - E-mail: **executives@adatum.com**
11. Double-click the **Managers** OU.
12. Hold down the Ctrl key, and click the following users:
 - **Aidan Delaney**
 - **Bill Malone**
13. In the Tasks pane, click **Add to group ...**
14. In the **Select Groups** dialog box, type **Executives**, and then click **OK**.
15. Close the Active Directory Administrative Center.
16. In Server Manager, click **Tools**, and then click **DNS**.
17. In the DNS Manager console, expand **LON-DC1**, and expand **Forward Lookup Zones**.

18. Select and then right-click **Adatum.com**, and click **New Host (A or AAAA)**.
 19. In the **New Host** dialog box, enter the following information, and then click **Add Host**:
 - o Name: **adrms**
 - o IP address: **172.16.0.21**
 20. Click **OK**, and then click **Done**, and close the DNS Manager console.
- **Task 2: Install and configure the AD RMS server role**
1. Log on to **LON-SVR1** with the **Adatum\Administrator** account and the password **Pa\$\$word**.
 2. In Server Manager, click **Manage**, and then click **Add Roles and Features**.
 3. In the Add Roles and Features Wizard, click **Next** three times.
 4. On the **Server Roles** page, click **Active Directory Rights Management Services**.
 5. In the **Add Roles and Features** dialog box, click **Add Features**, and then click **Next** four times.
 6. Click **Install**, and then click **Close**.
 7. In Server Manager, click the **AD RMS** node.
 8. Next to Configuration required for Active Directory Rights Management Services at LON-SVR1, click **More**.
 9. On the **All Servers Task Details** page, click **Perform Additional Configuration**.
 10. In the **AD RMS Configuration: LON-SVR1.adatum.com** dialog box, click **Next**.
 11. On the **AD RMS Cluster** page, click **Create a new AD RMS root cluster**, and then click **Next**.
 12. On the **Configuration Database** page, click **Use Windows Internal Database on this server**, and then click **Next**.
 13. On the **Service Account** page, click **Specify**.
 14. In the **Windows Security** dialog box, enter the following details, click **OK**, and then click **Next**:
 - o Username: **ADRMSSVC**
 - o Password: **Pa\$\$w0rd**
 15. On the **Cryptographic Mode** page, click **Cryptographic Mode 2**, and then click **Next**.
 16. On the **Cluster Key Storage** page, click **Use AD RMS centrally managed key storage**, and then click **Next**.
 17. On the **Cluster Key Password** page, enter the password **Pa\$\$w0rd** twice, and then click **Next**.
 18. On the **Cluster Web Site** page, verify that **Default Web Site** is selected, and then click **Next**.
 19. On the **Cluster Address** page, provide the following information, and then click **Next**:
 - o Connection Type: **Use an unencrypted connection (http://)**
 - o Fully Qualified Domain Name: **adrms.adatum.com**
 - o Port: **80**
 20. On the **Licensors Certificate** page, type **Adatum AD RMS**, and then click **Next**.
 21. On the **SCP Registration** page, click **Register the SCP now**, and then click **Next**.
 22. Click **Install**, and then click **Close**.
 23. Click to the Start screen, click **Administrator**, and then click **Sign Out**.



Note: You must sign out before you can manage AD RMS.

► Task 3: Configure the AD RMS Super Users group

1. Log on to **LON-SVR1** with the **Adatum\Administrator** account and the password **Pa\$\$w0rd**.
2. In Server Manager, click **Tools**, and then click **Active Directory Rights Management Services**.
3. In the Active Directory Rights Management Services console, expand the **LON-SVR1** node, and then click **Security Policies**.
4. In the **Security Policies** area, under **Super Users**, click **Change super user settings**.
5. In the Actions pane, click **Enable Super Users**.
6. In the **Super Users** area, click **Change super user group**.
7. In the **Super Users** dialog box, in the **Super user group** text box, type **ADRMS_Superusers@adatum.com**, and then click **OK**.

Results: After completing this exercise, you should have installed and configured AD RMS.

Exercise 2: Configuring AD RMS Templates

► Task 1: Configure a new rights policy template

1. Ensure that you are logged on to **LON-SVR1**.
2. In the Active Directory Rights Management Services console, click the **LON-SVR1\Rights Policy Templates** node.
3. In the Actions pane, click **Create Distributed Rights Policy Template**.
4. In the Create Distributed Rights Policy Template Wizard, on the **Add Template Identification information** page, click **Add**.
5. On the **Add New Template Identification Information** page, enter the following information, and then click **Add**:
 - Language: **English (United States)**
 - Name: **ReadOnly**
 - Description: **Read only access. No copy or print**
6. Click **Next**.
7. On the **Add User Rights** page, click **Add**.
8. On the **Add User or Group** page enter **executives@adatum.com**, and then click **OK**.
9. When **executives@adatum.com** is selected, under **Rights**, click **View**. Verify that **Grant owner (author) full control right with no expiration** is selected, and then click **Next**.
10. On the **Specify Expiration Policy** page, choose the following settings and then click **Next**:
 - Content Expiration: Expires after the following duration (days): **7**
 - Use license expiration: Expires after the following duration (days): **7**
11. On the **Specify Extended Policy** page, click **Require a new use license every time content is consumed (disable client-side caching)**, click **Next**, and then click **Finish**.

► Task 2: Configure the rights policy template distribution

1. On LON-SVR1, on the taskbar, click the Windows PowerShell® icon.
2. In the Windows PowerShell window, issue the following commands, each followed by Enter:

```
Cmd.exe
mkdir c:\rmstemplates
net share RMTEMPLATES=C:\rmstemplates /GRANT:ADATUM\ADRMSSVC,FULL
mkdir c:\docshare
net share docshare=c:\docshare /GRANT:Everyone,FULL
```

3. To exit the Windows PowerShell window, type **exit** twice.
4. Switch to the Active Directory Rights Management Services console.
5. Click the **Rights Policy Templates** node, and in the Distributed Rights Policy Templates area, click **Change distributed rights policy templates file location**.
6. In the **Rights Policy Templates** dialog box, click Enable Export.
7. In the Specify Templates File Location (UNC), type **\\LON-SVR1\RMTEMPLATES**, and then click **OK**.
8. On the taskbar, click the Windows Explorer icon.
9. Navigate to the **C:\rmstemplates** folder, and verify that **ReadOnly.xml** is present.
10. Close the Windows Explorer window.

► Task 3: Configure an exclusion policy

1. Switch to the Active Directory Rights Management Services console.
2. Click the **Exclusion Policies** node, and then click **Manage application exclusion list**.
3. In the Actions pane, click **Enable Application Exclusion**.
4. In the Actions pane, click **Exclude Application ...**
5. In the **Exclude Application** dialog box, enter the following information, and then click **Finish**:
 - Application File name: **Powerpnt.exe**
 - Minimum version: **14.0.0.0**
 - Maximum version: **16.0.0.0**

Results: After completing this exercise, you should have configured AD RMS templates.

Exercise 3: Implementing the AD RMS Trust Policies**► Task 1: Export the Trusted User Domains policy**

1. On LON-SVR1, on the taskbar, click the Windows PowerShell icon.
2. In the Windows PowerShell window, issue the following commands, and then press Enter:

```
Cmd.exe
mkdir c:\export
net share export=c:\export /GRANT:Everyone,FULL
```

3. To close the Windows PowerShell window, type **exit** twice.
4. In the Active Directory Rights Management Services console, expand the **Trust Policies** node, and then click the **Trusted User Domains** node.

5. In the Actions pane, click **Export Trusted User Domains**.
6. In the **Export Trusted User Domains As** dialog box, navigate to `\\LON-SVR1\export`, set the file name to **ADATUM-TUD.bin**, and then click **Save**.
7. Log on to **MUN-DC1** with the **TREYRESEARCH\Administrator** account and the password **Pa\$\$w0rd**.
8. In Server Manager, click **Tools**, and then click **Active Directory Rights Management**.
9. In the Active Directory Rights Management Services console, expand **MUN-DC1**, expand the **Trust Policies** node, and then click the **Trusted User Domains** node.
10. In the Actions pane, click **Export Trusted User Domains**.
11. In the **Export Trusted User Domains As** dialog box, navigate to `\\LON-SVR1\export`, set the file name to **TREYRESEARCH-TUD.bin**, and then click **Save**.

► **Task 2: Export the Trusted Publishing Domains policy**

1. Switch to LON-SVR1.
2. In the Active Directory Rights Management Services console, under the **Trust Policies** node, click the **Trusted Publishing Domains** node.
3. In the Actions pane, click **Export Trusted Publishing Domains**.
4. In the **Export Trusted Publishing Domain** dialog box, click **Save As**.
5. In the **Export Trusted Publishing Domain File As** dialog box, navigate to `\\LON-SVR1\export`, set the file name to **ADATUM-TPD.xml**, and then click **Save**.
6. In the **Export Trusted Publishing Domain** dialog box, enter the password **Pa\$\$w0rd** twice, and then click **Finish**.
7. Switch to MUN-DC1.
8. In the Active Directory Rights Management Services console, under the **Trust Policies** node, click the **Trusted Publishing Domains** node.
9. In the Actions pane, click **Export Trusted Publishing Domains**.
10. In the **Export Trusted Publishing Domain** dialog box, click **Save As**.
11. In the **Export Trusted Publishing Domain File As** dialog box, navigate to `\\LON-SVR1\export`, set the file name to **TREYRESEARCH-TPD.xml**, and then click **Save**.
12. In the **Export Trusted Publishing Domain** dialog box, enter the password **Pa\$\$w0rd** twice, and then click **Finish**.

► **Task 3: Import the Trusted User Domain policy from the partner domain**

1. Switch to LON-SVR1.
2. In the Active Directory Rights Management Services console, under the **Trust Policies** node, click the **Trusted User Domains** node.
3. In the Actions pane, click **Import Trusted User Domain**.
4. In the **Import Trusted User Domain** dialog box, enter the following details, and then click **Finish**:
 - Trusted user domain file: `\\LON-SVR1\export\TREYRESEARCH-TUD.bin`
 - Display Name: **Trey Research**
5. Switch to MUN-DC1.

6. In the Active Directory Rights Management Services console, under the **Trust Policies** node, click the **Trusted User Domains** node.
7. In the Actions pane, click **Import Trusted User Domain**.
8. In the **Import Trusted User Domain** dialog box, enter the following details, and then click **Finish**:
 - o Trusted user domain file: **\\LON-SVR1\Export\ADATUM-TUD.bin**
 - o Display Name: **Adatum**

► **Task 4: Import the Trusted Publishing Domains policy from the partner domain**

1. Switch to LON-SVR1.
2. In the Active Directory Rights Management Services console, under the **Trust policies** node, click the **Trusted Publishing Domains** node.
3. In the Actions pane, click **Import Trusted Publishing Domain**.
4. In the **Import Trusted Publishing Domain** dialog box, enter the following information, and then click **Finish**:
 - o Trusted publishing domain file: **\\LON-SVR1\export\ TREYRESEARCH-TPD.xml**
 - o Password: **Pa\$\$w0rd**
 - o Display Name: **Trey Research**
5. Switch to MUN-DC1.
6. In the Active Directory Rights Management Services console, under the **Trust policies** node, click the **Trusted Publishing Domains** node.
7. In the Actions pane, click **Import Trusted Publishing Domain**.
8. In the **Import Trusted Publishing Domain** dialog box, provide the following information, and then click **Finish**:
 - o Trusted publishing domain file: **\\LON-SVR1\export\adatum-tpd.xml**
 - o Password: **Pa\$\$w0rd**
 - o Display Name: **Adatum**

► **Task 5: Configure anonymous access to the AD RMS licensing server**

1. Switch to LON-SVR1.
2. In Server Manager, click **Tools**, and then click **Internet Information Services (IIS) Manager**.
3. In Internet Information Services (IIS) Manager, expand **LON-SVR1\Sites\Default Web Site\wmcs**.
4. Right-click **licensing**, and then click **Switch to Content View**.
5. Right-click **license.asmx**, and then click **Switch to Features View**.
6. Double-click **Authentication**, click **Anonymous Authentication**, and in the Actions pane, click **Enable**.
7. Right-click **licensing**, and then click **Switch to Content View**.
8. Right-click **ServiceLocator.asmx**, and then click **Switch to Features View**.
9. Double-click **Authentication**, click **Anonymous Authentication**, and in the Actions pane, click **Enable**.

10. Close Internet Information Services (IIS) Manager.

Results: After completing this exercise, you should have implemented the AD RMS trust policies.

Exercise 4: Verifying the AD RMS Deployment

► Task 1: Create a rights-protected document

1. Log on to **LON-CL1** as **Adatum\Aidan** using the password **Pa\$\$w0rd**.
2. On the Start screen, type **Word**. In the Results area, click **Microsoft Word 2010**.
3. In the **User Name** dialog box, click **OK**.
4. In the **Welcome to Microsoft Office 2010** dialog box, click **Don't make changes**, and then click **OK**.
5. In the document, type the following text:
This document is for executives only, it should not be modified.
6. Click **File**, click **Protect Document**, click **Restrict Permission by People**, and then click **Manage Credentials**.
7. In the **Windows Security** dialog box, enter the following credentials.
 - User name: **Aidan**
 - Password: **Pa\$\$w0rd**
8. Enable **Remember My Credentials**, and then click **OK**.
9. In the **Select User** dialog box, click **OK**.
10. In the **Permission** dialog box, enable **Restrict Permission to this document**.
11. In the **Read** text box, type **bill@adatum.com**, and then click **OK**.
12. Click **Save**.
13. In the **Save As** dialog box, save the document to the **\\lon-svr1\docshare** location as **Executives Only.docx**.
14. Click to the Start screen, click the **Aidan Delaney** icon, and then click **Sign out**.

► Task 2: Verify internal access to protected content

1. Log on to **LON-CL1** as **Adatum\Bill** using the password **Pa\$\$w0rd**.
2. On the Start screen, click **Desktop**.
3. On the taskbar, click the **Windows Explorer** icon.
4. In the Windows Explorer window, navigate to **\\lon-svr1\docshare**.
5. Double-click the **Executives Only** document.
6. In the **User Name** dialog box, click **OK**.
7. In the **Microsoft Word** dialog box, click **Yes**.
8. In the **Windows Security** dialog box, enter the following credentials, select **Remember my credentials**, and then click **OK**.
 - Username: **Bill**

- o Password: **Pa\$\$w0rd**
9. In the **Select User** dialog box, ensure that **bill@adatum.com** is selected, and then click **OK**.
 10. In the **Microsoft Office** dialog box, click **OK**.
 11. In the **Welcome to Microsoft Office 2010** dialog box, click **Don't make changes**, and then click **OK**.
 12. When the document opens, verify that you are unable to modify or save the document.
 13. Select a line of text in the document.
 14. Right-click the text, and verify that you cannot make changes.
 15. Click **View Permission**, review the permissions, and then click **OK**.
 16. Click to the Start screen, click the **Bill Malone** icon, and then click **Sign out**.
- ▶ **Task 3: Open the rights-protected document as an unauthorized user**
1. Log on to **LON-CL1** as **Adatum\Carol** using the password **Pa\$\$w0rd**.
 2. On the Start menu, click **Desktop**.
 3. On the taskbar, click the **Windows Explorer** icon.
 4. In the Windows Explorer window, navigate to **\\lon-svr1\docshare**.
 5. Double-click the **Executives Only** document.
 6. Verify that Carol is unable to open the document.
 7. Click to the Start screen, click the **Carol Troup** icon, and then click **Sign out**.
- ▶ **Task 4: Open and edit the rights-protected document as an authorized user at Trey Research.**
1. Log on to **LON-CL1** as **Adatum\Aidan** using the password **Pa\$\$w0rd**.
 2. On the Start screen, type **Word**. In the Results area, click **Microsoft Word 2010**.
 3. In the document, type the following text:
This document is for Trey Research only, it should not be modified.
 4. Click **File**, click **Protect Document**, click **Restrict Permission by People**, and then click **Manage Credentials**.
 5. In the **Select User** dialog box, click **OK**.
 6. In the **Permission** dialog box, enable **Restrict Permission to this document**.
 7. In the **Read** text box, enter **april@treymresearch.net**, click **OK**, and then click **Save**.
 8. In the **Save As** dialog box, save the document to the **\\lon-svr1\docshare** location as **TreyResearch-Confidential.docx**.
 9. Click to the Start screen, click the **Aidan Delaney** icon, and then click **Sign Out**.
 10. Log on to **MUN-CL1** as **TREYRESEARCH\APRIL**.
 11. On the Start screen, click **Desktop**.
 12. On the taskbar, click the **Windows Explorer** icon.
 13. In the Windows Explorer window, navigate to **\\lon-svr1\docshare**.
 14. In the **Windows Security** dialog box, enter the following credentials, and then click **OK**:

- Username: **Adatum\Administrator**
 - Password: **Pa\$\$w0rd**
15. Copy the file **TreyResearch-Confidential.docx** to the desktop.
 16. Double-click the file.
 17. In the **User Name** dialog box, click **OK**.
 18. In the **Microsoft Word** dialog box, click **Yes**.
 19. In the **Windows Security** dialog box, enter the following credentials, select **Remember my credentials**, and then click **OK**:
 - Username: **April**
 - Password: **Pa\$\$w0rd**
 20. In the **Select User** dialog box, ensure that **april@treymresearch.com** is selected, and then click **OK**.
 21. In the **Microsoft Office** dialog box, click **OK**.
 22. In the **Welcome to Microsoft Office 2010** dialog box, click **Don't make changes**, and then click **OK**.
 23. When the document opens, verify that you are unable to modify or save the document.
 24. Select a line of text in the document and verify.
 25. Right-click the text, and verify that you cannot make changes.
 26. Click **View Permission**, review the permissions, and then click **OK**.

Results: After completing this exercise, you should have verified that the AD RMS deployment is successful.

► To prepare for the next module

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

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

MCT USE ONLY. STUDENT USE PROHIBITED

Module 12: Implementing Active Directory Federation Services

Lab: Implementing AD FS

Exercise 1: Configuring AD FS Prerequisites

► Task 1: Configure DNS forwarders

1. On LON-DC1, in Server Manager, click **Tools**, and then click **DNS**.
2. Expand **LON-DC1**, and then click **Conditional Forwarders**.
3. Right-click **Conditional Forwarders**, and then click **New Conditional Forwarder**.
4. In the **DNS Domain** dialog box, type **TreyResearch.net**.
5. Click in the **IP address** column, and type **172.16.10.10**. Press Enter, and then click **OK**.
6. Close the DNS Manager.
7. On MUN-DC1, in Server Manager, click **Tools**, and then click **DNS**.
8. Expand **MUN-DC1**, and then click **Conditional Forwarders**.
9. Right-click **Conditional Forwarders**, and then click **New Conditional Forwarder**.
10. In the **DNS Domain** box, type **Adatum.com**.
11. Click in the **IP address** column, and type **172.16.0.10**. Press Enter, and then click **OK**.
12. Close the DNS Manager.

► Task 2: Exchange root certificates to enable certificate trusts

1. On LON-DC1, access the **Search** page.
2. In the **Search** box, type `\\MUN-DC1.treyresearch.net\certenroll`, and then press Enter.
3. In the CertEnroll window, right-click the **MUN-DC1.TreyResearch.net_TreyResearchCA.crt** file, and then click **Copy**.
4. In the left pane, click **Documents**, and then paste the file into the **Documents** folder.
5. Open a Windows PowerShell® command prompt, type **MMC**, and then press Enter.
6. In the Console1 window, click **File**, and then click **Add/Remove Snap-in**.
7. Click **Group Policy Management Editor**, and then click **Add**.
8. In **Group Policy Object**, click **Browse**.
9. Click **Default Domain Policy**, and then click **OK**.
10. Click **Finish**, and then click **OK**.
11. Double-click **Default Domain Policy**. In the console tree, expand **Computer Configuration > Policies > Windows Settings > Security Settings > Public Key Policies > Trusted Root Certification Authorities**.
12. Right-click **Trusted Root Certification Authorities**, and then click **Import**.
13. On the **Welcome to the Certificate Import Wizard** page, click **Next**.
14. On the **File to Import** page, click **Browse**.

15. In the Open window, click **MUN-DC1.TreyResearch.net_TreyResearchCA.crt**, click **Open**, and then click **Next**.
16. On the **Certificate Store** page, verify that **Place all certificates in the following store** is selected, verify that the **Trusted Root Certification Authorities** store is listed, and then click **Next**.
17. On the **Completing the Certificate Import Wizard** page, click **Finish**, and then click **OK**.
18. Close the Group Policy Management Editor without saving changes.
19. On MUN-DC1, access the **Search** page.
20. In the **Search** box, type `\\LON-DC1.adatum.com\certenroll`, and press Enter.
21. In the CertEnroll window, right-click the **LON-DC1.Adatum.com_Adatum-LON-DC1-CA.crt** file, and then click **Copy**.
22. In the left pane, click **Documents**, and paste the file into the **Documents** folder.
23. Open a Windows PowerShell command prompt, type **MMC**, and then press Enter.
24. In the Console1 window, click **File**, and then click **Add/Remove Snap-in**.
25. Click **Certificates**, and then click **Add**.
26. Click **Computer Account**, and click **Next**.
27. Verify that **Local computer** is selected, click **Finish**, and then click **OK**.
28. Expand **Certificates**, and then click **Trusted Root Certification Authorities**.
29. Right-click **Trusted Root Certification Authorities**, point to **All Tasks**, and then click **Import**.
30. On the **Welcome to the Certificate Import Wizard** page, click **Next**.
31. On the **File to Import** page, click **Browse**.
32. In the Open window, click **LON-DC1.Adatum.com_Adatum-LON-DC1-CA.crt**, click **Open**, and then click **Next**.
33. On the **Certificate Store** page, verify that **Place all certificates in the following store** is selected, verify that the **Trusted Root Certification Authorities** store is listed, and then click **Next**.
34. On the **Completing the Certificate Import Wizard** page, click **Finish**, and then click **OK**.
35. Close Console1 without saving changes.

► **Task 3: Request and install a certificate for the web server**

1. On LON-SVR1, in Server Manager, click **Tools**, and then click **Internet Information Services (IIS) Manager**.
2. In the console tree, click **LON-SVR1 (Adatum\Administrator)**. Click **No** to dismiss the message that displays.
3. In middle pane, double-click **Server Certificates**.
4. In the Actions pane, click **Create Domain Certificate**.
5. On the **Distinguished Name Properties** page, enter the settings as listed below, and then click **Next**.
 - Common name: **LON-SVR1.adatum.com**
 - Organization: **A. Datum**
 - Organization unit: **IT**

- City/locality: **London**
 - State/province: **England**
 - Country/region: **GB**
6. On the **Online Certification Authority** page, in **Specify Online Certification Authority**, click **Select** to search for a certification authority (CA) server in the domain.
 7. Select **Adatum-LON-DC1-CA**, and then click **OK**.
 8. In **Friendly name**, type **LON-SVR1.adatum.com**, and then click **Finish**.

► **Task 4: Bind the certificate to the claims-aware application on the web server, and verify application access**

1. On LON-SVR1, in Internet Information Services (IIS) Manager, expand **Sites**, click **Default Web Site**, and then in the Actions pane, click **Bindings**.
2. In the **Site Bindings** dialog box, click **Add**.
3. In the **Add Site Binding** dialog box, under **Type**, select **https**, and under **Port**, verify that **443** is selected. In the **SSL Certificate** drop-down list, click **LON-SVR1.adatum.com**, and then click **OK**.
4. Click **Close**, and then close **Internet Information Services (IIS) Manager**.
5. On LON-DC1, open Windows® Internet Explorer®.
6. Connect to **https://lon-svr1.adatum.com/adatumtestapp**.
7. Verify that you can connect to the site, but that you receive a 401 access denied error. This is expected because you have not yet configured AD FS for authentication.
8. Close Internet Explorer.

Results: In this exercise, you configured DNS forwarding to enable name resolution between A. Datum and Trey Research, and you exchanged root certificates between the two organizations. You also installed and configured a web certificate on the application server.

Exercise 2: Installing and Configuring AD FS

► **Task 1: Install and configure AD FS**

1. On the LON-DC1, in Server Manager, click **Manage**, and then click **Add Roles and Features**.
2. On the **Before you begin** page, click **Next**.
3. On the **Select installation type** page, click **Next**.
4. On the **Select destination server** page, click **Next**.
5. On the **Select server roles** page, select the **Active Directory Federation Services** check box, click **Add Features**, and then click **Next**.
6. On the **Select features** page, click **Next**.
7. On the **Active Directory Federation Services (AD FS)** page, click **Next**.
8. On the **Select role services** page, click **Next**.
9. On the **Confirm installation selections** page, click **Install**, and wait for the installation to finish. Do not close the window.

► Task 2: Create a standalone federation server using the AD FS Federation Server Configuration Wizard

1. On the **Installation progress** page, click **Run the AD FS Management snap-in**.
2. In the Overview pane, click the **AD FS Federation Server Configuration Wizard** link.
3. On the **Welcome** page, ensure that **Create a new Federation Service** is selected, and then click **Next**.
4. On the **Select Stand-Alone or Farm Deployment** page, click **Stand-alone federation server**, and then click **Next**.
5. On the **Specify the Federation Service Name** page, ensure that the **SSL certificate** selected is **LON-DC1.Adatum.com**, the **Port** is **443**, and the **Federation Service name** is **LON-DC1.Adatum.com**, and then click **Next**.
6. On the **Ready to Apply Settings** page, verify that the correct configuration settings are listed, and then click **Next**.
7. Wait for the configuration to finish, and then click **Close**.

► Task 3: Verify that FederationMetaData.xml is present and contains valid data

1. Log on to the **LON-CL1** virtual machine as **Adatum\Brad** using the password **Pa\$\$w0rd**.
2. Click the Desktop tile, and then open Internet Explorer.
3. Click the **Tools** icon at the top right corner, and then click **Internet options**.
4. On the **Security** tab, click **Local intranet**.
5. Click **Sites**, and clear the **Automatically detect intranet network** check box.
6. Click **Advanced**, and in the **Add this website to the zone** box, type **https://lon-dc1.adatum.com**, and then click **Add**.
7. Type **https://lon-svr1.adatum.com**, click **Add**, and then click **Close**.
8. Click **OK** twice.
9. Connect to **https://lon-dc1.adatum.com/federationmetadata/2007-06/federationmetadata.xml**.
10. Verify that the xml file opens successfully, and scroll through its contents.
11. Close Internet Explorer.

Results: In this exercise, you installed and configured the AD FS server role, and verified a successful installation by viewing the Federation Meta Data .xml contents.

Exercise 3: Configuring AD FS for a Single Organization

► Task 1: Configure a Token-signing certificate for LON-DC1.Adatum.com

1. On the LON-DC1 virtual machine, in Server Manager, click **Tools**, and then click **Windows PowerShell**.
2. At the prompt, type **set-ADFSProperties -AutoCertificateRollover \$False**, and then press Enter. This step is required so that you can modify the certificates that AD FS uses.
3. Close the Windows PowerShell window.

4. In Server Manager, click **Tools**, and then click **AD FS Management**.
5. In the AD FS console, in the left pane, expand **Service**, and then click **Certificates**.
6. Right-click **Certificates**, and then click **Add Token-Signing Certificate**.
7. In the **Select a token signing certificate** dialog box, click the first certificate with the name **LON-DC1.Adatum.com**, and then click **Click here to view certificate properties**.
8. Verify that the certificate purposes includes **Proves your identity to a remote computer** and **Ensures the identity of a remote computer**, and click **OK**. The certificate may also have other purposes, but these two are required. If the certificate does not have the intended purposes, view the properties of the other certificates until you find one with the intended purposes. Click **OK**.
9. When the AD FS Management warning dialog box displays, click **OK**.



Note: Verify that the certificate has a subject of **CN=LON-DC1.Adatum.com**. If no name displays under the Subject when you add the certificate, delete the certificate, and add the next certificate in the list.

10. Right-click the newly-added certificate, and then click **Set as Primary**. Review the warning message, and then click **Yes**.
11. Select the certificate that has just been superseded, right-click the certificate, and then click **Delete**. Click **Yes** to confirm the deletion.

► **Task 2: Configure the Active Directory claims provider trust**

1. On LON-DC1, in the AD FS console, expand **Trust Relationships**, and then click **Claims Provider Trusts**.
2. In the middle pane, right-click **Active Directory**, and then click **Edit Claim Rules**.
3. In the Edit Claims Rules for Active Directory window, on the **Acceptance Transform Rules** tab, click **Add Rule**.
4. In the Add Transform Claim Rule Wizard, in the **Select Rule Template** page, under **Claim rule template**, select **Send LDAP Attributes as Claims**, and then click **Next**.
5. On the **Configure Rule** page, in the **Claim rule name** box, type **Outbound LDAP Attributes Rule**.
6. In the **Attribute Store** drop-down list, select **Active Directory**.
7. In the **Mapping of LDAP attributes to outgoing claim types** section, select the following values for the LDAP Attribute and the Outgoing Claim Type:
 - E-Mail-Addresses = **E-Mail Address**
 - User-Principal-Name = **UPN**
 - Display-Name = **Name**
8. Click **Finish**, and then click **OK**.

► **Task 3: Configure the claims application to trust incoming claims by running the Windows Identity Foundation Federation Utility**

1. On LON-SVR1, click to the Start screen, and then click **Windows Identity Foundation Federation Utility**.

2. On the **Welcome to the Federation Utility wizard** page, in **Application configuration location**, type **C:\inetpub\wwwroot\AdatumTestApp\web.config** for the location of the web.config file of the Windows Identity Foundation sample application.
3. In **Application URI**, type **https://lon-svr1.adatum.com/AdatumTestApp/** to indicate the path to the sample application that will trust the incoming claims from the federation server. Click **Next** to continue.
4. On the **Security Token Service** page, select **Use an existing STS**, type **https://lon-dc1.adatum.com/federationmetadata/2007-06/federationmetadata.xml** for the STS WS-Federation metadata document location, and then click **Next** to continue.
5. On the **Security token encryption** page, select **No encryption**, and then click **Next**.
6. On the **Offered claims** page, review the claims that will be offered by the federation server, and then click **Next**.
7. On the **Summary** page, review the changes that will be made to the sample application by the Federation Utility wizard, scroll through the items to understand what each item is doing, and then click **Finish**.
8. Click **OK**.

► **Task 4: Configure a relying party trust for the claims-aware application**

1. On LON-DC1, in the AD FS Management console, click **AD FS**.
2. In the middle pane, click **Required: Add a trusted relying party**.
3. In the Add Relying Party Trust Wizard, on the **Welcome** page, click **Start**.
4. On the **Select Data Source** page, select **Import data about the relying party published online or on a local network**, and then type **https://lon-svr1.adatum.com/adatumtestapp**.
5. Click **Next** to continue. This action prompts the wizard to check for the Metadata of the application that the web server role hosts.
6. On the **Specify Display Name** page, in the **Display name** box, type **ADatum Test App**, and then click **Next**.
7. On the **Choose Issuance Authorization Rules** page, ensure that the **Permit all users to access this relying party** is selected, and then click **Next**.
8. On the **Ready to Add Trust** page, review the relying party trust settings, and then click **Next**.
9. On the **Finish** page, click **Close**. The Edit Claim Rules for ADatum Test App window opens.

► **Task 5: Configure claim rules for the relying party trust**

1. In the **Edit Claim Rules for Adatum Test App** properties dialog box, on the **Issuance Transform Rules** tab, click **Add Rule**.
2. In the Add Transform Claim Rule Wizard, on the **Select Rule Template** page, under **Claim rule template**, select **Pass Through or Filter an Incoming Claim**, and then click **Next**. This action passes an incoming claim through to the user by means of Integrated Windows authentication.
3. On the **Configure Rule** page, in **Claim rule name**, type **Pass through Windows Account name rule**. In the **Incoming claim type** drop-down list, select **Windows account name**, and then click **Finish**.
4. Click **Add Rule**.
5. On the **Select Rule Template** page, under **Claim rule template**, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.

6. On the **Configure Rule** page, in **Claim rule name**, type **Pass through E-mail Address rule**, in the **Incoming claim type** drop-down list, select **E-mail Address**, and then click **Finish**.
7. Click **Add Rule**.
8. On the **Select Rule Template** page, under **Claim rule template**, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
9. On the **Configure Rule** page, in **Claim rule name**, type **Pass through UPN rule**, in the **Incoming claim type** drop-down list, select **UPN**, and then click **Finish**.
10. Click **Add Rule**.
11. On the **Select Rule Template** page, under **Claim rule template**, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
12. On the **Configure Rule** page, in **Claim rule name**, type **Pass through Name rule**, in the **Incoming claim type** drop-down list, select **Name**, and then click **Finish**.
13. Click **Apply**, and then click **OK**.

► **Task 6: Test access to the claims-aware application**

1. On LON-CL1, open Internet Explorer.
2. Connect to **https://lon-svr1.adatum.com/AdatumTestApp/**



Note: Ensure that you type the trailing forward slash (/).

3. If you are prompted for credentials, type **Adatum\Brad** with password **Pa\$\$w0rd**, and then press Enter. The page renders, and you see the claims that were processed to allow access to the web site.

Results: In this exercise, you configured a Token signing certificate and configured a claims provider trust for Adatum.com. You also should have configured the sample application to trust incoming claims, and configured a relying party trust and associated claim rules. You also tested access to the sample Windows Identity Foundation application in a single organization scenario.

Exercise 4: Configuring AD FS for Federated Business Partners

► **Task 1: Add a claims provider trust for the TreyResearch.net AD FS server**

1. On LON-DC1, if required, in Server Manager, click **Tools**, and then click **AD FS Management**.
2. In the AD FS console, expand **Trust Relationships**, and then click **Claims Provider Trusts**.
3. In the Actions pane, click **Add Claims Provider Trust**.
4. On the **Welcome** page, click **Start**.
5. On the **Select Data Source** page, select **Import data about the claims provider published online or on a local network**, type **https://mun-dc1.treyresearch.net**, and then click **Next**.
6. On the **Specify Display Name** page, click **Next**.
7. On the **Ready to Add Trust** page, review the claims provider trust settings, and then click **Next** to save the configuration.
8. On the **Finish** page, click **Close**.

9. In the **Edit Claim Rules for mun-dc1.treyresearch.net** properties dialog box, on the **Acceptance Transform Rules** tab, click **Add Rule**.
10. In the **Claim rule template** list, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
11. In the **Claim rule name** text box, type **Pass through Windows account name rule**.
12. In the **Incoming claim type** drop-down list, select **Windows account name**.
13. Select **Pass through all claim values**, and then click **Finish**. Click **Yes**.
14. Click **OK**, and then close the AD FS console.
15. On LON-DC1, in Server Manager, click **Tools**, and then click **Windows PowerShell**.
16. At the command prompt, type the following command, and then press Enter:

```
Set-ADFSClaimsProviderTrust -TargetName "mun-dc1.treyresearch.net" -  
SigningCertificateRevocationCheck None
```


17. Close the Windows PowerShell window.

► **Task 2: Configure a relying party trust on MUN-DC1 for the A. Datum claims-aware application**


1. On the MUN-DC1, in Server Manager, click **Tools**, and then click **AD FS Management**.
2. In the AD FS console, on the **Overview** page, click **Required: Add a trusted relying party**.
3. On the **Welcome** page, click **Start**.
4. On the **Select Data Source** page, select **Import data about the relying party published online or on a local network**, type **https://lon-dc1.adatum.com**, and then click **Next**.
5. On the **Specify Display Name** page, in the **Display name** text box, type **Adatum TestApp**, and then click **Next**.
6. On the **Choose Issuance Authorization Rules** page, select **Permit all users to access this relying party**, and then click **Next**.
7. On the **Ready to Add Trust** page, review the relying party trust settings, and then click **Next** to save the configuration.
8. On the **Finish** page, click **Close**. The Edit Claim Rules for Adatum TestApp window opens.
9. In the Edit Claim Rules for Adatum TestApp window, on the **Issuance Transform Rules** tab, click **Add Rule**.
10. In the **Claim rule template** list, select **Pass Through or Filter an Incoming claim**, and then click **Next**.
11. In the **Claim rule name** box, type **Pass through Windows account name rule**, in the **Incoming Claim type** drop-down list, select **Windows account name**.
12. Select **Pass through all claim values**, and then click **Finish**.
13. Click **OK**, and then close the AD FS console.

► **Task 3: Verify access to the A. Datum test application for Trey Research users**

1. On MUN-DC1, open Internet Explorer, and connect to **https://lon-svr1.adatum.com/adatumtestapp/**.

 **Note:** The logon process has changed, and you must now select an authority that can authorize and validate the access request. The Home Realm Discovery page (the Sign In page) appears and you must select an authority.

2. On the **Sign In** page, select **mun-dc1.treyresearch.net**, and then click **Continue to Sign in**.
3. When prompted for credentials, type **TreyResearch\April** with the password **Pa\$\$w0rd**, and then press Enter. You should be able to access the application.
4. Close Internet Explorer.
5. Open Internet Explorer, and connect to **https://lon-svr1.adatum.com/adatumtestapp/** again.
6. When prompted for credentials, type **TreyResearch\April** with password **Pa\$\$w0rd**, and then press Enter. You should be able to access the application.
7. Close Internet Explorer.

 **Note:** You are not prompted for a home realm again. Once users have selected a home realm and been authenticated by a realm authority, they are issued an **_LSRealm** cookie by the relying party federation server. The default lifetime for the cookie is 30 days. Therefore, to log on multiple times, you should delete that cookie after each logon attempt to return to a clean state.

► **Task 4: Configure claim rules for the claim provider trust and the relying party trust to allow access only for a specific group**

1. On MUN-DC1, open the AD FS console, expand **Trust Relationships**, and then click **Relying Party Trusts**.
2. Select **Adatum TestApp**, and in the Actions pane, click **Edit Claim Rules**.
3. On the Edit Claim Rules for Adatum TestApp window, on the **Issuance Transform Rules** tab, click **Add Rule**.
4. On the **Select Rule Template** page, under **Claim rule template**, select **Send Group Membership as a Claim**, and then click **Next**.
5. On the **Configure Rule** page, in the **Claim rule name** field, type **Permit Production Group Rule**.
6. Next to **User's Group**, click **Browse**, type **Production**, and then click **OK**.
7. Under **Outgoing claim type**, click **Group**.
8. Under **Outgoing claim value**, type **Production**, and then click **Finish**. Click **OK**.
9. On LON-DC1, if required, open the AD FS Management console.
10. In the AD FS console, expand **Trust Relationships**, and then click **Claim Provider Trusts**.
11. Select **mun-dc1.treyresearch.net**, and in the Actions pane, click **Edit Claim Rules**.
12. On the **Acceptance Transform Rules** tab, click **Add Rule**.
13. On the **Select Rule Template** page, under **Claim rule template**, select **Pass Through or Filter an Incoming Claim**, and then click **Next**.
14. On the **Configure Rule** page, in **Claim rule name**, type **Send Production Group Rule**.
15. In the **Incoming claim type** drop-down list, click **Group**, and then click **Finish**. Click **Yes**, and then click **OK**.
16. In the AD FS console, under **Trust Relationships**, click **Relying Party Trusts**.

17. Select the **Adatum Test App**, and in the Actions pane, click **Edit Claim Rules**.
18. On the **Issuance Transform Rules** tab, click **Add Rule**.
19. Under **Claim rule template**, click **Pass Through or Filter an Incoming Claim**, and then click **Next**.
20. Under **Claim rule name**, type **Send TreyResearch Group Name Rule**.
21. In the **Incoming claim type** drop-down list, click **Group**, and then click **Finish**.
22. On the **Edit Claim Rules for Adatum Test App** window, on the **Issuance Authorization Rules** tab, select the rule named **Permit Access to All Users**, and then click **Remove Rule**. Click **Yes** to confirm. With no rules, no users are permitted access.
23. On the **Issuance Authorization Rules** tab, click **Add Rule**.
24. On the **Select Rule Template** page, under **Claim rule template**, select **Permit or Deny Users Based on an Incoming Claim**, and then click **Next**.
25. On the **Configure Rule** page, in **Claim rule name**, type **Permit TreyResearch Production Group Rule**, in the **Incoming claim type** drop-down list, select **Group**, in **Incoming claim value**, type **Production**, select the option to **Permit access to users with this incoming claim**, and then click **Finish**.
26. On the **Issuance Authorization Rules** tab, click **Add Rule**.
27. On the **Select Rule Template** page, under **Claim rule template**, select **Permit or Deny Users Based on an Incoming Claim**, and then click **Next**.
28. On the **Configure Rule** page, in the **Claim rule name** field, type **Temp**, in the **Incoming claim type** drop-down list, select **UPN**, in the **Incoming claim value** field, type **@adatum.com**, select the option to **Permit access to users with this incoming claim**, and then click **Finish**.
29. Click the **Temp** rule, and click **Edit Rule**.
30. In the **Edit Rule –Temp** dialog box, click **View Rule Language**.
31. Press **Ctrl + C** to copy the rule language to the clipboard, and then click **OK**.
32. Click **Cancel**.
33. Click the **Temp** rule, click **Remove Rule**, and then click **Yes**.
34. On the **Issuance Authorization Rules** tab, click **Add Rule**.
35. On the **Select Rule Template** page, under **Claim rule template**, select **Send Claims Using a Custom Rule**, and then click **Next**.
36. On the **Configure Rule** page, in the **Claim rule name** field, type **ADatum User Access Rule**.
37. Click in the **Custom rule** box, and then press **Ctrl + V** to paste the clipboard contents into the box. Edit the first URL to match the following text, and then click **Finish**.

```
c:[Type == "http://schemas.xmlsoap.org/ws/2005/05/identity/claims/upn", Value =~
"^(?i).+@adatum\.com$"]=> issue(Type =
"http://schemas.microsoft.com/authorization/claims/permit", Value = "PermitUsersWithClaim");
```



Note: This rule enables access to anyone who presents a claim that includes the User Principal Name (UPN) of @adatum.com. The Value line in the first URL defines the attribute that must be matched in the claim. In this line, ^ indicates the beginning of the string to match, (?i) means that the text is case insensitive, .+ means that one or more characters will be added, and \$ means the end of the string.

38. Click **OK** to close the property page, and save the changes to the relying party trust.

► **Task 5: Verify restrictions and accessibility to the claims-aware application**

1. On MUN-DC1, open Internet Explorer, and connect to **https://lon-svr1.adatum.com/adatumtestapp/**.
2. When prompted for credentials, type **TreyResearch\April** with password **Pa\$\$w0rd**, and then press Enter. April is not a member of the Production group, so she should not be able to access the application.
3. Close Internet Explorer.
4. Re-open Internet Explorer, click the **Tools** icon in the top right corner, and then click **Internet options**.
5. Under **Browsing history**, click **Delete**, click **Delete** again, and then click **OK**.
6. Connect to **https://lon-svr1.adatum.com/adatumtestapp/**.
7. On the **Sign In** page, click **mun-dc1.treyresearch.net** and then click **Continue to Sign in**.
8. When prompted for credentials, type **TreyResearch\Morgan** with password **Pa\$\$w0rd**, and then press ENTER. Morgan is a member of the Production group, and should be able to access the application.
9. Close Internet Explorer.

Results: In this exercise, you configured a claims provider trust for TreyResearch on Adatum.com. and a relying party trust for Adatum on TreyResearch. You verified access to the A. Datum claim-aware application. Then you configured the application to restrict access from TreyResearch to specific groups, and you verified appropriate access.

► **To shut down the virtual machines**

When you finish the lab, revert the virtual machines to their initial state.

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

MCT USE ONLY. STUDENT USE PROHIBITED